MARK ORTIZ CHASSIS NEWSLETTER

ÍNDICE DE LOS ARTÍCULOS PUBLICADOS

2000.5 – Rear sway bars; Effects of wheel and axle offsets; Making handling more consistent; Relationship of gear ratio and RPM
2000.6 – Effects of rear spring split; Using brake floaters
2000.7 – Brake failure in pavement Late Models; Polar moment of inertia (yaw inertia)
2000.8 – Shock dynamics – What the shock dyno does and doesn't tell you; Effects of gas pressure; Control ratio; Natural frequency, damping strength, and roadholding
2000.9 – How to string a car for 4-wheel alignment; Effects of caster
2000.10 – Checking rear axles for straightness; Torque arms versus pull bars
2000.11 – Suggested off-season reading
2000.12 – Springs, roll, and cornering balance; Shorty Panhard bars versus long ones
2001.1 – Wind tunnel testing for short track cars; Rear caster; Stacked coilover springs
2001.2 – Required frame stiffness; Making ballast weights
2001.3 – Safety issues – The HANS device; Soft wall design requirements
2001.4 – 5th coil location and rate; Soft wall update; Soft noses on cars
2001.5 – Chrome moly in stock cars; Rear stagger versus cross

2001.6 – Weight transfer in winged-over sprint cars, with discussion of the various components of weight transfer (wheel load transfer)
2001.7 – Chassis troubleshooting chart, with discussion of the 5 phases of the cornering process and assumptions on which the troubleshooting chart is based
2001.8 – Things that make spring changes work backwards

2001.9 – Stock car safety issues, including discussion of NASCAR's Earnhardt Report; Tire data – its availability and its limitations
2001.10 – Balancing the car with camber and tire pressure
2001.11 – The big tracks and their big wrecks; Factors in spring selection
2002.1 – Suspension natural frequencies
2002.2 – Winston Cup harness installation; Things that make spring changes work backwards, revisited; Basic shock questions
2002.3 – Update/correction on Winston Cup harness installation; Racing front-drive cars
2002.4 – Quads on pavement
2002.5 – Suspension frequencies revisited; Quads on pavement revisited; Space frame or monocoque?; Tire care
2002.6 – Oil accidentally or deliberately applied to tires; Diagonal percentage and tire stagger for road racers; Narrower rear track than front
2002.7 – Braking/downshifting technique
2002.8 – Steering geometry variables; Importance of steering rack placement; Bump steer as Ackermann modification?; Relationship o horsepower and engine torque to acceleration and axle torque
2002.9 – Additional remarks on bump steer and horsepower/torque; Finding centers of gravity for whole car, unsprung masses, and sprung mass
2002.10 – Effects of wheelbase, CG height, and static rear percentage on traction and corner exit behavior
2002.11 – Push rods, pull rods, or both; More on trail and steering feel; Kart scaling; Another tire care tip
2002.12 – Ackermann effect, with recommendations for autocross and hillclimbing; Torque and power distribution with various types of differentials
2002.12 – Analyzing shock traces

2003.1 – Shocks wanted for research; Tire warmers found; Spring placement on triangulated 4-link rear axle; Roll center with a j-bar
2003.2 – Shock research update; More on rear wheel placement and traction, including discussion of gyroscopic precession
2003.3 – Tire warmers less expensive than I though; Independent rear suspension for dirt racing?
2003.4 – Special issue on shock and spring forces and their effects on tire loading
2003.7 – Right rear spring stiffest, on a Cup car? Pushrod angle, motion ratio, wheel rates, and roll angle
2003.8 – Load transfer basics
2003.9 – Shock research update – end-of-stroke phenomena due to fluid compressibility, as opposed to acceleration sensitivity
2003.11 – Weight distribution and tire size; Car width versus drag
2003.12 – More on racing front-drive cars, and on load transfer; Single or dual rear brakes for FSAE?
2004.1 – Reverse Ackermann or toe-in on ovals; Setting up road racing cars that are heavier on one side than the other; Independent or beam axle front suspension for ovals?
2004.3 – Types of differentials and their effects on car behavior
2004.4 – Effect of brake caliper mounting position; More on weight distribution; Spring splits with big bars

2004.5 – Correction to April remarks on caliper mounting (archived April issue incorporates the correction); New Nextel Cup tire and implications for control arm geometry; control arm angles in F1; Arm angles in triangulated 4-bar beam axle rear suspension
2004.6 – More on F1 control arm angles, with discussion of "lateral anti" concept and shortcomings of using force line intersection as roll center; Use of clutch when downshifting
2004.7 – Setup for lower grip; Lowering blocks on truck arms; Ohlins high-frequency shock piston; Which front shock to soften for improved exit; soft links on sway bars
2004.8 – Further discussion of roll center assignment based on lateral anti, with discussion of lateral migration of the roll center as traditionally defined

2004.9 – Tire load sensitivity and weight transfer in trailbraking; Why are wide tires better?; Relative importance of front/rear weight distribution versus yaw inertia; Roll and ride spring preload in monoshock suspensions; Why so few rear monoshock suspensions?
2004.12 – Roll center below ground; Roll moments from longitudinal anti
2005.1 – Steering axis inclination and scrub radius
2005.3 – Finer points of anti-roll bars on stock cars
2005.5 – Can an oval track car have too much left percentage?

2005.7 – Dynamics of three-wheeled vehicles
2005.9 – Load transfer case with added rear ride compression
2005.11 – Optimizing engine-over-drive-wheels cars
2006.1 – Tires in the snow
2006.2 – Roll axis inclination
2006.3 – Worm gear vs. clutch pack differentials
2006.4 – Springs, bars, and load transfer
2006.5 – Rear percentage or polar moment?
2006.6 – Effects of spring rates and damper settings during cornering
2006.7 – Beam axle front suspension pros, cons, and design issues
2006.8 – Effects of Panhard bar fore-aft location, slope, and offset
2006.9 – Roll center in trailing arm front suspension; More on roll axis and related analytical topics
2006.10 – Twist beams and jointed De Dion tubes
2006.11 – Stiff rebound damping
2006.12 – Coil binding setups in stock cars – friend or foe?; Falling force shocks
2007.1 – Tall vs. short sidewalls
2007.2 – Modern driveshaft joints vs. rubber doughnuts; Zero-droop setups
2007.3 – Diamond layout
2007.4 – Roll steer and anti-squat/anti-lift in strut suspension
2007.5 – Roll axis for live rear axle road/hillclimb car; Roll axis of the axle
2007.6 – Another triangulated four-link axle; Large vs. small tire and wheel diameter
2007.7 – More thoughts on zero-droop setups, and related oval track spring split issues; High or low panhard bar
2007.8 – Is it bad to lift a wheel?; Hollow vs. solid anti-roll bars
2007.9 – Droop-limited rear suspension; Satchell link rear suspension; A word from the inventor
2007.10 – What exactly are shaker rigs for?
2007.11 – On the march again; Radial tire camber; A radical radial notion
2007.12 – Ins and outs of toe and ackermann
2008.1 – More exotic vehicle layout ideas
2008.3 – Another exotic car layout
2008.4 – Late Corvair rear suspension
2008.5 – Another pony car suspension from the 60's; Really quick steering
2008.6 – Why not go for 100% camber recovery?; Gyros for roll stability
2008.7 – Front to rear track ratio
2008.8 – Dirty tricks; Porsche 993 rear suspension
2008.9 – Monoshock suspensions; Z-bars; Stock car front ends with bump rubbers
2008.10 – Differentials and yaw control
2008.12 – Controlling AWD front/rear torque split on a budget; Is there such a thing as too big (in tires, that is)?
2009.1 – Damping ratio
2009.3 – Trailing arms with torsional rubber springing; Ride and roll rates with monoshocks
2009.4 – "Minding Your Anti" video now on DVD; Raised lower control arms on 2009 F1 cars; Pullrod suspensions in F1; Rate and length numbers on torsion bars
2009.5 – The anti controversy
2009.7 – The unsprung component in load transfer
2009.8 – Chasing a push
2009.9 – Leading and trailing wheels for oval track
2010.1 – Use ride height to tune cornering balance, or anti-roll bars?; Roll center and anti-lift in triangulated 4-link beam axle; Measuring grip
2010.2 – Effect of changing location of just the wheel center; Effects of droop limiting
2010.3 – Old-school rocker suspension; De Dion tube design and location; Lateral location of live axle
2010.4 – Banked turn puzzle; Reverse-cant rear leaf springs
2010.5 – Selecting rear end ratio; Torque tube for front engine IRS car
2010.6 – Torsion bar tube height; Twin I beam suspension, Relationship of tire size and weight distribution
2010.7 – Tuning transient behavior of front-wheel-drive autocross car
2010.8 – When lots of front roll stiffness helps, and when it doesn't
2010.9 – When one is a very small animal: Considering an RC car
2010.10 – Roll center theory
2010.11 – High and low speed balance, and torque arm length
2010.12 – Hybrid or semi-independent suspensions
2011.1 – Coil bind and bump stop setups in stock cars
2011.2 – Sway bar/spring rate equivalency; What wheel rate for damper calculations?
2011.3 – Followup information to January issue on stock car anti-roll bar preload and related asymmetries; Followup question to the February issue on anti-roll bars
2011.4 – Negative camber only on the front wheels, or more on the fronts than on the rears, on some cars; Roll center/front view geometry with dual ball joints
2011.5 – Why use sway bars at all?
2011.6 – Effects of caster and scrub radius; Camber gain recommendations
2011.7 – Mumford Linkage; Toe out for turn in
2011.8 – Toe, Ackermann, cornering force, turn-in and related puzzles
2011.9 – Anti-roll bars on Formula Vees revisited; General advice on setting dampers
2011.10 – The DeltaWing car
2011.11 – First-generation Mazda RX7 four-link; Unified theory of suspension setup for particular track?
2011.12 – More/less anti-squat – or what? – to fix wheelspin/rear tire wear
2012.1 – Equal wheel rates left/right, or equal frequencies?; Two-wheeler that doesn't lean, but shifts weight laterally?
2012.2 – Anti-dive, and the Lotus reactive ride height system
2012.3 – More on reactive anti-dive; Forces on control arms when cornering

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2000

WELCOME

This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Selected questions will be presented, at my discretion. Readers are invited to submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net

Mark Ortiz

We are nearing completion on a street stock 79 Camaro to run on ½ mile asphalt track with 10 degree banking in looong corners. We are considering trying a rear sway bar. We're being told to toss it, but I can't help thinking there's something there if we're willing to work on it. Any thoughts on what we might expect to encounter if we try running it?

Go ahead and try it.

If it is simply added with no other changes, a rear anti-roll bar will loosen the car. Used with softer springs to get similar overall roll resistance, it will result in a softer wheel rate in pitch and heave.

Be aware that your setup will require different springs than other people's. Of course, if you're running a rear anti-roll bar, you've moved beyond tuning by imitation anyway.

On a street stock, which usually cannot run jacking screws, anti-roll bars with drop links offer a way to adjust wedge and tilt (by changing link length), provided that bar preloading is not prohibited. Having bars at both ends expands the scope of such adjustments. You also have the ability to quickly loosen or tighten the car by disconnecting a bar.

All the tech books and articles I've seen tell you to kick out the left rear tire and tuck in the right rear tire to tighten a dirt car and give it better forward bite. This is commonly done for a dry slick track. I understand why the right rear bites better but how can moving the left rear out from the car give you better forward bite? You are taking weight off of it when you do this.

The reason rear wheel lateral positioning works as it does relates to the line of action of the forward thrust of the tire, rather than the load on the tire.

To illustrate, imagine your car was nose-heavy enough so you could take off one rear wheel and drive it as a tricycle. If you had only the left rear wheel, the car would try to turn right under power because all the thrust would be acting left of the center of mass. If you ran only a right rear instead, it would try to turn left under power. (Motorcycle sidecar rigs really display this effect.)

With two rear wheels, moving either or both of them left or right has a similar effect, only more subdued .

Moving both rear wheels to the left, then, makes the car try to turn right more (or left less) under power. Exiting a left turn, this tightens the car. Consequently, the rear tires don't have to use as much of their grip for cornering and have more available for propulsion.

Toeing both rear wheels leftward (leading the right rear, with a beam axle) has a similar effect, and doesn't cost you left percentage.

How do you make a car less sensitive to track variations? On a dirt track modified, I am constantly having to adjust for track changes and many times am one notch behind the track. Is this the sport in it or is there a technical solution?

Here are two tools to help control variation in a car's balance as grip varies:

1) Static diagonal percentage / roll resistance distribution. If the car is close to right in average conditions but goes loose on slick and tight on tacky, add diagonal and use more rear roll resistance and/or less front. Idea here is to increase dynamic (running) diagonal at moderate lateral acceleration (low grip) yet decrease it at high lateral acceleration (high grip). For a car that goes tight on slick, reverse this strategy.

2) Tire stagger. Tire stagger has more effect as grip increases. Therefore, if the car has little or no rear stagger – meaning stagger effect is tightening the car – the car will feel this effect more when grip is good, and act tighter. As grip diminishes, the locked axle push effect decreases and the car goes toward loose. If the car runs generous stagger, you see an opposite effect. Of course, changes in stagger require compensating changes elsewhere in the setup, and in some cases the tire sizes you'd like aren't obtainable – but at any rate that's how it works.

If I change from 6.26 final drive to a 6.56 final drive, how much RPM gain can I expect, and will the lap times be faster or slower? I have good bite, but am getting beat out of the corners. Currently turning 6200 – 406 engine, 87" rear tires, 1/3 to 3/8 mile dirt track.

You'll be around 6500 at equal speed. If you're above that, it's because you're going faster! Assuming that 406 can survive, you should gain some speed.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2000

WELCOME

This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Selected questions will be presented, at my discretion. Readers are invited to submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net

Mark Ortiz

We race on 1/3 mile high-banked and ½ mile flat asphalt tracks. Car is Chevy metric with rear weight jacks, 3200# minimum, stock lower and upper A-frames, stock 4-link rear end with Ford 9" axle, coil springs all around.

Without really getting into specific spring rates, my question has to do with the relative stiffness of the rear springs. Isn't it a generally accepted principle that the LR spring is stiffer than the RR spring, by 25 pounds or so?

I thought that on dirt the opposite was true, but on asphalt the LR is the stiffer spring.

Comments? What characteristics would a stiffer RR produce?

> Actually, stiffer LR is more common on both dirt and pavement.
>
> Let's assume we're comparing setups using the same pair of springs, and just swapping them side to side.
>
> On a truly flat track (no banking at all), mid-turn behavior should be about the same either way. The steeper the banking, the more the right-stiff split will loosen the car mid-turn.
>
> In general, the right-stiff split will loosen exit on any track. Steeper banking will intensify the effect. This assumes that the rear suspension compresses under power. Rear anti-squat will reduce the effect. If anti-squat is great enough so the rear end rises under power, the effect reverses.
>
> The right-stiff split will tighten entry on a flat track, assuming that the rear lifts in braking and that the car has a reasonable amount of front brake. Extreme amounts of anti-lift or rear brake can cause the effect to reverse. Steeper banking will diminish the normal effect and can reverse it in extreme cases.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

I have seen brake floaters being used in several ways. Some use a floater only on the right side. I've also seen them with the floater rods mounted on the top or mounted on the bottom of the right side. What would be the proper position of the rods, and should you have them on both left and right sides?

> There is really no right or wrong way to set up brake floaters – just predictable effects.
>
> Most often, brake floaters are set up to produce pro-lift. For links above the axle running forward, this means the forward end of the link is above the rear end. For a link below the axle, the link runs downhill toward the front or uphill toward the rear.
>
> Floaters can also be set up for anti-lift, which really makes more sense for shortening braking distances as it lowers rather than raises the center of gravity. Pro-lift has the advantage of allowing you to use lots of rear brake without wheel hop. In some cases, rear lift is used to promote rear steer effects.
>
> When a floater is used on the right only, it is usually intended to give more pro-lift on the right, de-wedging the car and loosening entry. (Note that de-wedging the car can tighten entry instead if the car has a lot of rear brake.) The left brake torque reacts through whatever means are provided to resist it at the axle housing.
>
> It is possible, especially with telescoping links on the axle, to get almost any desired properties using only one floater. The one thing you can't do is separately tune engine braking torque reaction and left brake torque reaction.
>
> I think two floaters are worth having, since the added expense is pretty small. But a car with one can be made to do most of the same things.

CORRECTION

Last month's newsletter contained a statement in the response to a question about rear wheel offset/lateral position that read: "If you had only one rear wheel, the car would try to turn right under power…" That should have been: "If you had only the left rear wheel, the car would try to turn right under power…"

My apologies for any confusion this may have created.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2000

WELCOME

This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Selected questions will be presented, at my discretion. Readers are invited to submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net

Mark Ortiz

BRAKE FAILURE IN PAVEMENT LATE MODELS

I've been getting both questions and comments from various people about brake failures in pavement Late Models. It appears that brakes on these cars go away a lot, especially the fronts. These failures occur when the car is running under green for long periods, and also following restarts after red flags. Very often, the car doesn't lose brakes entirely. It just loses the fronts, sometimes gradually, sometimes suddenly, and spins on entry.

Although pads and rotors can fail from extreme heat, it appears that the majority of problems in these cars are related to fluid boiling in the front calipers. This is certainly the case when the front brakes suddenly don't work following a red-flag heat soak.

Better ducting helps, especially under green-flag conditions. Juggling weight distribution and brake balance to make the rears do more of the work helps a little. Making sure the fluid is fresh is essential. But the clients who say they've really licked the problem are using both heat-shielded calipers and/or reduced-conductivity pistons, and brake fluid recirculating systems. In general, applying these measures on just the front brakes seems to be sufficient. This will of course depend somewhat on brake bias.

POLAR MOMENT OF INERTIA (YAW INERTIA)

Mark, one of my racing buddies used the term "polar moment of inertia" in a conversation we were having the other day. I have heard this expression before, but do not understand what it is. Could you explain what it is, and what effect it has on race cars – also, can you measure it somehow, and how does it relate to suspension design and/or coilover placement?

> The way racers use the term, it means polar moment of inertia in yaw. A car also has a polar moment of inertia in roll, and in pitch.
>
> Yaw is rotational, or angular, motion about a vertical axis (i.e. rotation as seen from above, or rotation that changes the direction the car points). To take a turn, we must accelerate the car in

yaw in the direction of the turn during entry, then decelerate it in yaw (accelerate it in the direction opposite the turn) during exit. We must start it rotating to make it turn, then stop the rotation to make it go straight again.

The car acts as a giant flywheel – its inertia opposes these accelerations. When it's running straight, it doesn't want to start rotating. Once it's rotating, it wants to keep rotating. This effect tightens entry and loosens exit.

The polar moment of inertia is the magnitude of this inertial effect. We increase it by moving masses away from the center of gravity. We decrease it by centralizing masses. A mid-engine car, like an Indy car, has a small polar moment of inertia in yaw. A stock car with the engine between the front wheels and 200 pounds of ballast, the battery, and the fuel load behind the rear axle has a large polar moment of inertia in yaw. So does a VW beetle, an Audi front-drive sedan, or a Porsche 911, with the engine outside the wheelbase.

Most people don't try to measure yaw inertia. GM built a giant turntable fixture to measure it. You can mathematically estimate it by breaking the car down into components, weighing these or calculating their mass, and multiplying the masses by the square of their distance from the CG. Most of us don't bother. For a pure racing car, we just try to put all the heavy stuff as close to the middle, or the expected CG, as possible.

For a production-based car, we often face the issue as a choice between placing components or ballast toward the rear bumper to get more rear percentage, or more centrally to reduce yaw inertia. In such cases I usually go for the rear percentage, especially for oval track applications. An exception would be where you can get more than enough rear percentage, and still fall short of legal minimum weight. Then it makes sense to centralize the ballast.

On an oval track car, we can use asymmetrical setups to make the car enter and exit as loose or tight as we want, even if the car has a lot of yaw inertia. Also, we don't encounter such abrupt changes of direction on an oval as we see in a chicane or sharp turn on a road course. Consequently, minimizing yaw inertia is more important in road racing than on oval tracks.

Both large and small polar moments of inertia are mixed blessings for any car. A car with a small polar moment and a short wheelbase will be twitchy (e.g. older Toyota MR2), unless it's set up very tight (e.g. Pontiac Fiero). When such a car encounters a slippery patch in mid-turn, it will do a big wiggle and possibly spin, whereas a car with more yaw inertia will be more stable.

So a car with a small polar moment should have a long wheelbase if possible. Suspension geometry requirements don't really change with yaw inertia. Moving coilovers toward the center of the car reduces yaw inertia, but not a lot since coilovers aren't very heavy.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2000

WELCOME

This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Selected questions will be presented, at my discretion. Readers are invited to submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net

Mark Ortiz

"THE SHOCK DYNO LIES!"

How come if I dyno a Bilstein shock (Roehrig dyno, using their software), and then duplicate the graph exactly on an Ohlins, the shock feels completely different to the driver? I mean different as in no side bite, and no forward bite – but the driver says the car doesn't feel any stiffer.

Also, if I vary the gas pressure, the dyno displays the same graph, but the driver can feel a difference in the shock. Why?

At what point in rebound/compression valving split will jacking occur? We valved an Ohlins to be exactly like our Bilsteins, and the Ohlins ratcheted down in the turn until it bottomed out and almost scared the driver to death.

How do I match the shocks to the track surface for maximum compliance and therefore maximum forward bite? The tires of the winning car in the 100 lap day show seemed to just roll over the washboard surface, and on all the other cars you could see the tires bouncing up and down.

The tracks get rough, and because the roughness is caused by the same type of vehicle, the roughness always has the same look. I would think that since the pattern of 3" to 4" holes and bumps is always similar, the "perfect" spring frequency and shock valving could be found. How do I test for this?

> The shock dyno doesn't lie, but it doesn't tell the whole truth.
>
> First of all, I believe your dyno cycles the shock through a 2" stroke, normally near mid-travel, at 100 cycles/min (1.67 Hz). This gives a peak velocity at mid-stroke of just over 10 in/sec. A rough dirt track may give you shaft speeds above that.
>
> Your dyno only generates simple harmonic motion. Other companies make more expensive dynos that can generate approximately square-wave (very high acceleration) motion, or can reproduce on-track motions recorded by electronic data acquisition. These modes of testing have value because shocks are sensitive to acceleration as well as velocity.
>
> Your software can be programmed to show you various outputs. Many people look only at one

end of the stroke, most often the extended end (rebound valving closing, compression valving opening). It helps to look at both ends. A plot showing both ends of the stroke, force versus absolute velocity, will have two points or noses at the left side, and four traces. I have seen instances where looking at the full stroke showed me acceleration sensitivities that I never would have known about if I had only looked at the extended end of the stroke.

The usual thing the program does with the gas force is to re-zero the load readout after stopping momentarily at mid-stroke and reading the gas force. However, the program should tell you, as a numerical readout, what that gas force is. It will vary with the pressure you put in at build, and also with the volume of the gas, which you control by varying the floating piston's position at build or the oil volume you pour in.

You don't always get the same trace with different gas pressures, even with re-zeroing. More gas pressure actually increases rebound damping force (with gas force omitted), due to reduced nucleate boiling (incipient cavitation) on the downstream side of the piston. This effect is greatest at high velocity, with a stiff rebound stack, low gas pressures, and hot oil. The effect is least – sometimes unnoticeable – at low velocity, with a soft rebound stack, high gas pressures, and cooler oil.

Valving split is sometimes expressed in terms of *control ratio* – the ratio between rebound and compression damping, at a particular shaft speed. As a rule of thumb, a control ratio of 1.3 to 2.5 is pretty normal; <1.3 is somewhat bump-stiff; >4.0 is likely to jack down. Jacking is also promoted by stiffer dampers, softer springs, or a bumpier track.

Tuning for a particular disturbance frequency is mainly a matter of making sure your natural frequencies *don't* match the excitation frequency. Since your unsprung masses and tires are similar to the other cars', this means using spring rates that don't match theirs, or using stiffer damping. Sometimes it helps to stiffen just rebound, but if you're jacking down to the bump stops you may be too far down that path now. Soft damping gives better roadholding, except when the bumps excite the system at one of its natural frequencies. Stiff damping raises natural frequencies, and also makes the system less frequency-sensitive.

You can also raise natural frequency by using somewhat stiffer springs than your competitors. In the days of cart-sprung cars with primitive dampers, this was a major reason people sprung race cars stiffly. Another approach is to run substantially softer instead. With everybody so soft nowadays, that can be difficult, but if you add a sway bar and good bump rubbers it can work.

Finding a good combination is mainly cut-and-try at the track. Electronic data acquisition can be a big help. I have an associate who specializes in that. His name is John Chapman. He's in Charlotte at 704-549-1309, e-mail jchap56756@aol.com. For shock dyno and build service, I recommend Scott Munksgard at Munksgard Technical Services in Concord, NC at 704-782-2611, e-mail MTSdyno@aol.com. He custom-builds Bilstein, Ohlins, Penske, and Carrera shocks, and sells AFCO and Pro shocks and Afcoil and Hypercoil springs.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
September 2000

WELCOME

This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Selected questions will be presented, at my discretion. Readers are invited to submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net

Mark Ortiz

WHEEL ALIGNMENT

I run an open-wheel mod (2400 lbs, asphalt). When squaring the chassis, is it best to align the right front to the right rear, or the left front to the left rear, and why? Our car is one inch wider in the front.

> You really shouldn't use the front wheels as a reference at all. If you change your front end settings, your wheel offsets, or your axle lateral position, you've lost the ability to recover your rear wheel alignment setting. Not good.
>
> Instead, you need a way to align your rear wheels with respect to the frame. Preferably, you should directly measure the alignment of both wheels – not one wheel, not the axle. That way, you are never caught out by a bent axle.
>
> To do this with string, you need to make yourself a set of stringing bars which let you run two parallel strings down both sides of the car. I buy two pieces of aluminum angle, clamp them together to form a "T", and match-file notches in the edges at the base of the "T". The spacing of the notches needs to be slightly greater than the width of the car.
>
> I set the bars on jack stands in front of and behind the car so the strings are about hub height, weight each end of each string with a nut, bolt, and pair of washers to hold them taut, make sure they lie in the filed notches so they're parallel, and measure from the strings to the wheel rims.
>
> On an asymmetrical oval track car, you have to pick a method of positioning the parallel strings relative to the frame – i.e. you have to decide what you call "straight ahead". I like to do this by measurement from two marked positions on a frame rail, preferably the left one since it gets bent less often. I pay attention to any existing marks or methods, and try to make my definition of "straight ahead" consistent with those.
>
> Once you've gotten to this point, you have a way of checking alignment of all four wheels that will not be thrown off by changes to front end settings, wheel offsets, axle position, or other

tuning variables. One other advantage: if you ever race cars with independent rear suspension, you'll be right at home.

Having said all this, I can tell you one thing about how the two methods you mention compare to each other for your car – assuming the axle is straight, and assuming you don't change anything else on the car. If you string on the left, the car will end up with more right rear lead than if you string on the right. That will tighten the car under power. Effect when braking depends on brake bias, driving style, and nuances of language. Many people report a looser car, especially those who mainly slow the car with the rear wheels.

If your car were wider in back than in front, you would get more right rear lead by stringing on the right instead.

Some people consider me an old lady on this wheel alignment stuff, but cars are extremely sensitive to rear wheel alignment. Undetected rear wheel alignment problems are a leading cause of mysterious handling quirks.

CASTER

We run asphalt tracks, 1/3 to ½ mile, with a 2800# car. I am looking for a little more front end bite to be able to "cut under" a competitor, starting at the center and coming out of the turn.

The very first inch of contact on the inside and outside of the RF tire seem to run 10 to 12 degrees hotter than the rest of the tire (157/145/155). We run caster settings of 1.5 degrees left/3.5 degrees right. Would raising the caster to 2.0/5.5 help? The driver steers the front wheels about 5 to 6 degrees to take these turns. Will we get too much caster-induced camber?

I'd try more caster. You don't need 3 ½ degrees split, unless you like the steering to pull left, which has nothing to do with making the front tires stick. I'd try 5 degrees both sides, or use whatever split the driver is comfortable with.

Based on your tire temperatures, the RF could use some more air. As for camber, many people will tell you that near-equal right and left shoulder temps mean your cornering camber is perfect. However, clients of mine who have measured tire temperatures while the car is running say that, as a rule of thumb, when the left and right shoulders of an oval track car's RF tire read similar in the turns, the left shoulder reads about 10 degrees hotter than the right in the pits after a run. So I don't think you'll end up with excessive camber. The car may even want a little more static camber. Try leaving it unchanged, and see what temps you get with more caster and more air.

Changing caster always changes bump steer, so you will need to check that.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2000

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

CHECKING REAR AXLES

Last month's piece on stringing cars to measure wheel alignment on the vehicle has proven very popular, and has also engendered further discussion. One topic that's come up has been how to check a rear axle before installing it. Here's a good method. You need a pair of v-block stands to support the axle, and a bump steer gauge with two dial indicators or the digital equivalent.

Install the bump steer plate where one wheel would mount. Position the indicators against the plate as you would when measuring bump steer, and zero the indicators. While you hold the bump steer plate so it can't rotate, have a helper rotate the axle housing, while you watch the dial indicators. Start at zero degrees pinion angle, or a pinion angle of your choice, and record readings at 90 degree intervals. Repeat for the other wheel. Compare the difference in indicator readings at positions 180 degrees opposite, then DIVIDE BY TWO to get camber and toe in inches at indicator span.

Many people use a somewhat similar method, turning the housing with the axle resting on its wheels, and just measure wheel-to-wheel toe, using a tape or a trammel bar. The problem with this is that both snouts can be cocked the same direction, and if you only measure wheel-to-wheel you won't pick that up. If you place a stationary object next to each tire, you can see any large individual wheel misalignment. An axle with both wheels aimed to the right, or both aimed left, can be made to work acceptably if you string the car as described in last month's issue. But if you just check wheel-to-wheel toe, then align the axle to the frame, you can get unexplained handling problems.

Since nothing is perfect in the real world, the question arises of what constitutes an acceptable axle. The answer is that a bit of overall toe-in is acceptable, perhaps even desirable. Toe-in adds drag, but stabilizes the car under power. Any amount of toe-out is unacceptable. Toe-out makes a car loose and directionally unstable under power. It is possible to use toe-out as a crutch to make a spool work on a road course, but this is a desperation measure.

Camber can be used to improve cornering. For road racing, we generally want a little negative camber on both wheels. Sometimes we may want more camber on one wheel than the other, for courses that are predominantly right or left. For oval track, we want positive camber on the left and negative on the right. On oval track cars, camber comes from both the axle and from tire stagger. With a track of 60", each inch of stagger adds about .15 degree of camber, negative on the right, positive on the left.

How much camber you can run will be limited either by the rules or by component life. On full-floater axles, with standard shafts and drive flanges, a prudent limit is .75 degree at the axle, plus whatever you get from tire stagger. This would correspond to a maximum camber reading of about .013" per inch of indicator span when checking the axle as described above. You can push this limit, depending on how much torque you are transmitting through the splines, at what wheel rpm, for how long. Axles with convex splines are available for more aggressive angles. In general, non-full-floating axles will not accept as much toe or camber as full-floaters. You only have one set of splines to work with, at the inboard end of the shaft. Cocking the bearing at the outboard end just gets you bearing failures.

TORQUE ARM OR PULL BAR?

Will a torque arm give more forward bite than a pull bar?

> Not necessarily. All either system does is generate an anti-squat jacking force, which only helps a little anyway. Either layout can be made to produce any amount of anti-squat desired.
>
> One advantage of pull bars is that they can be made short enough to go behind the driver, although many you see are too long for that. If you mount the pull bar left of center, its lift force adds wedge under power. In most classes where torque arms are legal, you can't fit an adequately long one behind the driver. Another advantage for the pull bar is that you usually save some overall and unsprung weight.
>
> In general, torque arms provide better damping of axle rotation. Their shocks act at a greater distance from the axle. Also, in most existing cars, torque arm layouts provide more anti-squat, and the anti-squat changes less with ride height. This doesn't always have to be the case, however. Actual geometry of the particular layout is very important.
>
> One interesting possibility with torque arms is to use two springs or coilovers - one ahead, one further back on the arm. You then use a conventional single spring for the rear one, and a conventional spring stacked with a very light "tender spring" for the front position. This makes the arm act short when grip is poor and axle torque is low, and effectively lengthen as axle torque increases. Most of the damping should be at the front spring, or forward of it. It probably is also possible to get a similar effect using two pull bars, or even a pull bar and a torque arm together. Bottom line: you can get good results with either layout, and both hold unrealized potential for those who are willing to reason from first principles and innovate.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2000

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

WINTER READING MATERIAL

As the racing season winds down, I'm getting lots of requests for recommendations for off-season reading. Therefore, I have decided to devote this issue of the newsletter to sources of information to help readers advance their self-education.

Please note that it is impossible to include all works of value in a brief list. Omission of a work does not imply condemnation. Likewise, inclusion of a work does not imply infallible accuracy, although I am making an effort to recommend documents that are largely accurate and widely well-regarded.

Let's start with some of the best general-purpose texts on vehicle dynamics.

Race Car Vehicle Dynamics by William & Douglas Milliken, SAE, 1995: Widely considered the standard reference work on the subject. Some portions are by other authors. Some chapters are highly mathematical; others are more conversational. Includes a chapter on the history of vehicle dynamics. Authors operate Milliken Research Associates, Inc. in Williamsville (Buffalo area), NY.

Tires, Suspension, and Handling by John Dixon, SAE, 1996: Another excellent reference. Good discussions of tire and suspension characteristics. Author is Senior Lecturer in Engineering Mechanics at The Open University, Great Britain.

Fundamentals of Vehicle Dynamics by Thomas Gillespie, SAE, 1992: A somewhat shorter and very readable vehicle dynamics text. Author teaches vehicle dynamics at University of Michigan and Ford Motor Company.

Car Suspension and Handling by Donald Bastow and Geoffrey Howard, SAE, 1993: Third edition of a work first published in 1980. More conversational than the above books, with a large appendix devoted to calculations. Authors are engineers with long and varied experience in the English automotive industry.

A couple of sources specifically on dampers (shock absorbers) deserve mention.

The Shock Absorber Handbook by John Dixon, SAE, 1999: Not really my idea of a handbook. More of an engineering text on dampers. Good coverage of principles and history of damper design and ride tuning theory. Not much on use of dampers as a tool to control transient behavior of race cars.

For a good discussion of how to use low-speed damping to influence transient handling, plus a good brief overview of damper function, see a pair of articles in the March and April 1996 *Sportscar*, entitled *Damper Physics: The Black Art of Adjustable Shock Absorbers*, by Neil Roberts. These articles, minus illustrations, can also be read and downloaded for free at www.gtf1.com under the titles of *Racing Dampers 201* & *202*.

On aerodynamics, a couple of good ones are:

Competition Car Downforce by Simon McBeath, Haynes/Foulis, 1998: A fairly brief, highly readable work specifically dealing with downforce – therefore readily applicable to racing.

Aerodynamics of Road Vehicles, ed. by Wolf-Heinrich Hucho, SAE, 1998: Compendium of papers and articles by a variety of authors on aerodynamics. Fourth edition, originally published in Germany. Good general-purpose text on the subject, including lift/downforce, drag, and directional stability.

Some lesser-known magazines I like include:

Racecar Engineering and *Race Tech*, published in England, both available by subscription in the US from EWA Magazines, 205 US Hwy. 22, Green Brook, NJ 08812, phone 1-800-392-4454, web site www.ewacars.com. More technical depth than most US car magazines.

Grassroots Motorsports, available from Motorsport Marketing, 555 West Granada Blvd., Suite B-9, Ormond Beach, FL 32174, phone (904) 673-4148, web site www.grmotorsports.com. Good coverage of amateur road racing, autocross, and vintage racing.

Speedway Illustrated, P.O. Box 37574, Boone, IA 50037, phone 1-888-837-3684, web site www.speedwayillustrated.com. New magazine from former *Stock Car Racing* editor Dick Berggren.

For a good catalog of automotive books, contact Classic Motorbooks at www.motorbooks.com or 1-800-826-6600. SAE sells books at www.sae.org/BOOKSTORE or (724) 776-4970.

I have published 6 articles in *Racecar Engineering*. These comprise a 2-part series on suspension interconnection and a 4-part series on asymmetrical cars. I have photocopies available at $2/article, check or money order to me (postage included; NC residents please add 6% sales tax).

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2000

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

SPRINGS, ROLL, AND CORNERING BALANCE

As I understand it, the stiffer the coil spring, the more weight is put on that corner, therefore planting the tire more. My question is, what is the difference between the stiffer coil vs.body roll? Example: A stiffer RF coil should put more weight on that tire, giving it more bite – therefore making the car steer better. However, you hear all the time about NASCAR teams taking spring rubbers out of the RF when they are tight to allow the chassis to roll over on the RF, making it turn. Is it because this allows the LR to lift, which makes it turn? Please explain, in simple terms.

Stiffening the right front spring, or adding preload to it, does load that tire more, and does make it produce more cornering force, in a left turn.

However, this comes at the expense of left front tire loading. The spring change can't change the total load on the front wheel pair, only the distribution of that total between the right front and left front. The spring change also can't change the total load on the rear wheel pair, the right wheel pair, or the left wheel pair, only the diagonally opposite wheel pairs.

So rear wheel loads when cornering are also affected by the front springs. The total rear wheel load doesn't change, but its right/left distribution changes, oppositely to the front wheels.

This means that with a stiffer RF spring, the front tires are loaded more unequally when cornering, and the rears are loaded more equally, than with a softer RF spring.

Now here's the key: When you concentrate the load on the outside tire, you lose more cornering power on the inside tire than you gain on the outside one. This is because grip from a tire increases with load, but at a DECREASING RATE.

Therefore, when you load the fronts more unequally and the rears more equally, that hurts available cornering force at the front and improves available cornering force at the rear: tighter car. That's the condition with the spring rubber in the RF. Take the rubber out, and you load the fronts more equally than before, and the rears more unequally. That helps stick the front, at the expense of the rear: looser car.

It works this way on dirt too, contrary to what some people will tell you.

SHORTY PANHARD BARS VS. LONG ONES

Can you help me understand the advantages and/or differences of the shorty Panhard bar (left chassis to left diff.) versus a long Panhard bar, for a dirt modified or Late Model? I know the shorter bars are more aggressive and plant the left rear more. Some say they're harder to drive. Generally speaking, what changes are needed when changing between these types of bars?

The shorty bar, especially when steeply inclined, jacks the left rear corner of the frame up in response to left-turn cornering force.

Since its angle increases as the car rolls and jacks, the effect increases exponentially as cornering force increases. The angle also varies a lot as the car goes over bumps, so the chassis becomes highly bump-sensitive.

Consequently, the shorty bar is at its worst on a rough, tacky, variable track, and at its best on a smooth, slick, highly consistent track.

Despite the popularity of shorty bars, I question their merit, because you can get as much dynamic diagonal as you want by other means, without the inconsistent behavior. I think a long bar is a clear advantage if the track changes and has bumps, as most tracks do. I doubt that a long bar is even a disadvantage on a smooth dry-slick track, if you set it up right.

Some people believe that the jacking force adds to the overall loading of the axle or tires. This doesn't really occur. The suspension can't push down any harder on the axle than the car pushes down on the suspension (if we ignore the weight of the suspension parts). Raising the rear of the car, even with cornering-induced jacking, doesn't increase total rear wheel load. Raising the left rear does increase diagonal percentage. This makes the rear stick better, at the expense of the front.

Things a long bar calls for, compared to a shorty bar on the same car, include some combination of the following:

1) More static diagonal
2) Lower Panhard bar height
3) Softer rear springs
4) Stiffer front springs

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

WIND TUNNEL TESTING FOR SHORT TRACK CARS

Many classes in short track racing involve cars that are aerodynamically very similar to each other, due both to tight bodywork rules and to copycat engineering. Most of these classes of cars are never tested scientifically to evaluate and improve their aerodynamic properties. This is due in considerable measure to the fact that wind tunnel time is expensive, and also to the fact that most short track racers don't know where to look for a wind tunnel that would have the time or inclination to work with them.

There is now a wind tunnel that is actually looking for short track cars to test. It's the tunnel at Langley, VA, which was originally built in the 1930's by NACA for running full-scale tests on fighters and similar-size aircraft. The tunnel is now leased to Old Dominion University. They have so far tested a DIRT modified, and are looking to test other types of dirt and pavement cars.

The DIRT mod test is described in a recent SAE paper by Drew Landman and Eric Koster. Among other things, the test revealed that the car generated net lift at the front axle, despite having a "snowplow" front end. This was of great interest to the team running the car, as they had been trying all sorts of chassis tweaks to cure an entry push.

This testing service is not free. In fact, it costs $1400/hour. For a one-day session long enough to do any good, you're looking at eight to ten thousand dollars. However, the payoffs are potentially enormous. Adding downforce can dramatically improve lap time. Small changes in the average pressures on the top or bottom side of a car can generate hundreds of pounds of vertical force, because the car's plan view area is so large.

The similarity of the cars in many classes means that racers can form groups to share the costs, provide one or more representative cars for test, and share the knowledge gained. If racers can get their friends – or maybe their car builders -- interested, there is a good chance that participants can make considerable gains at reasonable cost. I have volunteered to serve as a coordinator and compile lists of interested racers by car class. E-mail or phone me if you might want to pursue this.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

REAR CASTER

What is rear caster, and how does it affect the handling of a vehicle?

The term "rear caster" is normally applied to the side view inclination of the rear upright in an independent rear suspension. The concept is not applicable to beam axle rears. The upright often will not have an identifiable steering axis, but it should have some agreed features by which we can measure its inclination on a particular car.

We ordinarily don't have a steering mechanism at the rear, but we have one or more links per wheel that provide toe location and give the suspension its bump steer properties. So on many cars setting rear caster is our main way of adjusting rear bump steer. Sometimes there is a factory-recommended spec, but it is best to establish the desired rear caster setting by dialing in rear bump steer properties using a bump steer gauge, then measuring what upright angle or rear caster we have when the bump steer is the way we want it. Measuring this angle then becomes a quick way to recover the setting without going through the whole bump steering process.

In general, tilting the upright back at the top (adding caster) adds roll understeer – makes the wheel toe in in bump and toe out in droop. Tilting the upright forward at the top does the opposite. In some designs rear caster has no effect at all, even though you can adjust it. This is uncommon nowadays since it is useful to be able to tune rear bump steer. On some other cars, you can't adjust rear caster but you can adjust bump steer with shims on the toe control link.

STACKED COILOVER SPRINGS

Why do people use two springs of different rates stacked on top of each other on a coilover?How do you figure the rate of a combination like that?

The usual reason for stacking a soft coil spring on top of a firmer one is to get a stepped rising rate. One spring, usually the softer one, coil-binds before the coilover bottoms. With one spring coil-bound, the rate of the combination is the rate of the spring that can still move.

The rate of two stacked coils is less than the rate of either one alone. If we call the rates of the two individual springs A and B, and the rate of the combination C, then C = (AB)/(A+B).

In many cases, neither spring is coil-bound at static ride height. The idea is to get a rising spring rate to avoid bottoming or to cope with increased aerodynamic downforce at higher speeds. In other cases, the softer spring is so soft that it is coil-bound in normal operation and only serves to take up clearance at full droop. This is sometimes called a tender spring. Finally, it is also possible to use a special stop that causes the assembly to go stiffer in extension rather than compression. This is most often used on the left front on dirt cars, to tighten exit.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

REQUIRED FRAME STIFFNESS

Is there a way of calculating how much torsional stiffness a car's frame needs to have?

In road car engineering, it is customary to calculate a torsional natural frequency for the sprung structure, and compare this to the damped natural frequencies of the front and rear suspension in roll and warp. What this calculation tells you is whether you have enough stiffness to make the dampers (shocks) work as intended.

Any good race car should exceed this minimum by a substantial margin. I have never attempted to calculate or measure torsional natural frequency for a race car, and I doubt that anybody else does. Natural frequency depends not only on torsional stiffness, but also the mass of the sprung structure, and even the placement of major masses within the structure.

Measuring just the torsional stiffness of the frame is fairly common. This involves anchoring one end of the frame, typically at spring or rocker mounting points, and twisting the other end with a lever and a jack at the spring or rocker mounts. People don't try to reduce this measured stiffness to a calculated minimum; more is better.

There are some broad rules regarding how stiff is "stiff enough".

1) The stiffer the suspension's wheel rate in roll and warp, the stiffer the frame needs to be. Cars with beam axles at both ends typically have the softest wheel rates in roll and warp, and require the least torsional frame stiffness.
2) The stiffer the frame is, the more responsive the car is to tuning via roll resistance distribution. A flexible car does respond to roll resistance variation, but it takes a bigger change in roll resistance to get a given increment of cornering balance adjustment.
3) A car that relies on unequal front and rear roll resistance (one that corners on three wheels, or nearly so) needs a stiff structure more than one that has similar roll resistance front and rear.

4) A flexible car is more difficult to mess up with adjustments. It's also more difficult to fix with adjustments.
5) Stiff cars need smoother drivers. Jerky drivers often prefer more flexible cars.

Torsional stiffness of the whole frame, with loads applied at the spring or rocker mounts, isn't the only kind of rigidity that matters. All the load points that absorb forces from the suspension and steering need to be as rigid as possible. We can set forth a few general rules to help assure this:

1) Triangulate the frame as well as possible.
2) Feed loads, especially suspension loads, into the structure at tube junctions, not in the middle of a span.
3) Design brackets to minimize local torsion and bending loads. Make forces pass right through tube centerline intersections when possible.
4) When you cannot avoid feeding a load into a span, or when you are using stressed panels, design brackets so they spread the load so you minimize localized deflection.
5) Mount spherical joints and rod ends in double shear (plates on both sides) whenever possible.

MAKING BALLAST WEIGHTS

What does lead weigh per cubic inch? I want to make ballast weights using lead-filled 3"x3" or 4"x4" square tubing, weighing 25 and 50 pounds, and I need to know how long to make them.

Lead weighs about .41 lb./cu.in. in pure form – let's say about .40 if you're melting down wheel weights. If you use .125" wall tubing, 3"x3" weighs .39 lb./in. and 4"x4" weighs .53 lb./in. Internal cross-sectional areas are about 7.56 sq.in. for the 3"x3" and 14.06 sq.in. for the 4"x4".

So one inch of 3"x3" filled with lead weighs about 3.41 lb., and one inch of 4"x4" weighs about 6.15 lb.

Therefore, a 25 lb. weight using the 3"x3" would be 7.33" long, and a 50 lb. one would be 14.66". Using the 4"x4", you'd need about 4.07" for a 25 lb. or 8.13" for 50 lb.

Mounting holes, mounting hardware, voids in the lead fill, and using other wall thicknesses may cause minor variations from these theoretical numbers. If you make the weights a little on the large side, you can lighten them more easily than you can add material.

CAUTION: Lead is toxic. Any time you melt lead, be sure to provide good ventilation, and avoid inhaling fumes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
March 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

SAFETY ISSUES

With the deaths of Dale Earnhardt and three other national-level NASCAR drivers in the last 12 months, all from basal skull fractures in frontal impacts with walls, much attention is suddenly being paid to this type of injury, and this type of impact. I would like to offer my own thoughts on the matter.

THE HANS DEVICE

The HANS (head and neck support) device has been around for some 20 years in various forms. It has had trouble gaining acceptance for reasons of bulk, mobility restriction, and appearance. There is little doubt that it will prevent basal skull fractures in frontal impacts, and NASCAR drivers are suddenly ordering the device in much larger numbers.

I do not doubt the device’s effectiveness. However, it should be pointed out that there is only a certain range of impact severity where it really helps. (Impact severity depends on magnitude, duration, and number of accelerations.) Below a certain threshold of impact severity, the driver will survive without the HANS. Above a higher threshold, the HANS will save the skull and neck, but the driver will be killed anyway by other internal injuries. Therefore, the HANS may be part of the answer, but there is also a need to reduce impact severity in frontal contact with the wall.

SOFT WALLS

The need to reduce impact severity in frontal collisions with walls has led to various ideas for cushions for concrete retaining walls. Some tracks in the northeastern US are using blocks of styrofoam (styrene foam). NASCAR has done a few experiments with encapsulated styrofoam, as used for marine dock bumpers. I do not claim to be an expert on impact absorption devices, but I would like to make some general observations on what is required of such a system.

First, the cushioning system must preserve the good points of a simple concrete wall as much as possible. The uncushioned wall performs very well in glancing impacts, which are in fact the vast majority. The car slides along the wall, loses speed gently, and either comes to rest or continues around the track. It is vital that any cushion present a hard, smooth, continuous surface in a glancing impact, and not snag the car.

Second, on frontal impact the cushion must yield in a controlled manner, and not spring back immediately.

These two objectives are absolutely crucial. Additionally, it is helpful if the cushion can recover its shape and absorb more than one impact. It should be replaceable quickly, in sections, when it cannot recover. It should do its job without making a mess. It should be as compact as possible, though there is an inevitable tradeoff here between compactness and impact absorption. It is desirable that the cushion be non-flammable, though this is a secondary consideration. Cost is inevitably a factor. Finally, it is a good thing if the cushion can be made of reused or recycled materials.

I have no patents on wall cushions, and I am not promoting anybody else's system, but there are some particular design features that can help accomplish the objectives described above.

The cushion needs to have a metal, plastic, or composite facing. This should probably be part of the individual segment of the cushion, for easy replacement of both the facing and whatever is behind it. To avoid the problem of a car deforming one segment and being snagged by the next segment, the facing of each segment should overlap the facing of the next segment, in the direction of vehicle travel.

Behind the facing, there must be some kind of deformable structure to absorb impact. Possibilities here include cellular or foam structures, telescoping hydraulic units, and bladders containing air or water, with blowoff valves.

One approach might be to mold cushion segments in one piece, with a relatively thick facing and a thinner-walled honeycomb crush structure behind that. A possible material would be polyethylene, sourced from recycled milk jugs. This would gently recover its shape after impact, if the damage is not too severe. When the segment is too severely crushed for that, it would be replaced, and recycled.

Bladders, made of reinforced rubber as in fuel cells, could be refilled with air or water after an impact and be ready for another. Water exiting through blowoffs absorbs impacts very nicely. The principle has been used for vehicle bumpers. The water would wet down the track, but absorbent compounds used for other spills, or simple evaporation, might cope with that. Alternatively, air-filled bladders could be restored to shape with a pressure hose. Foams such as those used for earplugs and cockpit padding can also be used in bladders.

To fasten segments to the wall, one solution is cables through the wall, with fork clips.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

5TH COIL LOCATION AND RATE

I'm running a dirt Late Model with a torque arm.The car is good overall but could use some more forward bite on slick corners. I have gotten a lot of different opinions about how to set up my torque arm and 5th coil. Please explain what moving the 5th coil mount forward or backward on the torque arm does to forward bite, and what softening or stiffening the spring does to forward bite.

First of all, there's a reason you're getting conflicting information on the effects of torque arm length and spring rate on "forward bite" (forward acceleration capability, propulsive traction, ability to put power down): the effects are mainly imaginary. There are real effects, but they don't amount to much where the tires meet the track.

The shorter the torque arm length is, the more upward jacking effect it has. Contrary to what people will tell you, this in itself does not increase the loading on the axle. It just makes a portion of the load go through the fifth coil and a correspondingly smaller portion go through the right and left rear springs. It does lift the car, however, and there is a small but real effect due to that. The higher sprung mass CG causes slightly more load transfer to the rear wheels.

Therefore, you should run the torque arm as short as you can without encountering wheel hop.

When the arm is short enough to cause the rear of the car to rise rather than squat when you get on the power (geometry of your axle locating linkages also affects this), the effect of rear wheel rate split is reversed. This means that softening the left rear spring adds wedge under power and tightens exit, whereas with a suspension that squats in forward acceleration you tighten exit by softening the right rear and/or stiffening the left rear. Effect of front spring split on exit is the same either way: stiffer left front for tighter exit.

The spring on the torque arm doesn't affect how much the car lifts. It just affects how much the axle rotates. This cushions the application of torque to the wheels. Whether this really does anything is questionable. I've had a client who did blind back-to-back tests with different spring rates, and no spring at all, on a torque link (not a torque arm, but the effect is similar). Neither the driver nor the stopwatch could detect any difference between different spring rates, or the rigid link. It may be that the spring makes some difference in a very jerky application of power.

We can also say for sure that if the spring is too light, the axle will rotate too much and you will destroy U-joints or other parts. The minimum spring rate required to prevent this increases a lot as you shorten the arm. It varies inversely with the square of the arm length (measured from axle center to spring center), plus a bit. That is, the rate required with a 30" arm is MORE THAN $4^2/3^2 = 1.78$ times as great as with a 40" arm. If the car didn't lift more with the shorter arm, the factor would be exactly 1.78. But it does lift more. How much more depends on the rest of the system, but it's safe to say you would need at least 2 times the spring rate.

There are all kinds of interesting possibilities with torque arms and torque links, involving offset links and arms, multiple links and arms, multiple springs, snubbers, dampers, and so on. However, these are beyond the scope of your question, and involve fabrication and advanced setup knowledge. I am interested in working with car owners or builders who would like to pursue such possibilities.

SOFT WALL UPDATE

Last month I offered some general remarks about soft wall technology. I was gratified to read on Jayski that Petty Enterprises tested a segmented, molded plastic wall cushion during March. They instrumented a couple of Adam's old cars with a recording accelerometer and crashed them into a wall with and without the cushion. Reportedly, the cushion, which is about 2 feet thick, reduced peak acceleration from 100 g to 40 g. I didn't have anything to do with this, and I don't have any more specifics on the system they used, but that's definitely enough difference to save a life. If the system performs well in glancing impacts, and is reasonably priced, this looks very promising.

SOFT NOSES

Another subject of recent interest is deformability of the front ends on stock cars. People have expressed concern that front clips have become too crush-resistant in the search for torsional stiffness. This may be true, but there are ways to make a front clip torsionally stiff without making it so hard in a crash. A real space-frame front clip, in mild steel, with triangulation that doesn't run so nearly lengthwise, and no boiler-plate frame rails, would help – if it were legal. Also, I think the nose structure forward of the frame could be made to absorb more energy. Right now, it collapses very easily, and then the deformation stops or slows abruptly once it reaches the frame. More sheet metal, honeycomb, and/or plastic crush structure inside the nose molding could help a lot, and this could be added to existing cars. The energy-absorbing nose cones in CART show what can be done.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
May 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering answers to chassis questions. Readers may submit questions by mail to: 155 Wankel Dr., Kannapolis, NC 28083; by phone at 704-933-8876; by e-mail to: markortiz@vnet.net. Topics are also drawn from my posts on the tech forums at www.racecartech.com and www.rpmnet.com. Readers are invited to check out these sites, and to subscribe to this newsletter by e-mail.

Mark Ortiz

CHROME MOLY IN STOCK CARS

Recent discussions of frontal impact deaths in stock cars have produced allegations that some of the more advanced teams in Winston Cup, which build their own chassis, are now using chrome-moly (4130) steel in their frames, supposedly including even the large rectangular main rails. I have no way of verifying these reports, but I have noticed that various other classes in racing are allowing thinner-walled tubing in roll cages when 4130 is used. I am told the material is making major inroads in dirt Late Model construction as a result of this, as it allows considerably lighter cars.

Most stock car classes, including Winston Cup, neither require 4130 nor give it a wall thickness break. Therefore, there is no weight saving to be had from it. The reported reason for using it in Cup cars is to gain stiffness.

These developments prompt me to offer some information about what 4130 will and won't do for a car, and some of the less-recognized properties of the material. All this information has been published before. I would particularly like to tip my hat to Carroll Smith, whose book *Engineer to Win* offers considerable insight into this and other materials-related matters.

STIFFNESS

First thing we need to emphasize, and this will come as a shock to many racers, is that 4130 offers *no significant stiffness advantage* over mild steel! The gain is 1% or less. That's assuming equal wall thickness and weight, and identical design. If the frame is built with lighter gauge tubing to save weight, it will actually flex MORE.

Note that *stiffness* means resistance to deflection under load, short of the point of permanent deformation. Stiffnesses of all steels are remarkably similar. The best spring steels are less than 3% stiffer than mild steel, and 4130 as normally used in frames is actually closer to mild steel than to spring steel. Resistance to permanent deformation is called *strength*, and that's where chrome-moly offers potential gains. In normal use, a car frame is a *stiffness-limited structure*, and therefore does not gain performance significantly from the use of 4130, unless the rules allow us a frame or cage weight reduction. This applies until we crash the car. At that point, we suddenly have a *strength-limited structure*.

STRENGTH

Steels may be very similar in stiffness, but they differ dramatically in strength. Strength can be expressed in terms of *ultimate tensile strength*, *yield strength*, and *impact strength*. Ultimate tensile strength refers to tension load required to pull the material completely apart. Yield strength refers to tension load required to permanently deform the material. Both ultimate tensile strength and yield strength are measured by applying a gradually increasing load to the material. Impact strength, on the other hand, is measured by striking a sudden blow, of known magnitude. If the material deforms, but doesn't fracture, it passes. If it actually breaks, it fails.

The distinctions between these different kinds of strength are of great importance in race car construction. For portions of the structure that are close to the driver, we mainly want to prevent the wall, track surface, or whatever the car hits, from intruding into the cockpit. This means we want yield strength. A little deformation may cushion an impact, but we don't have room for large deformations. For portions further from the driver, we want a structure that deforms in a controlled manner, preferably outer portions first, absorbing energy as it crumples. To achieve this, we need graduated yield strength – more as the deformation approaches the driver. In both cases, we need impact strength; the structure can only do its job if it doesn't come apart on impact.

Mild steel does not respond to heat treatment; its carbon content is too low. Consequently, its properties remain about the same no matter how we weld it, heat it, or cool it. It never becomes brittle, as long as we don't introduce impurities while welding. Its yield strength is moderate, but its impact strength is good. In a crash, it deforms and absorbs energy, while resisting being torn apart or fracturing. This makes mild steel a good choice for most portions of a stock car frame. There is no significant stiffness penalty compared to 4130, and greater deformability when the car takes a hit.

Therefore, there may be a case for using 4130 for the cage in the driver's compartment of a stock car, but not for the front or rear clips.

4130, on the other hand, will harden when heated above its *critical temperature* and cooled rapidly. Cooled slowly, it remains soft. In the hardened condition, 4130 has great tensile and yield strength, but POOR IMPACT STRENGTH. It's hard and strong, but brittle. 4130 tubing, as supplied, is not in the hardened condition. It's *cold drawn and normalized* – fairly soft. In this condition, it will usually bend on impact and not fracture, just like mild steel, and it is somewhat stronger. But when we TIG weld 4130, we get a hard zone, not at the weld but half an inch or a little more from the weld. This happens because this region is heated enough to produce hardening, and is close enough to cool metal to be cooled abruptly after welding. The result of this is the failure pattern commonly seen in crashed 4130 frames: the tubes bend, the frame diamonds and twists, but there are some fractures near the welds. Not at the welds, but an inch or less from them.

Most people who build 4130 frames simply live with this. But there is a solution. After welding, heat the joint and the nearby metal to a dull cherry red with an oxy-acetylene torch, and allow to air-cool slowly. In warm ambient temperatures, just letting the area cool naturally is often sufficient. In colder ambient temperatures, or for insurance, a sheet metal or foil shield loosely fitted around the joint will slow the cooling sufficiently to avoid hardening.

Another approach, seldom considered nowadays, is to gas weld the joint instead. 4130 tubing was originally invented about 75 years ago for the aircraft industry. TIG welding was unknown. Aircraft structures were gas welded. Gas welding heats the metal surrounding the weld more than TIG welding does. This compounds distortion problems, but it does alleviate the problem of having very hot metal near cool metal that can quench it. In many cases, this will automatically eliminate the brittle zone.

REAR STAGGER VERSUS STATIC CROSS

We run a Late Model on pavement and have a question about the relationship of cross weight to stagger. In 1999, we used about 1.25 inches of rear stagger with 54% cross. In 2000, we increased our stagger to 2 inches, thus having to jack more cross into the car to keep it the same. How do you know when you have hit this balance right? I hear guys talking at the track about how much stagger they run, and it seems to vary widely.

> Stagger loosens the car through the entire turn. Most people report the biggest effect on exit, at least on pavement. Stagger has greatest effect when the rear tires are loaded the most, but the effect goes away when the wheels spin.
>
> Static cross has more effect when cornering force is moderate, meaning it does affect entry and exit more than mid-turn. Springs, conversely, affect mid-turn the most. Static cross tightens the car, except that the effect can reverse on entry if you slow the car mainly with the rear wheels.
>
> Stagger has more effect when grip is good. Static cross has more effect when grip is poor. So the way you blend stagger and cross affects how the car's balance varies as track conditions change. A car with modest stagger and modest cross will go loose on slick more than the same car with more stagger and more cross.
>
> So usually, consistency improves with more stagger and cross. The penalty comes in tire drag, especially down the straights, which increases as you add stagger.

2001 CONSULTING RATES

Hourly: \$40/hour. Monthly retainer: \$240/month. Season retainer: \$1200/year. Retainer option gets you unlimited phone time, plus e-mail within reason. Payment is by check or money order. I am also available to join your team for tests or races on an hourly + cost + per diem basis.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Topics and questions are also drawn from my posts on the tech forum at www.racecartech.com , where readers can see chassis consulting done for free. Readers are invited to subscribe to this newsletter by e-mail.

WEIGHT TRANSFER IN WINGED-OVER SPRINT CARS

I see winged sprint cars rolling to the left in left turns when cornering on dirt. How does this happen? What is going on with the wheel loads when the car does this? How does this behavior impact the car's response to spring or torsion bar changes?

Roll and load transfer in winged sprint cars make a very interesting analysis exercise, and a very good lesson in the distinction between roll and wheel load transfer ("weight transfer"). When equipped with a high wing with big side plates, a sprint car can transfer wheel load rightward while rolling leftward. Or, on a very slippery surface, the car theoretically could actually transfer wheel load leftward, though this would be unusual.

Disregarding aerodynamic forces for the moment, total wheel load transfer in any car is simply the car's weight, times lateral acceleration in g's, times its overall CG height, divided by track. This is the number of pounds by which the inside wheel pair loading is reduced, and the identical number of pounds by which the outside wheel pair loading is increased.

To estimate front and rear wheel load changes, this total load transfer is customarily broken down into the following components:

1) **Load transfer of the front and rear unsprung masses.** These act independently at the front and rear, and act through the tires but not through the suspension. The rear unsprung mass load transfer is substantial when the car has a live axle. Each of the two unsprung masses has its own center of mass, or center of gravity, approximately at hub height.

2) **Geometric load transfer** (sprung mass load transfer through the suspension members). For the front or rear wheel pair, this component is equal to the pounds of sprung weight at that end of the car, times lateral acceleration in g's, times roll center height at that end of the car, divided by track width at that end of the car.

For a car with high roll centers, such as a sprint car, this component is the largest one. For a car with low roll centers at both ends, such as a car with four-wheel independent suspension, this component is small. For a stock car or IMCA-style modified or dirt Late Model, with independent front suspension and a live axle in back, this component is small at the front and large at the rear. For a roll center below ground level, this component is negative (leftward, in a left turn).

3) **Elastic load transfer** (sprung mass load transfer through the springs, including anti-roll bars or other roll-resisting interconnective springs if present). To calculate this component, we need to know the wheel rate in roll for each wheel, and the wheel's distance, laterally, from the sprung mass CG plane (the longitudinal, vertical, plane containing the sprung mass CG). Using these wheel rates and moment arms, we calculate the elastic (spring-derived) roll resistance, or angular anti-roll rate, for each end of the car, in lb-in per degree of roll. Adding the front and rear angular anti-roll rates, we have the angular anti-roll rate for the whole car.

   That angular anti-roll rate resists the sprung mass roll moment about the roll axis. This roll moment is equal to the sprung weight, times lateral acceleration in g's, times the moment arm of the sprung mass CG about the roll axis. (This is measured perpendicular to the roll axis in side view, not vertically in side view or diagonally in three dimensions.)

   If we divide the roll moment (lb-in) by the overall elastic anti-roll rate (lb-in/deg), we get the amount of roll (deg). Working backwards, we divide the front and rear elastic anti-roll rates (lb-in/deg) by the roll (deg) to get the elastic anti-roll moments for the front and rear wheel pairs (lb-in). We then divide the wheel pair elastic anti-roll moment (lb-in) by the wheel pair track (in) and we have the elastic component of the load transfer (lb) at each wheel of the pair.

   Sprint cars have a soft wheel rate in roll, and a small moment arm from sprung mass CG to roll axis. Therefore, their elastic load transfer is a fairly small component. Their roll angles can be considerable, despite the high roll axis, due to the soft springing. Any force acting at a large distance from the roll axis can roll a sprint car a lot, without creating a lot of wheel load change. If you have a sprint car on wheel scales, and you push laterally on the roll cage, you can easily rock the car, while the wheel scale readings change relatively little, compared to rocking a stock car or a road racing car the same amount (which would also take a much stronger push).

   In a car with a low roll axis, elastic load transfer is the largest component. In a stock car, elastic load transfer is the largest component at the front, and relatively small at the rear.

4) **Load transfer due to lateral CG movement.** This component is very small in cars, but can be highly significant in heavy trucks and other tall vehicles, especially when the cargo can shift or slosh.

These four components comprise the factors in lateral load transfer, assuming a flat (unbanked) turn, constant speed, constant turn radius, and rigid tires – and ignoring aerodynamic influences. This is of course a simplified model – complex enough already though, right?

When the car is not in steady state cornering, there are additional factors. Actual cornering involves longitudinal accelerations as well as lateral ones. These effects are highly significant, especially in oval track racing. But for simplicity, let's just consider the effect of lateral acceleration increasing on entry and decreasing on exit.

When the lateral acceleration is changing, the car's roll angle will be changing. That is, the car has a roll velocity, not just a roll position. The suspension is in motion. This is called a transient condition. In transient cornering, we have additional load transfer factors:

5) **Frictional load transfer.** This includes forces generated by the shocks and also any other friction in the suspension. These frictions create front or rear frictional moments that act in parallel with the elastic moments, and may either add to the elastic moments or subtract from them. In general, frictional moments add to elastic moments any time the suspension is moving away from static position, and subtract from elastic moments when the suspension is moving toward static position. Hydraulic forces are velocity dependent. Dry friction forces are more or less independent of velocity.

6) **Minor inertial effects.** These arise from accelerations (changes of velocity) of the various elements of the car, particularly the sprung mass in roll. Properly understood, the centrifugal force and all other forces the tires have to overcome when cornering are also inertial effects. Roll inertia is minor compared to these. However, it can have significant effects in cars that roll a lot, in sudden maneuvers. Roll inertia explains why a vehicle may corner in a stable slide on a skidpad, yet overturn in a lane change or slalom test.

Breaking down the components of roll resistance like this allows us to predict the effects of changes to the various design and tuning factors that govern these components, in a sprint car or any other.

Now let's look at the forces from the wings. On a sprint car, we actually have two wings, but the rear one generates the most significant forces; it's big, and it's far enough from the bodywork to get ample airflow. It's also up high, and has big side plates. These are the factors that allow it to roll the car leftward, when the car is passing through the air at an angle (i.e., is in aerodynamic yaw).

Each wing generates a downward force. This acts approximately in the CG plane, or a little more strongly on the left wheels than the rights. It loads all four wheels. This increases the longitudinal and lateral forces available from all the tires, which of course is why people use wings. This, in turn, increases all accelerations generated by tire forces.

The downward forces from the wings, in themselves, have little effect on roll or amount of lateral load transfer. The added tire forces affect roll, but if anything they increase the tendency to roll rightward.

The wings also generate a rearward drag force. This unloads the front tires and loads the rears, and causes some rearward pitch.

Finally, the wings, especially the side plates on the main one, generate a leftward drag force when the car is in aerodynamic yaw. That's what rolls the car left. This force is small compared to the leftward force from the tires or the rightward inertia force (centrifugal force) at the sprung mass CG, but it acts on a huge moment arm about the roll axis. It's like your hand on the roll cage in the earlier example of easily rocking a sprint car on wheel scales without affecting the scale readings a great deal, compared to other kinds of cars.

In round numbers, a sprint car's roll axis is about a foot above the ground, the overall and sprung mass CG's are about a foot and a half above the ground, and the wing is about 6 feet up.

So the car sees a net leftward load transfer if the side force at the wing is more than ¼ as great as the centrifugal inertia force at the CG. This means the wing has to make ¼ of the lateral force. It has to make 1/3 as much force from air as the tires make with four sticky rubber footprints on the ground, aided by downforce. That's not impossible, but it takes a lot of air speed and a very slippery track.

However, the car will roll leftward if the side force at the wing is more than about 1/10 as great as the centrifugal inertia *of the sprung mass only*. This is a much easier condition to achieve.

In this state, the elastic component of the load transfer is negative (leftward), but the geometric load transfer is still positive (rightward) and is greater than the elastic load transfer. The unsprung mass load transfer is still positive (rightward). The load transfer due to sprung mass CG movement is negative (leftward), but very small. Thus, the big components of the load transfer are not reversed, just the small ones. The car is transfering wheel load rightward despite rolling leftward.

Since the elastic load transfer is reversed, the usual effect of spring changes is also reversed. More rear spring adds wedge when cornering and tightens the car.

Effect of roll center height is as usual. Raising the front roll center tightens the car; raising the rear roll center makes it looser.

Effect of static cross (diagonal percentage) is as usual. More load on the right front and left rear tightens the car, except perhaps on entry when slowing mainly with the rear wheels .

Shock tuning, assuming the track is smooth enough for low speed damping of roll motion to matter, works backwards. To add wedge on entry, add RR rebound damping or LR compression damping, or reduce RF rebound or LF compression. To add wedge on exit, add RF compression or LF rebound, or reduce RR compression or LR rebound.

Bump rubbers on the RF or LR tighten the car, when the car is leaning on the rubber.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Topics and questions are also drawn from my posts on the tech forum at www.racecartech.com , where readers can see chassis consulting done for free. Readers are invited to subscribe to this newsletter by e-mail.

CHASSIS TROUBLESHOOTING GUIDE

From time to time, people ask me to write a simplified chassis troubleshooting guide, as some other writers and car builders have done. I have shied away from this because so many things can alter the way chassis variables work. For example, changes to the left springs of an oval track car work one way in steady-state cornering on a flat track, and the opposite way in steady-state cornering on a steep banking. Rear spring split works one way if the rear suspension squats under power, and the opposite way if it lifts. Anything that adds diagonal percentage tightens the car (adds understeer), except on entry when the car is mainly being slowed by the back wheels, if the retarding force is strong and cornering force is moderate. So I always ask a client about the car, the track, and the driver's style before trying to solve problems (although in some instances a question does have a quick, simple answer).

However, it is possible to create a simple troubleshooting guide *for a certain set of conditions and assumptions*. I will offer such a guide here, but I want to be very explicit about the assumptions:

1) The inside suspension is assumed to extend rather than compress in steady-state cornering. That is, the turn is assumed to be fairly flat, grip is assumed to be fairly good, and relationship between ride and roll rates is assumed to be fairly conventional. This will make the guide applicable to relatively flat ovals. It will also be applicable to most road course corners, but I will assume for this discussion that we are examining a left turn. Road racers will have to "think mirror image" when applying the rules to right turns.
2) The suspension is assumed to be free of large jacking forces. In braking, the front suspension compresses and the rear suspension extends. Under power, the front suspension extends and the rear suspension compresses.
3) The front brakes are assumed to do at least half of the braking. The driver is not assumed to be tossing the tail out with the brakes.
4) The surface is assumed to be smooth enough so that sprung mass motion creates most of the shock movement, rather than bumps. This means we are looking at low-speed damping.

We will need to break the turn down into five portions, rather than the customary three:

Early entry: Braking is hard, and brake application is either steady or increasing. Cornering force is present, and increasing, but still moderate compared to rearward force from braking. This phase of the cornering process may not exist in many corners on a road course, or a severely paperclip-shaped oval. In such cases, the driver will do the hard braking in a straight line, and start to ease out of the brakes as he/she begins to turn in. But on most ovals, this phase will usually be present. I quite often see oval track drivers turn before they lift, or about the same time. This phase may also be present in road course corners that are fast, last a long time, or require an in-fast-out-slow line.

In this phase, roll position is rightward from static (left turn, remember), and increasing. Roll velocity is rightward, and increasing. Pitch position is forward from static, and increasing. Pitch velocity is forward, and may be increasing or decreasing.

The most active corners of the car are the right front and the left rear. The right front suspension's position is compressed from static, and its velocity is in the compression direction. The left rear suspension's position is extended from static, and its velocity is in the extension direction.

Late entry: Braking is diminishing, and ends at the completion of this phase. With a capable driver, cornering force should build as braking force diminishes.

Roll position is rightward from static, and increasing. Roll velocity is rightward, and may be increasing or decreasing early in this phase. Late in this phase, roll velocity will be rightward and decreasing. Pitch position is forward from static, and decreasing (because braking is diminishing). Pitch velocity is rearward.

The most active corner in terms of position is the right front. It will generally see its greatest compression somewhere early in the late entry phase. (This varies depending on several factors, including anti-dive, anti-roll, and roll rate/ride rate relationship.) The left rear is also active in terms of position. It will see its greatest extension. The most active corners in terms of velocity are the left front and right rear. The left front is extending; the right rear is compressing.

Mid-turn: Braking has ended. The driver feeds in at least enough power to overcome drag. The car either holds steady speed or gently begins to gain speed. The car is approximately in steady-state cornering. Forward acceleration is negligible. Lateral acceleration is at its maximum. Duration of this phase may be considerable with a smooth driver in a long turn, or it may be negligible if the turn is brief or the driver is abrupt.

Roll position is rightward from static, and stable. Roll velocity is near zero. If the mid-turn phase lasts a noticeable length of time and steady-state cornering is closely approximated, pitch position will be close to static, and pitch velocity will be near zero.

If steady-state cornering is approximated, all corners of the suspension are active in terms of position, and none are active in terms of velocity. The more the turn is banked, the more the rights are compressed and the less the lefts are extended. In corners around 15 degrees, the lefts neither compress nor extend much, and at steeper angles the lefts compress. As stated earlier, we are not considering such cases here.

The car is sensitive to all of its springs, especially the rights, and none of its shocks.

Early exit: The driver begins to increase power application and allow the car to widen its arc. Lateral acceleration diminishes and forward acceleration increases.

Roll position is rightward from static, and decreasing. Roll velocity is leftward, and increasing. Pitch position is rearward from static, and increasing. Pitch velocity is rearward.

The most active corner in terms of position is the right rear. It will see its greatest compression during this phase. The left front is also active in terms of position. It will see its greatest extension. In terms of velocity, the most active corners are the right front and left rear. The right front is extending (de-compressing); the left rear is compressing (de-extending).

Late exit: Similar to early exit, except that forward acceleration is now the dominant factor and lateral acceleration is fading into insignificance. Lateral acceleration will be zero at the conclusion of this phase, or very nearly zero, and forward acceleration will be at its maximum.

Roll position is rightward from static, less than before, and diminishing. At the conclusion of this phase, roll position reaches approximately static (car is going straight). Roll velocity is leftward, and decreasing. Pitch position is rearward from static, and increasing. Pitch velocity is rearward.

The most active corners in terms of position are still the right rear and left front, but the relative significance of right rear compression is diminishing. At some point in this phase, right front and left rear positions reverse from earlier phases: the right front goes into an extended position and the left rear goes into a compressed position. This means that spring changes on these two corners work backwards from the way they worked in previous phases. The most active corners in terms of velocity are still the right front and left rear.

We pay attention to suspension position because it is the key to spring tuning. We pay attention to suspension velocity because it is the key to shock tuning. Note that early and late exit are similar in terms of suspension velocity, but qualitatively different in terms of suspension position.

Now that we know what the suspension is doing in the turn, we are in a position to predict the effects of spring and shock changes. Remember that the rules which follow are only as good as your situation’s match-up to the one we’re modeling here. If your rear suspension lifts under power or compresses in braking, or you run on steep banking, the rules change.

I am also including rules relating to other tuning variables such as tire stagger, brake bias, and so on. I have tried to keep the chart to a single page, with reasonable size print, so it is a useful basic guide but cannot be a comprehensive reference work.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

CHASSIS TROUBLESHOOTING CHART

CAUTION: See the rest of this publication for important information on applicability of these rules.
Tuning factors listed are the *most* influential ones for the phase of cornering specified,
but are not the *only* influential ones.

	TO TIGHTEN LOOSE CAR	TO LOOSEN TIGHT CAR
Early Entry	Stiffer right springs, especially RF	Softer right springs, especially RF
	Softer left springs, especially LR	Stiffer left springs, especially LR
	Stiffer RF compression damping	Softer RF compression damping
	Softer LR extension damping	Stiffer LR extension damping
	More front brake	More rear brake
	Less front and rear tire stagger	More front and rear tire stagger
	More static diagonal percentage	Less static diagonal percentage
Late Entry	Stiffer RF spring	Softer RF spring
	Softer LR spring	Stiffer LR spring
	Stiffer LF extension damping	Softer LF extension damping
	Softer RR compression damping	Stiffer RR compression damping
	More front brake	More rear brake
	Less front and rear tire stagger	More front and rear tire stagger
	Higher front/lower rear roll center	Lower front/higher rear roll center
	Stiffer front/softer rear anti-roll bar(s)	Softer front/stiffer rear anti-roll bar(s)
	More static diagonal percentage	Less static diagonal percentage
Mid-turn	Stiffer front springs, especially RF	Softer front springs, especially RF
	Softer rear springs, especially RR	Stiffer rear springs, especially RR
	Stiffer front/softer rear anti-roll bar(s)	Softer front/stiffer rear anti-roll bar(s)
	Higher front/lower rear roll center	Lower front/higher rear roll center
	Less rear tire stagger	More rear tire stagger
	More static diagonal percentage	Less static diagonal percentage
Early Exit	Softer RR spring	Stiffer RR spring
	Stiffer LF spring	Softer LF spring
	Softer RF extension damping	Stiffer RF extension damping
	Stiffer LR compression damping	Softer LR compression damping
	Less rear tire stagger	More rear tire stagger
	Higher front/lower rear roll center	Lower front/higher rear roll center
	Stiffer front/softer rear anti-roll bar(s)	Softer front/stiffer rear anti-roll bar(s)
	Aim rear wheels leftward	Aim rear wheels rightward
	More static diagonal percentage	Less static diagonal percentage
Late Exit	Softer right springs, especially RR	Stiffer right springs, especially RR
	Stiffer left springs, especially LF	Softer left springs, especially LF
	Softer RF extension damping	Stiffer RF extension damping
	Stiffer LR compression damping	Softer LR compression damping
	Less rear tire stagger	More rear tire stagger
	Aim rear wheels leftward	Aim rear wheels rightward
	More static diagonal percentage	Less static diagonal percentage

For help with cars or situations not covered here, call Mark Ortiz at 704-933-8876 or e-mail markortiz@vnet.net.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Topics and questions are also drawn from my posts on the tech forum at www.racecartech.com , where readers can see chassis consulting done for free. Readers are invited to subscribe to this newsletter by e-mail.

COMPUTER PROBLEMS AT RACECARTECH

The tech forum mentioned above at www.racecartech.com has been down for about two weeks at this writing. Some people are under the impression that I own the site, but actually I just post there regularly. The timing of the problem leads me to suspect it may be related to a reported long-running tunnel fire in Baltimore which has damaged fiber optic cables. One would think the internet would have enough circuit redundancy to make damage to one cable inconsequential, but news reports are saying this fire has massively disrupted internet communications. Whether this is the problem or not, the forum will hopefully be up and running again soon. Meanwhile, my services remain available privately, as always.

THINGS THAT MAKE SPRING CHANGES WORK BACKWARDS

Last month I presented a chassis troubleshooting chart. I took care to point out that the recommendations in that chart apply only for a certain set of assumptions, including a fairly flat track and suspension with no large jacking forces. This month I'm going to supplement last month's information by discussing some factors that make spring changes work differently.

TRACK BANKING

In a flat turn, on most cars the inside suspension (left side, for a left turn) extends and the outside compresses. As the track banking gets steeper, the inside suspension extends less and the outside suspension compresses more. The car still rolls outward, but the entire chassis is pressed down due to the banking. Steep bankings are generally only encountered on oval tracks, so we will be discussing left-turn situations here.

Beyond a certain banking angle, the left suspension no longer extends, but compresses instead. This reverses the effect of left side spring changes: stiffer left front reduces instantaneous diagonal percentage and loosens the car, while a stiffer left rear adds instantaneous diagonal and tightens the car, in steady-state cornering. Right spring changes still work the same as on a flat track, except they have *greater* effect, due to the greater deflection.

I like to speak of a *critical angle* for track banking. This refers to the angle at which the left suspension neither compresses nor extends in steady-state cornering. The critical banking angle varies with springs, anti-roll bars, suspension geometry, aerodynamics, and amount of grip. It is usually somewhat different for the front and rear of the car. As a rule of thumb, critical banking angle for stock cars on asphalt is around 15 degrees.

The slipperier the surface, the smaller the critical banking angle. Or more accurately, the slipperier the track/tire combination, the smaller the critical banking angle.

The less the car relies on its springs for its roll resistance, the less the critical banking angle. If you increase the anti-roll bar stiffness and decrease the spring stiffness, that makes the car corner at a lower ride height on the banking, and reduces the critical banking angle. This has been a major issue in Winston Cup lately. Some teams have tried outrageously soft springs on the front, with very stiff bars, to make the car corner lower on relatively flat tracks. This was the reason for the bump rubbers which NASCAR recently outlawed.

Raising the roll center on a beam axle, and softening the springs, also reduces the critical banking angle. Raising the roll center on an independent suspension can have a similar effect, although we may also encounter jacking effects that can reduce or reverse this. On beam axles, we can have jacking effects that are separable from roll resistance. For example, if we raise the left end of an "across-the-car" (or long) Panhard bar, and lower right end an equal amount, we make the car jack up in a left turn, with little effect on the roll center. Such a change increases the critical banking angle.

It is difficult to calculate the critical banking angle precisely, but it is quite easy to know when we're there if we have electronic data acquisition. When we are close to the critical banking angle, the ride height traces from the left wheels will correlate heavily with longitudinal acceleration, throttle position, and brake pressure, and will be largely insensitive to lateral acceleration. In this situation, the car's steady-state cornering balance is insensitive to left spring changes, but left spring changes do affect its entry and exit characteristics. This means we can tune mid-turn properties with the right springs, and tune entry and exit with the lefts.

LARGE JACKING FORCES

Note that steep bankings reverse the effect of left spring changes because they reverse the usual direction of suspension motion on the left side of the car. It is a basic rule that anything that reverses the usual direction of suspension motion at a particular corner of the car reverses the effect of spring changes at that corner. The other very common cause of reversed suspension motion is large jacking forces: forces that try to extend or compress the suspension when the tire generates horizontal forces.

Designers deliberately build jacking properties into suspensions to resist roll and pitch, and to raise the center of gravity under power, which increases load transfer to the rear wheels.

At the front end, we call upward jacking forces (ones that try to extend the suspension) in braking *anti-dive*. At the rear, we call downward jacking forces (ones tending to compress the suspension) in braking *anti-lift*. In rear-wheel-drive cars, upward rear suspension jacking forces under power are called *anti-squat*. It is also possible to have *front anti-lift* under power when the front wheels are driven.

The front suspension is said to have *100% anti-dive* if the jacking force is exactly sufficient to prevent the suspension from compressing under braking. Most cars have less anti-dive than this, and many have none at all. When the anti-dive is zero, jacking forces are absent in braking, and the forces tending to compress the front suspension are resisted entirely by the springs. If downward jacking forces are produced in braking, anti-dive is said to be negative. Negative anti-dive is also referred to as *pro-dive*.

If a car has exactly 100% anti-dive at the front, the left and right front suspensions neither compress nor extend in braking, regardless of spring rates. This means that, in pure braking, front spring choices have no effect on instantaneous diagonal percentage. If anti-dive exceeds 100%, the front of the car actually lifts in braking, and instantaneous diagonal percentage increases if we soften the right front spring or stiffen the left front – opposite of the usual.

Note that these are mainly hypothetical cases, since most cars have far less than 100% anti-dive. Most stock cars nowadays have moderate anti-dive at static, and lose anti-dive rapidly as the suspension compresses, sometimes going to pro-dive. When a car has pro-dive, front spring changes affect entry balance in the usual way, only their effect is greater. When a car has moderate anti-dive, front spring changes affect entry balance in the usual way, only their effect is less. These comments also apply to individual corners of the car: when we have pro-dive on the right front and anti-dive on the left front, entry is highly sensitive to right front spring changes, and much less sensitive to left front spring changes.

Similar effects occur at the rear in braking. If the car has 100% anti-lift, the rear suspension neither extends nor compresses in braking, and spring choices have no effect on instantaneous diagonal percentage in pure braking. Of course, to meaningfully say that the car is loose or tight, we must have some cornering, and therefore some roll, along with our braking, and front and rear springs will have effects on instantaneous diagonal percentage due to their effect on front and rear roll resistance, even in the case of a car with 100% anti-dive and 100% anti-lift at all four corners.

Unlike 100% anti-dive, 100% anti-lift (or more) is common in road cars, or in production-based road racing sedans and sports cars. Cars that react rear brake torque through a simple trailing arm or semi-trailing arm generally have more than 100% anti-lift. Examples include C2 and C3 Corvettes, many BMW's, Porsche 911's and 356's, and all but the first Mazda RX-7's. Some dirt modifieds and Late Models also have more than 100% anti-lift, though others have pro-lift. Anti-lift exceeding 100%

will be evident in data acquisition outputs or trackside observation: the rear will drop rather than rise when the car slows. The anti-lift effects may be different for engine braking than for actual brake forces, and a car can have anti-lift on decel yet have pro-lift on the brakes. A typical dirt car 4-bar rear with a torque arm, and calipers on the birdcages, usually exhibits this mix of properties. Most trailing arm independent rears are the opposite: pro-lift on decel, anti-lift on the brakes.

Under power, the right and left rear suspensions may either extend or compress. In addition, live axle suspensions transmit driveshaft torque, which tends to extend the left rear suspension and compress the right rear, adding instantaneous diagonal percentage. The effect on suspension position is called *torque roll*; the effect on wheel loads is called *torque wedge*.

If the rear suspension as a whole neither compresses nor extends under power, that is 100% anti-squat. In this case, rear spring changes have little effect on wheel loads in pure forward acceleration, except that in live axle rears, torque roll and torque wedge still occur unless the suspension is carefully designed to eliminate this. Softening either rear spring, or both, increases torque roll and torque wedge, regardless of overall anti-squat.

Stiffening the front anti-roll bar decreases torque roll but increases torque wedge. Stiffening the left front spring likewise decreases torque roll but increases torque wedge. Stiffening the right front spring also decreases torque roll and increases the *torque-related component* of wedge change. However, in most cases the unloading of the front end under power extends the right front suspension more than torque roll compresses it, so the net effect of a stiffer right front spring is to de-wedge the car in pure forward acceleration.

We can also speak of anti-squat effects at each rear wheel individually, even in live axles, and we may include driveshaft torque effects when considering these, or not – as long as we don’t forget that the driveshaft torque is there. When either rear spring extends under power rather than compressing, the effect of spring changes at that corner of the car is reversed. A common instance of this occurs on the left rear of typical 4-bar dirt Late Models, where a softer LR spring will tighten exit.

INTERACTION OF THESE EFFECTS

As if we didn’t have enough complexity just considering these effects in isolation, in the real world we often have banking effects and jacking effects acting together. Without electronic data acquisition, it may be difficult to know or predict whether, or when, a particular corner of the car compresses or extends. However, we do know this much: if the actual direction of suspension motion is opposite to what we’d get on a flat track with small jacking forces, effects of spring changes will be opposite too. If motions are bigger, effects of spring changes are bigger. If little motion occurs, spring rate will have little effect.

With electronic data acquisition, we can use these principles to predict effects of spring changes in particular parts of the turn, even with complex jacking/banking combinations. And even if we don’t have electronic data acquisition, these principles can still help us make sense of our observations.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net.

BAD NEWS, AND GOOD NEWS

The tech forum at www.racecartech.com, arguably the hottest chassis forum on the internet, has now been down for a month and a half. Over the past year or more, I have answered questions for free on this forum, and have gotten quite a few paying clients as a result. Since people have come to think of that site as the place to find me, I would prefer to stay there. But if it stays down much longer, I'm going to have to start my own message board, or possibly find another one to hang out on. I have e-mailed the owners of the racecartech site, asking them what we should expect. I have received no answer. Stay tuned. Meanwhile, I am definitely still in the consulting business, though the "economic slowdown" has hit me hard.

On a brighter note, I have recently come to an agreement with *Racecar Engineering* magazine to publish a monthly column based on this newsletter. For those unfamiliar with that magazine, it is published in England and features deeper and better tech articles than the car magazines at your local supermarket. Subscriptions for the US, Canada, and Mexico are handled by EWA Magazines, 205 US Hwy. 22, Green Brook, NJ 08812; phone 1-800-392-4454; website www.ewacars.com.

STOCK CAR SAFETY ISSUES

The month of August has seen the release of NASCAR's report on the death of Dale Earnhardt, and the first public demonstration of the Humpy Bumper. There was also one independent study released which advocated crush structures in stock car noses.

On the whole, I find NASCAR's official conclusions, as presented by Dr. James Raddin, to be logical and consistent with the physical evidence as presented. The only real room for doubt lies in the difficulty of verifying that physical evidence, particularly in view of the long delay before the separated belt was announced, the rescue worker's insistence that the belt was not separated, and the refusal to allow public access to the autopsy photos.

Concerning the belt, I agree that it was installed improperly, if the illustrations we've seen are correct. I also agree that the shoulder harness was installed with an inordinately long run to the anchor point. I agree that the lap belt's mounting could cause failure, at approximately the point where the belt passed through the hole in the seat, which is also where the adjuster was. But we are being asked to believe that this happened during the crash; that the separated portion of the belt was displaced 4 to 8 inches after the crash (therefore more than that during the crash); that the driver moved far enough to severely deform the steering wheel and break 8 ribs and his sternum against it - and that somehow the harness was still tight enough to make it hard to release the buckle, and the separated belt was somehow invisible to everybody until some days after the accident. Maybe I'm missing something, but I don't see how that's physically possible. I can see how the belt could have separated, but I don't see how it could have done so without the separation being immediately apparent.

Looking at published photographs of the belt, I agree that it appears to have been pulled apart, not cut. However, this appears to have happened in two stages, not all at once. I say this because the bottom third of the break has a different appearance than the top two thirds. The bottom third of the break appears to have been initially torn over a sharp edge, such as the adjuster, and operated in this condition for a while, with the tear slowly growing. If this were not so, the fibers on the left (in the photo), or working, segment would not be bent back and flattened the way they are, and the edge of the right, or unstressed tail, segment would not be so cleanly cut. The entire break would be ragged, without the fibers being bent back and flattened, like the top two thirds.

Here's a possible explanation: the belt did fail, or was in the process of failing, due to improper installation, *but not to the point of complete separation*. It stretched more than one would normally expect, and the shoulder straps did too. But the belt still held together. Then some time after the crash, somebody discovered the damage to the belt, and somebody pulled it apart the rest of the way.

I realize the enormity of this suggestion, and I realize that there is no direct proof that this occurred. However, this theory would explain the initial absence of any mention of the belt failure, and the 5-day delay in announcing it. I have heard no other plausible explanation of these things.

This theory does require a concerted attempt to deceive the public, and this would require a motive. Such a motive is not hard to discern, however. The Earnhardt crash was the fourth in a string of similar ones, resulting in similar fatal injuries. Without belt failure as an issue, the focus inevitably shifts to car and wall construction. Changing cars and walls costs big money, and is fraught with the perils of developing new technology, which may not always work as intended. Remember that the France family owns big interests in a number of tracks, as well as a controlling interest in NASCAR. With the belt failure controversy raging, the waters are muddied sufficiently that changes to cars and walls can be deferred for some time – perhaps indefinitely.

I do not prefer conspiracy theories over simpler, less dramatic explanations. But I will believe in a conspiracy more readily than I will believe in a miracle. And what we're seeing with this belt issue appears to me to be unexplainable except as a conspiracy or a miracle. I await the explanations of those who think otherwise.

NASCAR's position on crush structures in the nose is that there may be "unintended consequences" in the form of cars moving each other around more when light contact occurs. This concern is not entirely unwarranted, but I certainly think the concept of energy-absorbing noses should be tested, not merely dismissed.

Also, it is my understanding that the object of nudging another car is to move it, at least when the nudge is deliberate. In the case of purely accidental contact, we are faced with the classic design tradeoff that we always encounter when designing any kind of cushioning system in a finite space: softer is better for light forces, but a soft cushion bottoms out when forces are greater. This dilemma is the same one we face with seat pads, roll bar pads, and even vehicle suspension systems.

I was not present for the demonstration of the Humpy Bumper, the carbon fiber structure designed to fit inside stock car noses. Published reports say a test car was driven into a wall at an angle, with an impact speed of 40 mph. I am not clear whether that is the car's speed, or the component of its velocity perpendicular to the wall, or the reduction of its velocity during impact. The last of these is the one that would create a similar condition to the Earnhardt crash. The car's velocity decreased 42 to 44 mph during the two impacts combined, from an initial speed of around 160 mph.. Reportedly, the test car's nose crushed visibly less than would normally be expected, and the car bounced off the wall, requiring some reporters to move fast to evade it.

I am not certain how much to make of this one instance, but it is not desirable for an energy-absorbing structure to be resilient: it should not spring back. It should crumple, and stay squashed.

For reasons of both cost and non-resilience, I would suggest looking hard at crush structures of mild steel, aluminum, or plastic such as ABS or polyethylene. How about "egg-crate" radiator ducts, with internal vanes forming a crush structure while also maintaining orderly air flow to the radiator?

The independent report that suggested crush structures proposed styrene foam, encapsulated in sheet aluminum. This has possibilities, especially in the fenders, ahead of the wheels and alongside the radiator duct. Aluminum honeycomb or egg-crate has potential in that area as well.

Whatever material is used, the structure should incorporate *graduated* rigidity or yield strength. The portions closest to the nose molding should deform easily, and the structure should offer progressively greater resistance as deformation moves further inboard. This increase in resistance should be as gradual and stepless as possible. With an egg-crate structure, this is easily achieved by using more panels, and/or thicker material, for the inboard portions.

I sent copies of my April newsletter, which addressed adding crush structures to both noses and walls, to Mike Helton and Gary Nelson. NASCAR sent them back, with a letter saying that their policy is to not accept any unsolicited suggestions! Remember that next time you see a NASCAR representative on TV saying they are receiving new ideas all the time. No doubt they are receiving them, but this evidently doesn't mean they read them.

Lest any one suppose that I am unconditionally critical of NASCAR, I do think they are making good decisions in establishing a new safety research center in Conover, NC, and putting impact data recorders in the cars. Hopefully, crush structures will be among the items investigated.

TIRE DATA

Why is it so difficult to find data for mathematical modeling of tire properties? Do tire companies have the ability to test the forces a tire generates at various slip angles, loads, camber angles, and so on? If so, why don't they make this information public?

> Tire companies do have machines that can test a tire against a simulated road surface, at controlled normal force (vertical load), camber angle, and slip angle, and measure the drag and lateral forces the tire generates. Sometimes they also contract out this work to a laboratory like Calspan, where much of the equipment used for this was first developed.
>
> The machines use a large wheel or a very strong belt to simulate the road surface. This of course most closely simulates pavement, not dirt. Early machines also were built that rolled the tire along the ground, attached to a heavy truck.
>
> The indoor machines with simulated road surfaces were developed to produce more accurate and repeatable measurements. The fact that this was necessary tells us something important about tire behavior: in the real world, what occurs doesn't just depend on the tire. It depends at least as much on the road surface and the weather. Therefore, tire data are only meaningful when taken under very carefully controlled conditions. To model tire behavior accurately for the real world, testing needs to be done under a variety of carefully controlled conditions, and the effects of these changes have to be included in the report.
>
> As if this weren't enough, tires themselves are highly variable in their behavior. Their properties vary with age, heat cycling (itself sensitive to amount and speed of temperature change), wear, inflation pressure, tire temperature, air temperature, road temperature, air flow to the tire, vehicle speed, combination of loadings in multiple directions, manufacturing variations, rim width, and other factors. The tire even has different properties as it proceeds through a turn, because it rapidly heats up. It heats faster at a high road speed than at a low road speed.
>
> Any attempt to test and analyze how these factors play off against each other results in a voluminous report. The best we can do for purposes of mathematical simulation is to assume a simplified tire model, preferably averaged from such a report, that we can use for comparative calculations when varying other factors. If we are dealing with incremental changes to the car, on known tracks, with many runs already logged, as in F1, then reasonably accurate lap time prediction is possible. If we are trying to predict less familiar situations, accuracy available from the simulation inevitably diminishes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net.

RACECARTECH DIES (I THINK)

The owners of www.racecartech.com , where I used to give free advice, have announced that they are closing the site due to unprofitability, effective Sept. 30. I would take this as final word, only a client of mine has recently expressed interest in possibly taking over the site. By next month I'll know if anything is going to come of this.

BALANCING THE CAR WITH CAMBER AND TIRE PRESSURE

Thanks for your July newsletter – the troubleshooting guide is very meaningful, even for a road racer such as myself. Could you comment on how camber and tire pressure might be used to address tight and loose conditions?

A tire has a preferred camber and pressure. If you go either direction from the optimum, you lose grip. Therefore, if you deliberately depart from best camber and pressure, you are throwing grip away. Ordinarily, you don't want to do that. It's better to balance the car by managing wheel loads, because that way you are redistributing the available "grip budget" rather than throwing some away at the end that sticks better.

Okay, that's the basic answer. But beyond that, there are complex and fascinating nuances.

To begin with, optimum camber and pressure for longitudinal acceleration (braking and propulsion) are different than for lateral acceleration (cornering). For longitudinal forces, we want the tire at zero degrees camber – straight up. For cornering, we want it leaning slightly into the turn if possible. The optimum pressure for longitudinal acceleration will generally be less than for lateral. So we have to compromise.

The compromise we strike has some effect on the tire's behavior in various parts of the turn. With a less aggressive camber setting, we may improve the tire's grip in early entry and late exit, when longitudinal forces are great and lateral ones are moderate, and reduce its grip in the middle of the turn. We can also do this with a moderate pressure reduction.

Reducing pressure will, in some cases, make the tire work better in the first part of a turn at the expense of the latter part, particularly on ovals. This happens because with lower pressure, the contact patch is longer. The rearmost portion of the contact patch has a form of interaction with the road surface which is characterized by cyclical sliding and reattachment. This builds huge amounts of heat in the outer layer of the tread rubber. Surface temperatures in normal racing use can exceed 400 deg. F. at corner exit, in a tire that shows half that temperature when we measure it after a run. This means the rubber temperature is well above optimum in the last part of a turn.

At lower pressure, the tire heats faster as it corners. Therefore, slightly lower pressure may improve grip early in the turn, and reduce grip through the last 2/3 of the turn. Lower pressure will also usually help the tire early in a run, at the expense of the rest of the run.

There is some difference between oval applications and road racing applications or street or autocross use. In oval racing, especially ovals of a mile or more, the turns are big and last a long time. On a 1.5 mile track, a 180 mph lap takes 30 seconds. The turns are around half the lap distance, and are taken at somewhat lower speed than the straights. This translates to sustained cornering for up to ten seconds, depending on the track and where we define the beginning and end of the cornering process to be. Steering inputs are small and gentle. This means that tire heating during the turn is a major factor, and tire rigidity on turn-in is not very important. Consequently, the driver may report better front grip on entry with slightly reduced front pressure on an oval. Conversely, when turns are tight and last only a second or so, steering inputs are large and abrupt, tire heating is less of a factor, and tire rigidity counts more. In such situations, turn-in may be improved with *higher* front pressures than optimum for steady-state cornering.

One problem with using these principles to tune the car is that we don’t really have very precise control over tire pressure. If the front tires are a little on the soft side and the rears are ideal when the sun is behind clouds, what happens when the sun comes out? What happens when the driver presses a little harder and the tires heat up? The fronts will optimize, and the rears will have too much pressure. The car will get looser. Tuning a race car with tire pressures makes it inconsistent. Therefore, when possible I try to make the fronts and rears optimize together, and be too soft or too hard together as conditions vary.

You hear a lot about tuning tire spring rates with pressures. This is mostly smoke, especially in cars with compliant suspensions, such as stock cars. Contact patch length, contact patch loading distribution, and tire heat buildup are the big factors that change when we vary pressure. When a team adds or subtracts a pound on one tire on a pit stop, they are mainly throwing away a little grip on that corner to balance the car. Often they can either add or reduce pressure to kill the grip, but the effect comes later in the run, and later in the turn, if underinflation is chosen, and earlier if overinflation is used. Using tire pressure instead of suspension adjustment during a race makes sense because balance is so important, and adjusting suspension during a pit stop costs time and track position.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
November 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

LIFE AFTER RACECARTECH

Since the Racecartech site has closed, I am making occasional posts at www.rpmnet.com and the Weekend Auto Racers forum at http://netpaths.com/cgi-bin/tech.pl.

THE BIG TRACKS AND THEIR BIG WRECKS

There is now widespread recognition of the prevalence of huge, multi-car wrecks at almost every race at Daytona and Talladega. Most observers agree that this relates to the large packs of cars running close together at high speeds. New rules against putting the left wheels on the apron to pass have not solved the problem. Restrictor plates on the engines, aerodynamic drag-inducing additions to bodies, and restrictions on spring and damper calibration have all been instituted with minimal impact on the problem.

I wish I had a simple solution myself, but I don't. I do, however, have some thoughts on the fundamental causes of the problem.

There is an "elephant at the dinner table" here – a problem that really should be clear to everybody, but that nobody wants to mention. That problem is the design of these tracks: they are 2 ½ miles or longer, with bankings of 30 degrees or more. This means that unless the cars can do well over 200 mph, they can run wide open all the way around the track. If they run fast enough to be grip-limited in the turns, they become really hard to contain when they get out of control, and impacts occur at pretty ferocious speeds. Concern about these factors is what has led to all the efforts to slow the cars.

But now that the cars can run wide open all the way around the track, drafting becomes the key to victory, just as it is in that other power-limited sport, bicycle racing. Bicycle racing is also subject to massive pileups, for very similar reasons. This fact should tell us something about the prospects for eliminating the pileups by slowing the cars. Slowing the cars further may reduce injuries and deaths somewhat, but the pileups will still occur.

It has been suggested that smaller engines be used rather than restrictor plates. This might improve the throttle response of the engines, but it will not change the fact that they will be running at full power except when drafting another car. We might expect a bit more passing, but these races already have an abundance of passing.

As long as there are no turns small or flat enough to slow the cars, there is no way out of the dilemma. Either the cars will run in huge drafting packs, or they will run so fast that they will be excessively dangerous to drivers and spectators when they crash, even though they will probably crash in smaller numbers at a time.

The only avenue that might offer some promise is to reduce the grip of the tires sufficiently so that even with modest power, the cars will have to slow significantly for the turns. Trouble is, that would require a pretty significant grip reduction. I hesitate to advocate making the cars really slidey when the wind already moves them around so much, and when they already take forever to stop following loss of control.

SELECTING SPRINGS

When a new car is built, is there a way to calculate best spring rates to try first, and thereby minimize testing?

> There is no method that works for all types of car. Spring rates are only one of many tuning and design variables, and all these variables interact with each other. You can measure some of these factors; others you can only guess at. With more resources, you can measure more, and guess at fewer. For example, with wind tunnel testing, aerodynamic effects become much less of a guess. The track or road conditions are another important variable. Any given car will want different setups, including springs, for different conditions.
>
> A common method is to simply borrow a spring combination from a car that is estimated to be similar, and test from there. Many race cars are built to existing class rules and are similar to ones already running, with small variations. When using this method or any other, it is vital to go by *wheel rate*, the rate at the wheel. Many suspensions have one wheel rate in ride and another in roll, and you need to look at both.
>
> In oval track racing especially, many different spring combinations can be made to work fairly decently, because we can vary diagonal percentage, stagger, and other factors to suit the springs. Even in road racing and figure 8 racing, the car can be adjusted to suit various spring packages by varying roll center heights, anti-roll bars, and aerodynamic surfaces. In many cases, the difference between a superior spring package and an inferior one will be seen in the car’s transient behavior, and its response to varying conditions.
>
> That said, we can state some definite rules for the ways spring choices relate to other factors.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

RELATIONSHIPS BETWEEN SPRING CHOICES AND OTHER FACTORS

OTHER FACTORS REMAINING EQUAL:

Stiffer springs at one end of the car call for stiffer springs at the other end.

Wider or stickier tires at one end of the car call for stiffer springs at that end and/or softer ones at the other end.

Rougher surfaces call for softer springs, with higher ride heights. This assumes that suspension travel is available to do this. If bottoming is a problem, stiffer springs may be needed to overcome that.

More aerodynamic downforce calls for stiffer springs.

A stiffer anti-roll bar calls for softer springs at that end of the car, and/or stiffer springs at the opposite end.

To correct an understeering car (tight condition), use stiffer rear springs and/or softer fronts.

To correct an oversteering car (loose condition), use softer rear springs and/or stiffer fronts.

More diagonal percentage calls for softer front springs and/or stiffer rears.

More rear tire stagger calls for softer rear springs and/or stiffer fronts.

A lower roll center at one end of the car calls for stiffer springs at that end. A lower roll axis calls for stiffer springs at both ends.

More weight at one end of the car calls for softer springs at that end and/or stiffer springs at the opposite end, to preserve similar cornering balance. This requirement is at odds with the need to add spring rate proportionally to weight, to preserve similar ride characteristics. In some cases it will not matter whether similar ride characteristics are preserved. In other cases, we may want to work with anti-roll bars and suspension geometry as well as springs, to reduce relative roll resistance at the end that's seeing the weight increase.

On banked ovals, there is a *critical angle* of banking, at which the left suspension neither compresses nor extends in steady-state cornering. For typical stock car chassis on asphalt, the critical angle will be around 15 degrees. On dirt, the critical angle will be somewhat smaller than this. Below the critical angle, left side spring changes have effects as described above. Near the critical angle, left spring changes have little effect on steady-state cornering balance. Above the critical angle, effects of left spring changes on steady-state cornering balance work *backwards*: a stiffer left front loosens the car; a stiffer left rear tightens the car.

On ovals, assuming the car does not generate large jacking forces (true for most pavement cars, but not for some dirt cars), and assuming that the car is not slowed primarily by the rear wheels, stiffer left springs and/or softer rights loosen the car when braking and turning left together (entry), and tighten it when turning left and applying substantial power (exit). Conversely, stiffening the rights tightens entry and loosens exit.

Anything that reverses direction of suspension position change reverses effects of spring changes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2001

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANALYZING SHOCK TRACES

I am an aspiring race engineer – currently work for a professional sports car team – as well as a mechanical engineering student. I am learning to use data acquisition. I was hoping you could give me some insight into reading the shock velocity traces from the data system. Should they be compared to the dyno charts to make sure they are operating in the correct range? What can be deduced as far as oversteer/understeer characteristics?

This is a fairly complex subject, and one that has seen its share of hype and mystification, particularly in stock car racing. A few years ago, shock technology went through a "trick of the week" phase. Teams suddenly discovered that it was worth paying attention to, and shocks went from almost a non-factor to a perceived panacea practically overnight. Suddenly people were trying to fix every problem with shocks. This is a wrong approach. The first thing to remember about shocks is that they are only one element of the total package, not a holy grail.

The second thing to remember is that you need a rapid sampling rate to get shock traces that are accurate enough to be meaningful. Once you have a good shock position trace, it can look quite different depending on how you filter it. This is because shock velocities and accelerations undergo dramatic, continual change, especially on a bumpy surface. What sampling rate is good enough? Opinions vary, but anything under 100/sec is definitely too slow. 250/sec is a common general-purpose recommendation. 500/sec is fast.

Sampling rate makes more difference on a rough surface than on a smooth one. This is also true of filtering. Sampling rate also becomes more critical if you want to differentiate the position trace to obtain a velocity trace, and even more if you want to differentiate the velocity trace and look at acceleration. Inaccurate information can be worse than none at all.

The shock traces, taken alone, actually tell you more about the track and the rest of the setup than they do about the shocks themselves, at least in normal circumstances. The wheel motion will be reduced some as you stiffen up the damping, but ordinarily the track surface and the spring rates mainly determine how the wheels move. (An exception occurs when a shock is grossly undervalved or overvalved, or is leaking, sticking, or otherwise malfunctioning.) The shock valving determines the forces the shock generates when it goes through those motions, and those forces affect wheel loadings.

To know the force a shock is generating at a particular instant, you need to know the instantaneous velocity, which you get from on-car data, and the force the shock can be expected to generate at that velocity, which you get from shock dyno data.

Shocks are also acceleration-sensitive. The degree of acceleration sensitivity can vary widely. To get the best assessment of what forces the shock is producing at a particular point on the track, we need to reproduce both the velocity and the acceleration on the shock dyno. To do this requires a high-cost shock dyno that can be programmed to either produce a particular acceleration, or to play back motion recorded on-car. This means the dyno must be controlled fully electronically – usually with with high-powered hydraulics providing the actual force – rather than the more common variety with a scotch yoke drive that always produces a sinusoidal motion.

If we don’t have access to such a dyno, we can at least get some idea of whether we have a highly acceleration-sensitive shock by looking at a full-cycle trace from a sinusoidal dyno. In a plot where absolute velocity is the horizontal coordinate and force is the vertical coordinate, a full cycle trace will have two points or noses at the left (zero velocity) edge of the graph. These represent the “turn-around” points in the cycle – full compression, and full extension. If the two noses have dramatically different shapes, the shock is probably displaying acceleration sensitivity, particularly if all shocks of the same design show this pattern. If just one shock shows such a pattern, that suggests a malfunction on that individual unit.

Another way to get some indication of a shock’s acceleration sensitivity is to look at its construction. If the valving includes relatively massive chunks of metal that see big accelerations when the unsprung mass does, that’s a strong indication that the valving will be acceleration-sensitive. Examples include: dual-tube shocks with foot valves having coil-spring-loaded discs or spools, where the shock body is down or moves with the wheel; and gas shocks with coil-spring-loaded discs or spools on the piston, where the shock body is up or the piston moves with the wheel. Deflective-disc valving has minimal acceleration sensitivity.

Acceleration sensitivity is not necessarily a performance disadvantage. It may even be helpful in some instances. However, it does complicate the process of inferring shock forces from on-car traces, based on sinusoidal dyno testing.

Merely knowing whether the shock is highly acceleration-sensitive does not allow us to know actual forces in combinations of acceleration and velocity that our dyno can’t reproduce, but it does at least let us make informed guesses as to whether we can assume dyno data to be representative for a specific instant picked from an on-car trace.

Fortunately, we can accomplish a lot without knowing exact shock forces. Provided we know what portion of the valving is active at the instant in question, we can at least qualitatively predict what will happen to our wheel loadings if we soften or stiffen that portion of the valving. And from that, we can predict whether such a change will move the car toward oversteer or toward understeer (loosen it or tighten it) at that point in the lap.

To meaningfully discuss such predictions, we have to break track surfaces down into rough ones and smooth ones. Rough surfaces are ones where there are enough bumps so that the bumps cause most of the suspension movement. Smooth surfaces are ones that are smooth enough so that most suspension movement is caused by sprung mass motion. Shock velocities will be much greater, and will change much more, on a rough surface than on a smooth one.

On a rough surface, damping affects tire loads in two basic ways. First, the amount of damping affects the wheel's ability to ride the bumps with minimum load variation at the contact patch, and minimum non-contact (airborne) time. Second, the balance between compression (bump) damping and extension (rebound) damping affects the suspension's tendency to jack up or down when riding a series of bumps. Jacking down is by far the more common of these two possibilities.

Looking at the first of these issues, the suspension behaves better over bumps when lightly damped *except when the bumps excite the suspension at the unsprung mass natural frequency, or a simple multiple or fraction (harmonic) of that frequency*. In such instances, the suspension performs better if stiffly damped. Excitation at a vulnerable frequency is the worst-case situation, the scenario most likely to send the car out of control due to being upset by bumps. Therefore, there is a strong case for erring on the stiff side when in doubt; the car will be worse on "friendly-frequency" bumps, but will be less upset by the patches to which it is most vulnerable.

Looking at the second issue, if all four corners of the car jack down a bit, this may have little effect on cornering balance. However, if only one wheel jacks its suspension down, or if three do and one doesn't, or if two diagonally opposite wheels do and the other two don't, then the jacking will change the car's instantaneous diagonal percentage (percentage of tire loading on the outside front and inside rear, or right front and left rear for oval track). If instantaneous diagonal is increased, that moves the car toward understeer (tightens it); if instantaneous diagonal is decreased, that moves the car toward oversteer (loosens it).

On bumpy surfaces, medium to high speed valving is at work. On smooth surfaces, only low-speed valving is relevant. These terms are relative. On a stiffly suspended car, as in F1 or CART, "low speed" might mean below 1 in/sec (.025 m/sec). In stock cars or moderate downforce sports cars, "low speed" is commonly taken to mean below 2 in/sec (.05 m/sec). In passenger cars or off-road cars, "low speed" can be a lot higher. In any of these contexts, low speed damping means the velocity range that can be attained by driving the car violently on a smooth surface. Usually, low speed damping is governed by bleeds and preloads in the valving.

On a surface smooth enough to allow sprung mass motion to be the main input, we can control corner entry and exit oversteer/understeer properties with the low-speed damping. The basic rules for this are:

1. Whenever a damper's velocity is in the extension direction, stiffer extension (rebound) valving reduces load on that tire and the diagonally opposite one, and increases load on the other two tires.
2. Whenever a damper's velocity is in the compression direction, stiffer compression (bump) valving increases load on that wheel and the diagonally opposite one, and decreases load on the other two.
3. We get the greatest effect from changing the shocks that are have the highest velocity at the instant or point on the track that we seek to affect.
4. Effects on oversteer/understeer balance can be predicted by examining the effect on instantaneous diagonal percentage; more diagonal = tighter car.

We can write troubleshooting charts for various types of cars and tracks, and these can get quite long and complex. However, the four principles above can be used for all cars, and with data acquisition, you don't have to infer or guess at shock velocities.

So, returning to your question regarding whether we can deduce oversteer or understeer from shock velocity, the answer is that we cannot, but knowing shock velocity can help us predict changes to balance that will result from changes in valving. We determine the presence of understeer or oversteer by examining steering position, or a calculated speed-corrected steer channel. We also make sure we talk to the driver, because what really counts is whether the car is looser or tighter than the driver wants it. Also, a tight car can exhibit oversteer if the driver is purposely driving it loose to make it turn. Data acquisition is a tool to supplement human senses and brains, not a substitute for them.

As to what we compare to dyno traces or tabular charts, comparison to dyno data is done to infer forces, as noted above. We compare our dyno data to the manufacturer's if we have data from the manufacturer for the build spec we're using. We also compare our dyno traces to traces from shocks in our own inventory with identical build specs. We use these comparisons to make sure we don't have leakage or sticking, and to make sure we really built what we thought we built. It's surprisingly easy to include an extra disc, or leave one out, or grab the wrong diameter or thickness.

Finally, we compare dyno data for different shocks to evaluate the effect of a change to the build spec – to see what velocity range it affected and how much. We then compare this to the on-car velocity data to predict the handling effect of the build spec change, applying the four rules listed above.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SUSPENSION NATURAL FREQUENCIES

What suspension frequency should I be looking for in a formula car? Should the front and rear suspension have equal or similar frequency?

For readers less familiar with the subject, this question concerns the sprung mass natural frequency in ride, for the front and rear wheel pair systems. Natural frequency is a measure of how stiff the suspension is, which expresses stiffness in a way that is independent of the vehicle size or weight. That is, a given frequency feels comparably stiff whether the car is big and heavy or small and light, unlike a given spring rate or wheel rate. If the car is twice as heavy and the wheel rate is also twice as stiff, the natural frequency is the same.

A natural frequency is a vibration frequency at which a system will vibrate or oscillate when displaced from its static position, and released. For a simple system consisting of a mass supported by a spring:

$$f = .159\ (S/m)^{1/2}$$

where

f = natural frequency in Hertz (Hz), or oscillations per second.
S = spring rate, in Newtons/meter (N/m).
m = mass supported by the spring, in kilograms (Kg).

The quantity S/m is the inverse of m/S, which is the *static deflection* – the amount of spring compression at static condition. We may say that the natural frequency is inversely proportional to the square root of the static deflection. We may also say that the static deflection is inversely proportional to the spring rate, for a given mass. Therefore, natural frequency is proportional to the square root of spring rate, other factors held constant. So if we put springs with twice the rate on a car, the front and rear natural frequencies increase by a factor of the square root of 2.

The simple formula above is not really a very good approximation of an automobile suspension, especially a racing vehicle with stiff springs, stiff damping, and aerodynamic downforce. The tire is compliant, and its compliance is in series with the wheel rate. This makes the system softer, and the natural frequency lower, than if the tire were rigid. This is true for all cars, but it assumes particular significance when the suspension is stiff and the tire compliance is a large percentage of the total. If we know a spring rate for the tire, and the wheel rate from the suspension, the rate of the system is:

$1/S_c = 1/S_s + 1/S_t$, or $S_c = (S_s S_t) / (S_s + S_t)$
where
S_c = spring rate of the combination.
S_s = wheel rate of the suspension.
S_t = spring rate of the tire.

So if the tire is as compliant as the suspension, the frequency is only about .71 times what it would be if the tire were rigid.

Note that the above formulas assume a constant spring rate. In actual cars, we often have a rising wheel rate. This means that the frequency varies depending on suspension position. Some authors have suggested that natural frequency diminishes as downforce increases, because deflection increases. This is a misconception. Natural frequency is sensitive to the relationship between weight (mass times gravity) and effective wheel rate, and nothing else. This does not mean that rising-rate suspension is bad.

We have so far been talking about *undamped* natural frequency. Damping a system raises its natural frequency. To calculate a damped natural frequency, we need a constant spring rate, and a constant damping coefficient. In actual cars, damping coefficient is not constant, so any value we use for calculations has to be an approximation based on shock dyno outputs for velocities in the range we are trying to model. Stiff damping can raise frequency by as much as 30%, compared to the same system completely undamped.

As if this were not enough complexity, a suspension system will usually have different wheel rates when absorbing a one-wheel bump (combined roll and ride, or oppositional and synchronous motion) than when absorbing a two-wheel bump (pure ride or synchronous motion). The unsprung masses also have their own natural frequencies, much higher than the sprung mass frequencies, and these also vary depending on whether the wheels are moving synchronously or oppositionally.

In passenger cars, we have additional natural frequencies for masses flexibly attached to the main sprung mass, such as the engine on its mounts and the driver on a springy seat. These can be tuned to interfere with the ride motions of the sprung structure, and some useful ride damping can be achieved this way. Conversely, if these frequencies reinforce, ride quality will be adversely affected.

Passenger car engineers also look at “bounce” frequency (front and rear suspensions moving in the same direction, or basically in heave, approximating car behavior when both ends are excited together, as when traversing long humps and dips) and “pitch” frequency (car oscillating about a node near its middle, front and rear suspensions moving opposite directions, relevant to excitation by a short bump struck by first the front wheels and then the rears).

Despite all this, we can actually tell quite a bit about a car’s characteristics by just looking at the front and rear undamped frequencies, calculated using the front and rear portions of the sprung mass and the combined wheel rates for the front and rear wheel pairs. As previously noted, these will be inversely proportional to the square root of the static deflection for the front and rear. With asymmetrical cars, that would be an average of left and right static deflections, for the front and rear.

The equations given earlier are for frequency in Hertz, which is cycles or oscillations per second. It has also been customary, especially in countries using English units, to express suspension frequencies in oscillations or cycles per minute (opm or cpm). This figure will of course be 60 times the frequency in Hertz. It may also be calculated from static deflection in inches as follows:

$F = 188 / (x^{1/2})$
where
F = frequency in opm or cpm
x = static deflection in inches

Note that static deflection is calculated from the point where the spring is completely unloaded, which may be different from full droop on the suspension if the spring is loose or preloaded at full droop. When wheel rate is not constant, the only way to obtain valid static deflection for frequency calculation is to determine the instantaneous wheel rate at static ride height and divide by sprung weight.

Ranges of frequencies commonly found in different types of vehicles are:
- Very soft passenger car: 0.5 to 0.8 Hz (30 to 50 opm)
- More sporting passenger car: 1.0 to 1.3 Hz (60 to 80 opm)
- Modern sports car: 1.1 to 1.5 Hz (70 to 90 opm)
- Pavement race car with modest downforce: 1.5 to 2.0 Hz (90 to 120 opm)
- Modern race car with ample downforce and ground effect: 5.0 Hz (300 opm) or more

As to the relationship between front and rear frequencies, the traditional answer from passenger car engineering is that the front static deflection should be about 30% greater than the rear. This would mean that the ratio of front frequency to rear frequency would be around .88. Such a relationship makes the front end and the rear end roughly go up and down together on the first bounce after a short disturbance, while allowing the car to ride long disturbances with a minimum of pitch. It also allows for some passengers and luggage in the rear.

This approach presents problems in rear-engine cars, however, and even in front-engine cars with aerodynamics that are sensitive to front ride height. The main difficulty is that front ride height changes too much in braking and forward acceleration. With tail-heavy cars, a very stiff front anti-roll bar is usually required if front roll resistance is to be adequate to assure proper load transfer allocation with soft front springs. In braking and under power, the front end will rise and fall dramatically, and if we have a front valance skirt, a front splitter, or a front wing, aerodynamic properties will be too sensitive to braking, power application, and disturbances from the track surface.

Consequently, race cars often have higher natural frequencies in front than in back. Front to rear frequency ratios may reach 1.5 or greater. This actually represents a return to the frequency relationships of the early 1930's, when cars had beam axles and longitudinal leaf springs, anti-roll bars were uncommon, and the front springs were spaced closer than rears for reasons of steering clearance. With these characteristics, and 50% or greater static rear percentage, the car will oversteer unless the front springs are a good deal stiffer than the rears.

It turns out that having the front frequency a lot higher than the rear is a good second choice in terms of ride, if it is impractical to make the front a bit softer. If you make the front and rear frequencies similar, or the rear a little softer than the front, the car pitches excessively on short bumps, although it rides long disturbances very nicely.

All of this is more important in softly damped passenger cars, and softly sprung and damped dirt cars, than in firmly damped pavement race cars. With firm damping, the first and second oscillations after a disturbance are not much of a concern, as the damping suppresses them in any case. Consequently, many race engineers pay little attention to frequency relationships unless the car exhibits ride motions bad enough to hurt lap time or make the driver complain.

Choice of frequencies then comes down to factors other than ride motion. There is no perfect set of springs for all tracks, as shown by the fact that almost everybody uses different combinations depending on where they're running. The most fundamental tradeoff is between ability to ride bumps (softer is better) and ability to limit ground clearance changes, reduce CG height, and control camber (stiffer is better). The bumpier the track, the softer and higher we have to run the car. The smoother the track, the more sensitive the car is to ground effects, the more sensitive the ground effects are to pitch and roll, and the more downforce the car generates, the stiffer we have to make the springs. On ovals, steeper banking calls for stiffer springs. In general, we can use softer springs if the car provides roll or pitch resistance from geometry or interconnective springs such as anti-roll bars. So we are faced with a complex compromise, which cannot be reduced to a simple formula, or even a simple rule about frequency relationships.

Finally, in production cars, we may stiffen up the springs just to crutch bad geometry or limit the antics of leaf-sprung axles. As Colin Chapman reportedly once said, "Any suspension will work if you don't let it."

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WINSTON CUP HARNESS INSTALLATION

NASCAR has announced that they are now going to require that safety harnesses be installed according to the harness manufacturer's instructions. Harness manufacturers, without exception, recommend that lap belts go downward or slightly rearward to their anchor points, and that all belts have as short a run as possible, and not be bent around any sharp edges.

As a consequence of my now being a columnist for *Racecar Engineering*, I got to visit Richard Childress Racing as part of the annual UAW-GM media tour. (Actually, Paul Van Valkenburgh was the magazine's designated representative, but he had to go to Atlanta the last day of the tour, so I snagged his seat on the bus.)

The folks at Childress gave us a really nice meal and presentation, after which they formally unveiled their new paint schemes, apparently on new cars. Following that, they let all of us check out the cars up close, and interview the drivers and other personnel.

Late during this period I was looking at Kevin Harvick's car. I noticed that the belts were mounted just like they used to be in Dale Earnhardt's cars: lap belt mounts almost straight back from the seat holes; shoulder harness run over a bar at shoulder height and then another foot and a half or so down to an anchor point around kidney level.

It occurred to me that maybe I was looking at an older show car, used to display the paint scheme. Yet the car appeared new. The lower seatback mount, a square steel tube a bit above the floor, which held the seat in place and carried the lap belt mounts, was unpainted, and definitely appeared to be brand new.

Following Dale Earnhardt's death, the belt mounting issue assumed great significance, and this is of course the reason for the new rule. It was a point of public contention whether the Childress organization had been informed that there was a problem with their belt installation. I don't recall anybody saying that the installation was correct, or superior to the recommended geometry. The only controversy was whether the incorrect installation had been brought to the team's attention. Surely they can't claim ignorance now.

I will be watching the press for mention of this matter in upcoming Daytona coverage. I also invite anybody from RCR who may see my remarks to comment.

THINGS THAT MAKE SPRING CHANGES WORK BACKWARDS, REVISITED

In the August 2001 issue of this newsletter, I pointed out that on oval tracks there is a *critical angle* of track banking for a particular end of a particular car, at which the left side suspension no longer extends as the car rolls. At this banking angle, the car becomes insensitive to left spring changes. Beyond this banking angle, the effect of left spring changes reverses.

It has recently come to my attention that in beam axle suspensions, there will be a range of banking angles within which the left suspension will extend, as measured at the wheel – yet the left spring will compress! This is possible because the spring is inboard of the tire. It is possible that a *node*, or point of zero motion, may exist between the wheel and the spring.

For example, suppose the springs are exactly ½ the track width apart, and suppose that the suspension is in a condition of 1" per wheel of roll (about 2 degrees) and 3/4" of "squash" or ride compression due to track banking. The right wheel then has 1 3/4" of compression, and the left wheel has 1/4" extension. The left *spring*, however, has 1/4" *compression*. The right spring has 1 ¼" compression. A node, or point of zero displacement, exists midway between the left tire center plane and the left spring.

If this is the rear suspension, and the front suspension is an independent system in the same state of roll and ride, we have a situation where stiffening *either* left spring tightens the car (adds understeer). Fortunately, such conditions will mainly be encountered close to the critical angle, where sensitivity to left spring changes is fairly small in any case. It is possible, however, that substantial reverse sensitivity to left spring changes may be encountered when compression at the left rear wheel is small.

For those with good simulation software or data acquisition, questions of whether a spring is compressed or extended from static at a particular point in the lap are more easily resolved. For those of us who rely largely on our butts and our brains, qualitatively understanding physical principles is the only way to make sense of our observations. And even for those with electronic help, qualitative insight is part of making sense of the data.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

SOME BASIC SHOCK QUESTIONS

Mark, could you help me understand some things about shocks?

- *What is hysteresis?*
- *What is the difference between force and absolute velocity?*
- *What are they referring to when they talk about compression open, compression closed, rebound open, and rebound closed?*
- *Why does a shock cycle have 4 strokes instead of just one compression and one rebound stroke?*

Webster's defines hysteresis as "a retardation of the effect when forces acting upon a body are changed (as from viscous or internal friction)". Applied to a damper, this means the damping force or work.

Most of us have a pretty good idea what a force is. Webster's definition is "an agency or influence that if applied to a free body results chiefly in acceleration of the body and sometimes in elastic deformation and other effects". In the case of a shock, the force we measure is the force acting at the damper shaft. The force can act in two directions: extension or compression.

Shocks generate two kinds of forces: damping forces and gas spring forces. Gas spring forces are always in the extension direction. Damping forces are always opposite to the direction of motion.

Velocity, in physics and engineering, means the rate *and direction* of position change. With a shock, we have a simple case where motion is either in the extension or the compression direction. We can therefore express a velocity simply as a magnitude (inches or millimeters per second) with a positive or negative sign to denote direction.

The absolute value of a number is the greater of the number and its opposite (or additive inverse). The absolute value of 4 is 4; the absolute value of –4 is also 4. The absolute velocity of a shock shaft is the magnitude of its rate of position change, irrespective of the direction. If the velocity is 5 in/sec in extension, the absolute velocity is 5 in/sec. If the velocity is 5 in/sec in compression, that is also an absolute velocity of 5 in/sec. Absolute velocity is therefore just a fancy way of saying shaft speed.

The earliest shock dynos were purely mechanical devices (except for the electric motor) that generated a plot of force (on the vertical axis) versus position (horizontal axis). The resulting curve would be a loop. Modern dynos can also produce such a plot. Such a loop is sometimes called a *hysteresis loop*. The area enclosed by the loop corresponds to the mechanical work done upon the shock by the dyno during one cycle. This mechanical work in turn corresponds to the watt-hours of electricity used by the dyno to produce the movement, and the heat energy that warms up the shock as we work it.

Modern shock dynos can calculate velocity from position. With most dynos, the shock is worked through a two inch stroke at 100 cycles per minute. This gives a peak velocity at mid-stroke of just over 10 in/sec, and a velocity at any other point in the cycle that can be directly calculated by multiplying the peak velocity by the sine of the crank angle.

The dyno can therefore display force versus velocity, in a number of different formats. The most popular format is force (on the vertical axis) versus absolute velocity (on the horizontal axis). This means that force is displayed as a negative or positive value, but absolute velocity is always positive. The axes take the form of a letter T rotated 90 degrees to the left. The horizontal axis extends across the middle of the screen or page, with zero at the left. The vertical axis runs along the left edge of the screen or page, with zero in the middle.

Force during the extension (rebound) stroke is customarily displayed as negative; force during the compression (bump or jounce) stroke is positive. Note that damping force always acts in opposition to motion, so the compression damping force actually acts in the extension direction, and extension damping force actually acts in the compression direction.

One cycle of the dyno does have just one compression stroke and one extension stroke. But on the type of plot described above, this produces four traces. On the compression stroke, the absolute velocity starts at zero, builds to just over 10 in/sec at mid-stroke, then decreases to zero again at the end of the stroke. Commonly, the absolute velocity scale reads up to ten. Higher values are off the screen or page. So the trace for this stroke starts at the left, runs across the screen or page and off the right edge briefly (at mid-stroke), then comes back to the left and ends at zero. The height of the trace at any given point corresponds to the force at that point in the stroke. A similar process occurs during the extension stroke, generating two more traces, running mostly below the horizontal axis.

During the first half of the compression stroke, the velocity is increasing, so the valves are opening. This is referred to as the compression, opening (not open) phase of the cycle. The second half of the compression stroke is the compression, closing phase. Correspondingly, the first half of the extension stroke is the extension, opening phase, and the second half is the extension, closing phase. Each of these phases corresponds to one of the four traces on the graph.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE ON BELT MOUNTING

Regular readers will recall that last month I mentioned seeing improperly mounted lap and shoulder belts on a car displayed to the press at Richard Childress Racing. I have recently talked with Brian Butler of Butler Built seats, who custom-builds seats for RCR and supervises their installation. He tells me that all actual race cars he has worked on lately at RCR have the belts mounted correctly, and we agreed on what “correctly” meant. Brian was sure I had seen a show car that did not reflect actual current race car practice. I am relieved to hear that. I like being proven wrong on something like this.

RACING FRONT-DRIVE CARS

I have had a number of inquiries lately from people racing front-wheel-drive cars, asking about literature or information sources about setting up front-drive cars. It appears there is a distinct lack of literature currently in print on this subject, at least beyond manuals dealing with particular cars, and brief passages in general chassis books. So I’m going to offer some comments on the characteristics of front-drive cars in racing and high-performance applications.

The idea of racing a front-drive car is a bit like the idea of teaching an elephant to dance. It can be done, and people do it, but the basic anatomy of the critter involved is not particularly conducive to the activity. If we drive only two wheels of a car, they really should be the rears, for a number of reasons. The most obvious reason is that under forward acceleration, and when going uphill, tire loading transfers from the front wheels to the rears.

A slightly less obvious reason is that we have better control of the car when we can control the front wheels with the steering wheel and the rear wheels with the throttle. With front wheel drive, the only thing we can use to influence the rear wheels is the brakes. Some drivers become fairly adept at left foot braking and using the hand brake to slide the rear end of a front-drive car, but this is fundamentally awkward compared to throttle-steering. Active yaw control, which selectively applies individual brakes according to a computer’s interpretation of car behavior and driver intent, also offers some promise, but it is essentially a band-aid. For racing, any use of the brakes to control the car's balance or yaw behavior is definitely a last-resort approach (except perhaps when trail-braking on entry), because it works the tires against each other and dissipates speed. Even the best electronics are better added to a chassis with good fundamental dynamics.

For these reasons, nobody builds race cars with front wheel drive nowadays, if they have free choice of layout. However, front wheel drive shows up on the race track in production classes, and production-based classes, because it is popular for passenger cars.

Let's examine the reasons for this popularity.

Front wheel drive, with front engine, is one of four possibilities in a four wheeled, two-wheel-drive vehicle, the others being rear engine/rear drive, front engine/rear drive, and rear engine/front drive. This last option is never used, although Buckminster Fuller designed a rear-engine/front-drive car (with rear steering!) in the 1930's. One example was actually built. It ran, but its handling properties did not attract imitators.

Front engine/front drive and rear engine/rear drive are commonly referred to as engine-over-drive-wheels layouts. Both layouts concentrate the entire powertrain at one end of the car. This saves weight and space. It also puts well over 50% of the car's weight on the drive wheels – usually somewhere between 55 and 67 percent. This is good for propulsive traction. However, it is at best a mixed blessing in terms of cornering behavior. Due to *tire load sensitivity* – the decrease in coefficient of friction as loading increases, which we have discussed in previous issues – nose-heavy cars tend to understeer in steady-state cornering, and tail-heavy ones tend to oversteer. The car tries to leave the desired path heavy-end-first.

If we are faced with a choice between heavy understeer and heavy oversteer, understeer is clearly the safer choice. This is one reason why rear-engine/rear-drive layouts have fallen from favor for passenger cars. Another reason is the simple fact that luggage or cargo space in the rear of a vehicle is more useful than luggage or cargo space in the nose, because the space can be filled or over-filled, and the lid left partly open, for short hauls with big loads. The rear seat can be made to fold to accommodate long objects.

Additionally, a nose-heavy car has better directional stability in crosswinds than a tail-heavy one, other things being equal. This is highly significant for a light sedan that spends much of its life on freeways. Tail-heavy cars can be made adequately stable in crosswinds, but this requires careful attention to aerodynamics, and places added constraints on styling.

These practical and safety-related concerns have driven the trend to front-engine/front-drive cars. I personally think that the rear engine/rear drive option has an unrecognized future, for cars with back seats, intended for drivers who place priority on performance. Existing front-drive powertrains, especially the larger V6 and V8 variety, could be adapted to such cars, placing the engine slightly ahead of the rear axle line, with a relatively long wheelbase. In other words, we are envisioning a stretched, larger-engined version of the Toyota MR2/Pontiac Fiero/Fiat X1/9 concept. However, such a layout would not match the practicality of a front-engine car for hauling bulky loads, and would probably have somewhat more interior noise. So front-engine/front-drive cars are here to stay, and will surely command a large share of the passenger car market for the foreseeable future.

I mentioned the tendency for a car to try to leave the desired path heavy-end-first. We have two principal tools we can use to control this tendency: use bigger and/or stickier tires at the heavy end of the car; and/or put most of the roll resistance at the light end. We can also play with secondary factors such as toe, camber, and tire inflation. Finally, we can simply reduce the nose-heaviness or tail-heaviness, by moving the engine toward the center of the car, or by moving heavy components to the light end as much as possible.

In rear-engine/rear-drive cars, it is common nowadays to use larger tires in back. The reliability of modern tires, the increased availability of road service, and the advent of space-saver spare tires have paved the way for this trend. When front-engine/front-drive cars pushed rear-engine/rear-drive cars out of the passenger car market, most manufacturers considered it essential to have a full-size spare that could be used at any corner of the car. This is still a practical advantage, but not the necessity that it once was.

It is also common in rear-engine cars to mitigate some of the tail-heaviness by using a mid-engine layout. Even if the tail-heaviness is modest, drive traction will be quite good, thanks to the rearward load transfer under power.

With front drive, this load transfer works against us. Consequently, we are faced with a dilemma: maximize front-heaviness so we can put power down, or minimize front-heaviness so we can corner. There is no way to achieve one objective without compromising the other. This is also true with tail-heaviness in rear-drive cars, but the compromise is less excruciating thanks to the help we get from rearward load transfer.

If we were designing for an imaginary set of rules that required us to use front wheel drive, but allowed us ample freedom otherwise, we might make the car extremely nose-heavy, use big tires in front and smaller ones in back, and be sure to provide power steering and huge front brakes. We would also make the wheelbase really long. This would be a funny-looking car, and less enjoyable to drive than a rear-drive racing car, but that would be the way to go fastest with front drive.

In real-world classes where front-drive cars compete, we are usually constrained by tire rules and limitations on modifying stock body configurations. Production front-drive cars invariably use the same size tires at both ends – partly for practicality, partly to provide for the occasional heavy load in the rear. Road-racing and oval-track front-drive cars consequently use equal-size tires all around, although big fronts and little rears are seen in drag racing.

We are also usually required to keep most of the stock body/frame structure and suspension, and prohibited from moving the engine. Our control of front/rear weight distribution is then limited to moving minor components, and placing ballast if we run any. To the extent that we can choose our CG location, the principles we want to follow with front wheel drive are these:

- If the track has high-speed turns; if a large portion of the lap is spent cornering; if grip is ample; if power is modest – try to move weight rearward.
- If the track has slow turns followed by straightaways; if a small portion of the lap is spent cornering; if grip is modest; if power is ample – try to move weight forward.
- For drag racing, standing starts, or hill climbing, try to move weight forward.
- If braking is especially important, try to move weight rearward.
- In all cases, try to place weight as low in the car as possible.

Regarding suspension setup, we are forced to work around the fact that the front wheels limit the car. If the car were not nose-heavy, it might make sense to give the front and rear suspension systems similar roll resistance, and try to work all four tires. A front-drive car done this way (if it were possible, which would only occur if we had lots of ballast to work with) would have very poor forward bite. Since a front-drive car is necessarily nose-heavy, it must be set up to work the *front* tires as evenly as possible. That means it must corner with the inside rear tire very lightly loaded or airborne. We trade away lateral grip at the rear to gain more at the front, where we need it.

We also gain drive traction on the inside front wheel. This is important in a front-drive car, because we cannot use limited-slip differentials that lock too firmly or abruptly, unless the driver has great tolerance for steering fight.

It is important to note that once the inside rear wheel is airborne, the rear suspension has contributed all the anti-roll moment it can, and any further roll resistance has to come from the front. Up to the point of rear wheel lift, rear load transfer builds faster than front load transfer. Beyond that point, rear load transfer is 100%, and front load transfer builds rapidly. So does roll angle. So does understeer.

As a general rule, to get a car that has good consistency as grip varies, we want the inside rear wheel to lift just a little in steady-state cornering, when grip is good. If it lifts more than that, we are likely to have a relatively loose car when grip is poor and a much tighter car when grip is good.

Many front-drive cars use MacPherson strut front suspensions. Most of these suspensions, especially when lowered for racing, have camber properties that produce little camber change in ride and substantial camber change in roll. This means we can improve the cornering camber on those overworked front tires by providing ample wheel rates in roll. On the other hand, allowing soft action in ride will not compromise camber control very much at all. This argues for fairly stiff anti-roll bars, even at the front, and relatively soft springs. That is, the front needs to be stiff in roll, and the rear needs to be stiffer yet, by a sufficient amount to make the inside rear wheel lift just a little when grip is good. In the real world, available suspension travel and rules regarding anti-roll bars and their mounting may constrain this approach, but the idea is to get the desired roll stiffness distribution, and do it with bars as much as possible.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

QUADS ON PAVEMENT

I am a student at the University of Ulster in Northern Ireland studying Mechanical Engineering. I am currently working on a project which involves designing a suspension for a purpose-built road going quad [quadricycle].

I am sure you are familiar with quads. The quad I am working on is similar to existing quads. However, it will be approved to go on the road, and it will handle better on the road than current quads, which are designed exclusively for off-road conditions.

Current sports quads have a single swing-arm at the rear which connects to a solid rear axle. This has long travel, and a relatively soft suspension. Re-bound is slow on these quads. At the front, a double wishbone set-up is used on each side. A similar long shock is used, however re-bound is slightly faster. The suspension set-up is mainly for high jumps and being used on rough terrain.

When this set-up is used on the road it produces a very soft ride. There is too much pitch during retardation and acceleration, and sufficient roll to lift the inside rear tyre, almost without fail, when there is even a slight change in direction. This wheel lifting problem is also attributed to the solid rear axle and single swing-arm set-up.

The quad I am developing has no differential, and is driven by the rear wheels. This can make acceleration from tight corners at slow speed very difficult. You can't put the power on until you are in a straight line. Can suspension set-up affect this? Often, the front wheels lift completely when the power is being applied. This results in massive understeer.

What would you suggest to be a cheap and effective method to design a new suspension set-up for a quad bike for the road? It should solve the following problems:

1) High levels of pitch and roll on hard surfaces.
2) Inside rear wheel lifting during cornering.
3) Understeer when applying power.

I don't know what laws you have in Northern Ireland that would allow licensing a quad for the street, but you definitely couldn't do that over here. Anything with an engine and four wheels is considered a car in the US, and has to comply with windshield, side impact, and other standards that a quad has no hope of meeting.

It would be legal to race quads on pavement almost anywhere in the world. As far as I know nobody has tried it, probably due to the vehicles' poor handling properties on pavement, which you have observed.

The ride and handling problems you describe are largely inherent in the layout of the vehicle. To get away from them, you would need a rear suspension with a wheel rate in roll comparable to the front, a differential, and a much lower CG and/or longer wheelbase and wider track. You could conceivably design such a vehicle, and still have the rider sit astride it instead of in it, but there would be safety and aerodynamic penalties, with the only advantage being that using body english would be somewhat easier. For good results, a pavement quad would have to look dramatically different from a dirt one. You can't just doctor the links and springs.

What you can do with the links and springs, keeping the existing layout, is make the front suspension very stiff in roll like the rear, and stiffen the wheel rate in pitch. If you do this with interconnective springing (anti-roll bars, anti-pitch bars, diagonal bars, or equivalent devices), it is still possible to keep the wheel rate in heave fairly soft. You would then essentially have a suspended go-kart with a motorcycle-style operator position. This will give you a vehicle that will bicycle, or flip, instead of tricycle, and will pitch less on its suspension due to longitudinal accelerations. However, there will be more sprung mass pitch due to bumps. The tendency to overturn -- laterally in cornering, rearward in forward acceleration, forward in rearward acceleration -- can only be controlled by the suspension up to the point of wheel lift. Beyond that, it's purely a matter of where the center of gravity resides relative to the four contact patches. To change that, you have to fundamentally redesign the vehicle.

The whole logic of the quad is that it gives you a very short, narrow package, allowing the machine to operate on many trails that would otherwise require a motorcycle or a mule. If you make the track and wheelbase sufficient to work well on pavement, you throw away the main advantage of the motorcycle riding position.

I have mentally toyed with an idea that poses some similar problems. I am a bicyclist, and I have noticed that the world could use a human-powered on-road vehicle that will work decently in snowy, icy conditions, and also on dry pavement. Such a vehicle would probably need four wheels, and would have to be narrow enough to allow cars to pass, and preferably narrow enough to fit through a door and be brought in the house like a bicycle. The realities of operation in snowy conditions would require a conventional riding position; a recumbent would expose the rider to cold, dirty baths of snow, ice, water, and slush thrown by passing vehicles, and would be too hard for motorists to see. The rider would have to sit slightly *higher* than on a bike, for reasons of pedal-to-ground clearance. The rider's legs must be straight enough to give good pedaling efficiency. So this would be a really tall, narrow device, with a really high CG -- even worse than what you're contemplating. Pitch stability would not be a big problem, due to the modest power, but roll stability would be a major issue.

In the UK, some people build high-tech pedal-driven tricycles, and even race them. These vehicles are almost never seen on my side of the Atlantic. They obviously have the same stability problem that has caused quads to replace trikes in the ATV market, only even worse. The riders coast through turns with a knee hooked over the top tube, hanging off the inside of the trike like a sidecar monkey -- and cornering speed is still limited mainly by overturning rather than grip.

I think the best approach to a tall, narrow vehicle with more than two wheels would be to give up on trying to make the vehicle corner flat like a car, and instead make the suspension extremely *soft* in roll. The rider would then lean the vehicle into the turns like a two-wheeler. The suspension would resist this, but only gently, up to perhaps 20 degrees of lean. The rider would also have to hang off the inside of the vehicle to corner hard on pavement. One thing that works in our favor with human power is that the rider is the majority of the mass, making the CG highly mobile.

I think you could build a self-propelled pavement quad that cornered like that. You'd need motorcycle tires if the wheels leaned with the vehicle, as they would with independent suspension. If you used beam axle suspension instead, that would call for shaft drive. You could then use quad or car tires. To my knowledge, nobody makes quad tires for pavement, so I expect you'd want to adapt car tires. Quad tires I've seen are not only unsuitable for pavement in terms of tread design, they are very likely incapable of coping with the speeds and temperatures they would see on pavement, even if you retreaded them.

The vehicle would dramatically outweigh the rider, so it would still be desirable to have a fairly wide track and put all masses as low as possible. Even at that, I would worry about having the vehicle "high-side" or flip toward the outside of the turn, as motorcycles sometimes do. In that case the rider would either be thrown off, or crushed by the vehicle.

A couple of years back, Daimler-Chrysler showed a three-wheeled vehicle that leaned into the turns. As I recall from published reports, this was accomplished via semi-active suspension on the two front wheels. If I understand correctly, the control system was a simple hydraulic valve operated by the steering; it leaned the way you turned the steering wheel, and the more you steered the more it leaned. I could be wrong about that, because such a control system would pose a problem if you needed to countersteer to correct oversteer. The vehicle would lean the wrong way, and immediately high-side. Also, the lean angle could not be optimized for both high and low vehicle speeds at a given steer angle. To get around these shortcomings, lean must be controlled by something other than steering position.

The way to make the vehicle lean itself would be to control the tilt electro-hydraulically from an acceleration sensor. Since the sensor would lean with the frame, the system could simply add tilt until the sensor no longer detected lateral force, or until active suspension travel was exhausted. For smaller vehicles, especially where we are trying to minimize cost and complexity, it makes more sense to let the operator sit astride the vehicle, and lean it with body english.

The power push you describe is a universal problem in all vehicles without differentials, except those that only turn one way and can use tire stagger. It is also a universal problem with vehicles that are easily capable of wheelstands, at available grip level. You have both of these problems at once. Therefore, you have to solve both of them to obtain satisfactory performance.

Quaife makes limited-slip differentials for chain drive, which are popular for motorcycle-engined cars. These are usually used with independent rear suspension.

Go-karts race without differentials, but only because the rules demand this. One strategy for getting the karts to turn sharp corners on pavement is to deliberately try to make them lift the inside rear. Since all suspension is prohibited, tuners can't achieve this with soft front wheel rates in roll, so they do it with large scrub radii (long front spindles) and lots of caster. When you have no differential, lifting that wheel isn't necessarily a vice. You can still put power down. It doesn't mean the vehicle is about to bicycle, although I do think a quad on pavement would bicycle pretty easily. It's actually a form of warning, and therefore arguably a safety feature. If neither inside wheel lifts until they both do, you have less warning of overturning. Of course, when you get on the power hard, and all the load is on the rear tires, the inside rear plants and you get the push. So for a quad on pavement, I think roll-compliant rear suspension and a differential would probably be the way to go.

The tendency to wheelstand depends on the ratio of CG height to the longitudinal distance from CG to rear axle. Suspension has little influence, except that upward or downward jacking can make the CG rise or fall a bit under power. Having ample load on the rear wheels is basically a good thing, but we also have to steer. Optimal balance between these two concerns can only be achieved for a fairly narrow range of forward acceleration. A vehicle with a high CG and short wheelbase is well suited to low-grip surfaces, but is wheelstand-limited when grip is good. When the vehicle is cornering hard and trying to gain speed, power understeer sets in before the front wheels will actually lift. This is a characteristic of dirt vehicles operated on pavement. Sprint cars on pavement have the same problem as your quad. In sprint cars, it helps to lengthen the wheelbase, adding most of the length between the engine and the rear axle. This results in less static rear percentage, in addition to a smaller ratio of CG height to CG-to-rear-axle distance. In moderation, the added front percentage doesn’t hurt the car, because sprint cars are so tail-heavy to begin with. To maintain static rear percentage while suppressing wheelstanding and power understeer, we would have to lengthen the front of the car as well.

The way to put power down best, over the widest range of grip levels, is to use generous static rear percentage, a low CG, and a long wheelbase -- think dragster. I don't mean that a pavement quad, or a pavement oval track or road racing car, needs a 300 inch wheelbase, but longer and lower is the direction you need to go when adapting dirt-optimized vehicles to pavement.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

NATURAL FREQUENCIES REVISITED

My January 2002 newsletter contained an error. I stated that damping a suspension system raises its natural frequency. Actually, damping a system reduces natural frequency. At damping coefficients exceeding critical damping (the amount of damping that makes the system return to static position in least possible time without overshoot, following displacement from static), the system is considered non-oscillatory. This means its damped natural frequency is undefined: it doesn't have one.

Usually, even relatively stiffly damped race cars are below critical damping, so they do have a damped natural frequency. This is given by the equation

$$\omega_d = \omega_n \sqrt{1 - \xi^2}$$

where:

ω_n	-	undamped natural frequency
ω_d	-	damped natural frequency
c	-	actual damping coefficient
$c_c = 2m\omega_n$	-	critical damping coefficient
$\xi = c/ c_c$	-	damping ratio
m	-	sprung mass

We can see that as the damping ratio approaches 1, the quantity under the radical approaches zero. At critical damping, the quantity under the radical is zero, and so is the damped natural frequency. If the damping ratio exceeds 1, the quantity under the radical becomes negative, and the square root of a negative number is undefined for real numbers; hence the damped natural frequency becomes an undefined quantity.

Stiffening the damping and stiffening the springing have qualitatively similar effects on amplitude of suspension motion, but opposite effects on frequency. Stiffening damping or springing reduces amplitude: the suspension moves a shorter distance in a given situation. But stiffening springing increases frequency, whereas stiffening damping reduces frequency.

My thanks to Professor Jorge Pinto Pereira of the Escola Superior de Tecnologia de Setubal, Portugal for advancing my education on this.

QUADS ON PAVEMENT REVISITED

The April 2002 newsletter addressed the dynamics of quads (4-wheel ATV's) on pavement. Newsletter recipient Chris Petersen reports meeting a person who has raced quads on pavement, in the American midwest (Missouri or Iowa, he thought), on road courses, with pavement tires. This person told Chris that there is now a small local racing series for these machines.

This doesn't sound safe to me, but it sounds like it would be entertaining to watch – in somewhat the same way as sidehack racing, or gladiatorial combat. I'd like to hear from anybody with further information.

SPACE FRAME OR MONOCOQUE?

For cars like a D Sports Racer (DSR), what kind of chassis stiffness would you recommend? For a home builder of DSR cars, what type of chassis construction (spaceframe, monocoque, or a hybrid) would be most appropriate? I am contemplating building a DSR at home – but it might be easier to have a car builder fabricate the actual chassis and suspension components since I am lacking in the welding department.

> I haven't been to an event where DSR cars have run for many years. I know from reading that some of the cars have become very professional, and some builders spend quite a lot on them. I believe it is still true that many people in this class build their cars rather than buy them, though.
>
> There is no such thing as too much chassis stiffness, as a general rule. You can have too much weight. You can have too little access to the engine and other components. You can have too little room for the driver. You can end up with these problems in the pursuit of stiffness. But stiffness is never a disadvantage in itself.
>
> You not only want a maximum of torsional stiffness as measured by applying a torsional load at the coilover or rocker pickup points, you also want good local stiffness at all the suspension and steering pickup points, and good rigidity in the suspension bits themselves.
>
> I really have no idea what torsional stiffnesses the best DSR's are attaining these days, but off the top of my head I would think 5000 lb ft / deg would be a reasonable target.
>
> Over the 40 years since the Lotus 25 demonstrated the superiority of monocoque construction by out-performing the space-framed but otherwise identical Lotus 24, a consistent pattern has

emerged. Monocoques beat space frames wherever they're allowed. Space frames prevail where the rules require them, or in some cases where the builders are wedded to tubular frames and monocoques are informally excluded through economics and culture.

The fact that space frames are still widely used, and that some people like them well enough to write rules requiring them, tells us that they must have some attractive features – and they do.

Advantages of tubular steel space frames include cost, repairability, improved access to components, and availability of instruction in the skills needed for construction and repair. There are a few places you can go to learn to work with carbon fiber or sheet aluminum monocoque structures, but it's much easier to find a welding class. It's also much easier to find a welder than than a composite technician, if you want to hire help.

My next door neighbor builds sprint car chassis. He does a lot of the hands-on work himself, including machining and fitting parts, but he hires others to do the welds. They do the welding in his shop, with his equipment. He maintains control of design, often without drawings of any kind, and the welds are professional quality. He currently has one welder who works day hours, more or less full-time, and another who fabricates for a Winston Cup team and moonlights in the sprint car shop in the evenings.

Partly, your decision will depend on precisely what cars you have to beat, and how much you really care about winning. I'm sure you understand that you are contemplating a hobby here, and a learning experience, rather than a career – at least in the short term, with this project. You can certainly have fun and gain worthwhile experience building a tube-frame car.

It's pretty certain that monocque construction is the path to best performance from the firewall bulkhead forward. For the rear of the structure, much will depend on your choice of powertrain. For some time, a Kohler was the engine to have. I think I heard a few years back that some guy in Wisconsin had built his own engine specifically for DSR and was having success with it. That is definitely not an approach for the faint of heart or the thin of wallet. Over the years, people have used an incredible variety of powerplants for DSR cars, including single and multiple motorcycle and snowmobile engines, Mercury outboard engines, SAAB two-strokes, rotaries, Fiat-Abarths, you name it.

Checking the rule book (mine is from 1999 and may not be 100% current) I see that SCCA is still encouraging diversity. DSR has five separate displacement limits: for two-strokes, anything-goes atmospheric four-strokes, two-valve-per-cylinder four-strokes, automotive-based four-strokes, and rotaries. There are two minimum weights – a lighter one for chain or belt drive cars. There is no limit on tire size – only a minimum wheel diameter of 10 inches. Obviously your first step will be to think through all these possibilities. You may find that different options give an edge on different tracks, so your choice may vary depending on where you plan to compete the most.

Ordinarily, the rear portion of a monocoque sports-racing chassis consists of two box structures running alongside the lower portions of the engine. These need to be as large in cross section as possible, to maximize rigidity. This inevitably conflicts with access to the sides of the engine. If you are using an automotive-based engine, mounted longitudinally, a monocoque can work pretty well. If you are using a motorcycle engine, and you need to be able to get the side covers off without pulling the engine, a space frame design starts to look more attractive.

So you are faced with a complex set of interdependent decisions. The decision of type of frame structure is one of them, and they all lean on each other. You need to come up with an integrated package that satisfies your personal design objectives and makes sense taken as a whole. The people who wrote your rules intended to give you an interesting puzzle. Have fun with it.

TIRE CARE

It is sometimes said that lap time is 50% tires and 50% everything else. There can be no doubt that those four little patches of rubber are vital to control of the car. Yet it is surprising how many racers ignore the advantages to be had from tire care.

Tires vary in their sensitivity to age and care. As a rule, real racing tires are more sensitive than street tires. Yet all tires will benefit to some degree from proper care.

All tires contain solvents. Over time, these evaporate, and this contributes to hardening of the rubber with age. One way to slow this process down is to store your tires in the heaviest plastic garbage bags you can find, and tie or twist-tie the bags. Plastic bags are somewhat porous, and they tend to get torn when used to store tires or wheel/tire assemblies, so it's not a bad idea to double-bag.

Tread compounds also harden due to polymerization. Heat cycling speeds this process dramatically. Effect of heat cycling is related not only to how extreme the temperatures get, but also how fast the tire is heated and cooled. Consequently, it helps to warm tires as gently as possible when you first go out on the track, and cool them as gently as possible at the end of a run. It helps to store tires at the coolest and most stable temperature possible.

Contrary to popular belief, moist air does not build more pressure per degree of temperature rise than dry air or nitrogen, *provided all water is in the vapor state*. Water only causes disproportionate pressure rise if it is in the liquid state when pressure is set. That said, it is surprisingly easy to have liquid water in a tire. Chief sources of this are mounting lubricants, condensate in air tanks and hoses, and condensation in the tire itself if pressure is set at low temperature and the tire was inflated in muggy weather. Inflation gas can even absorb atmospheric moisture through the tire rubber. So there is a point of diminishing returns on keeping moisture out of tires, but it's hard to know when we've reached it. Therefore, it makes sense to err on the side of dryness. This means using either dried air or dry nitrogen for inflation, and purging after mounting and before racing.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

OILY TIRES

I recently blew an engine and got oil all over a nearly new set of racing tires. Are they junk? Is there a way to clean them?

It may surprise you to learn that many people put oil and related substances on tires deliberately, to soften them. Substances used for this include automatic transmission fluid, diesel fuel, paint thinner, and WD-40.

Does it work? The answer, in Pogo's words, is "an unequivocable maybe". There is no doubt that the rubber will get softer. Whether this will actually improve grip is harder to predict. There is a difference between soft and sticky and soft and greasy. In some cases you really do get a grip improvement. In some cases you don't. The only way to really know is to use a test day to try out a particular tire/solvent combination.

When you put oil on a tire, the rubber starts absorbing the oil immediately. But it takes a while to absorb a lot, especially deeper into the rubber. So if you get oil on a rubber item that you don't want to soften, you can usually avoid any major impact on longevity if you wash the item off promptly with detergent. You can even just wipe it off thoroughly with a paper towel, depending on the importance of the item and the situation. Wiping first, and then washing, is the most thorough. If you're on the road, a car wash is generally the venue of choice. Use the foam brush, then rinse.

If, on the other hand, you're trying to soften the tread compound of a tire, you try to maximize absorption instead. Some people use electrically driven tire rotisseries to apply softening agents to tires. If the softener contains volatile compounds (ones that tend to evaporate), it helps to apply the liquid to the tread, then bag the tire for days or weeks.

Are such procedures safe? Again, maybe. Certainly many racers use tread softeners and usually don't cause the tires to come apart. Then again, there are no guarantees, and inevitably tire manufacturers discourage this sort of tampering. Once a tire has failed, it is uncommon for any systematic investigation to be conducted to determine what may have caused the failure, particularly in small-time racing. Even if a tire does fail structurally, and even if it can be determined that it was softened, there is usually room for doubt as to whether the softening really was the cause, as opposed to inflation pressure, aggressive driving, suspension settings, or track conditions.

Finally, there is the question of legality. Rules on tire soaking, and enforcement of those rules, vary widely. This is between the racer and the officials. All I can say is investigate before you experiment, and proceed at your own risk. Most petrochemicals do have an odor and are detectable with a sniffer.

One other thought, for situations where an engine lets go and oils down most of the car. Try to think of all the other rubber items that may have gotten oiled: silicone insulation, motor mounts, suspension bushings, vibration isolators for the radiator or electrical components . Usually we don't want to soften any of those pieces. Any foam rubber will tend to be especially susceptible to damage from oil, and especially hard to rid of oil.

SOME BASIC OVAL TRACK CONCEPTS FOR ROAD RACERS

Could you please clarify what you mean by "diagonal percentage" and "stagger"?

"Diagonal percentage" and "stagger" are oval track terms. However, the concepts are useful in road racing as well. Diagonal percentage is the combined loading of the outside front and inside rear tires (RF and LR for left turns, or if no turn direction is stated or implied), expressed as a percentage of the total for all four wheels. We can speak of *static* diagonal percentage, which is diagonal percentage as measured on wheel scales when setting up the car. We can also speak of *dynamic* or *instantaneous running* diagonal percentage, which is the percentage of total wheel loading on the outside front and inside rear, at a particular instant while running.

As a general rule, more diagonal tightens the car (adds understeer).

Spring rates do not affect static diagonal, since we can set that wherever we want with any spring combination. Springs do affect the way dynamic diagonal varies as the car runs.

Stagger is generally understood to mean the difference in tire circumference between the right and left tires at one end of the car. A car can have tire stagger at the front and/or rear, but if not otherwise stated, we mean the rear. Many oval track cars use locked, or spool, rear ends, so rear stagger has a big effect on cornering balance. With significant amounts of stagger on an axle that forces both wheels to run at identical rpm, the smaller tire drags. This tends to rotate the car toward that side. At the front end, or at the rear when the wheels can turn at different speeds, stagger does not cause the smaller tire to drag, but it does affect left/right brake bias, and also left/right propulsion force bias with an open differential. If brake or drive torques are equal, on two unequal-size tires, the smaller tire produces a greater rearward or forward force.

Ordinarily, we think of road racing cars as symmetrical. We tend to assume that if the car has to turn both ways, we want close to 50% diagonal, and no stagger. However, that isn't necessarily strictly true.

Many road courses are predominantly right-turn tracks. A few are predominantly left-turn. It is not uncommon for the car to spend more than three times more seconds per lap cornering one way than it spends cornering the other way. In such cases, we can often gain lap time by optimizing the car for the dominant turn direction. We want good handling balance in both right and left turns, but it may pay to sacrifice some speed turning one way to gain some speed turning the other way.

This most commonly involves moving ballast to the inside for the dominant turn direction. When we do that to a significant degree, in a nose-heavy or tail-heavy car, a funny thing happens. 50% diagonal doesn't give us equal left percentage at the front and rear. If we want similar cornering balance in both right and left turns, a good starting point is to have equal left percentages for the front and rear wheel pairs, rather than 50% diagonal. The diagonal percentage will not be really far from 50%, but there will be some difference. For a car with 60% rear and 55% left, for instance, we have 55% left at both front and rear when the diagonal is 51%.

Another reason to depart from 50% static diagonal in a road racing car is driveshaft torque in a live-axle rear suspension. Driveshaft torque acts through the suspension with a live axle. With conventional engine rotation, it rolls the car rightward, and unloads the right rear and left front tires. To compensate for this, we may want to run less than 50% static diagonal.

Ordinarily, we do not deliberately use tire stagger in road racing, but we can if we want to. We can also have tire stagger inadvertently. Oval track racers often use manufacturing variation in tire diameter to tune stagger. Road racers often do not even measure tire size to see if they have stagger. Even if your desired stagger is zero, the only way to be sure you have that is to measure. This is particularly important if you race on bias-ply tires, which in general are more prone to diameter variation than radials, both when new and due to change after they've run. We not only need to measure our tires to avoid inadvertently having undesired stagger, we also need to understand what stagger does to car behavior so we can recognize the clues that we might have undesired stagger.

It is also worth noting that there is a difference between stagger as usually measured (difference in unloaded circumference) and *effective* stagger, which is the difference in *effective circumference*, or 2π times *effective radius*. Effective radius is neither unloaded radius nor loaded radius (distance from hub center to ground, under load). It is approximately the unloaded radius minus 1/3 of the deflection under load (difference between unloaded and loaded radii). One important consequence of this is that we can achieve considerable effective stagger with radial tires through pressure variation, even though pressure changes have little effect on unloaded circumference.

NARROWER REAR TRACK THAN FRONT

Why do rear-engine and mid-engine cars so often have narrower track widths at the rear than at the front?

The reasons for this are more practical than theoretical. In racing, we are usually designing to an overall width limit, rather than a track width limit. Likewise, for road use overall width is realistically the most important constraint. If the car is tail-heavy, we will often use wider wheels and tires at the rear than at the front. Track width is measured between the centerplanes of the right and left tires. So it’s mathematically inevitable that we will end up with the rear track narrower than the front if the overall width is the same front and rear, and the rear tires are wider than the fronts.

There is also a rational case to be made for having the overall width slightly greater at the front than at the rear, simply to place the widest portion of the car in the driver’s field of view, and discourage the tendency to hit curbs or other obstructions with the inside rear wheel in tight turns.

There is one more reason to make the rear track narrow, which applies more often to road cars than race cars, but may apply to race cars as well if the car is power-limited and has full-width bodywork. For least aerodynamic drag, it is best to have the car taper toward the rear. Therefore, when top speed and/or fuel economy are top-priority design objectives, it is common to see a narrower rear track than one would otherwise expect.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

BRAKING/DOWNSHIFTING TECHNIQUE

I'm track-engineer on a Porsche 996 Biturbo. When I talk with our drivers, and also drivers of other teams and cars, and we are talking about braking, the drivers are always telling me that they brake on the engine, by downshifting and releasing the clutch. I ask them why they are doing it that way. They explain that they can brake harder with the engine and brakes together than with the brakes alone.

I disagree with them on this. In my opinion, the maximum braking force of a car is determined by the coefficient of friction between the tires and the track surface. How this force is generated doesn't matter. It can be generated by the brakes or the engine. But if the brakes are big enough, they can lock all four wheels by themselves, or get the ABS working. So why lose time downshifting, when the shortest braking distance possible can be achieved using the brakes alone?

Braking technique is controversial, among drivers, driving instructors, and engineers. There is no single best way for all cases. However, there are factors that will logically cause us to favor one approach or another, with a particular car, for a particular turn or situation. The tradeoffs are fascinating, and are a good example of the complexities that make racing a thinking person's sport.

Factors governing braking technique, including downshifting, include:

- The adequacy of the brakes, for the actual situation. That is, are we considering only one application, or a long race? Is fade a factor? Is pad wear a factor?
- The amount of flywheel inertia in the engine, relative to compression, friction, and pumping work on closed throttle, and relative to the car's ability to slow using the brakes and aerodynamic drag. Does the engine actually try to slow the car, or drive it?
- Brake bias, relative to the optimum for the level of deceleration being attempted. Does adding engine braking at the driven wheels improve brake bias, or impair it? Does ABS make the issue irrelevant?
- The need to be in a particular gear when we finish braking. Also, the need to be in gear, with revs matched to road speed, at the end of braking.

- Handling properties of the car when trailbraking. Does it like to brake and corner at the same time, or not? There may be more than one answer to this, depending on speed range.
- Properties of the transmission. Will it tolerate being being disengaged during braking, and simply stuffed directly into the gear needed after braking, or will this destroy the dogs or synchros? Is the shifter H-pattern or sequential?
- What is before and after the turn. Is an in-slow-out-fast line the best, or an in-fast-out-slow line? Or is there no turn? Are we trying to bring the car to a halt for a pit stop, or for some other reason?
- Tactical considerations. Are we trying to pass under braking? Pass after braking? Avoid being passed under braking? Avoid being passed on exit?
- Driver preference. Race driving isn't easy. It is difficult to use any braking technique to the absolute limit of its potential, and this makes it difficult for a driver to change from one technique to another as circumstances vary, or even to adapt to different cars. In general, a less-than-optimal technique executed well will beat a theoretically better technique executed poorly.

Until the advent of disc brakes, it was quite easy to use up any car's brakes in almost any road race, and often on short ovals as well. Stability in braking was uncertain, especially with the self-energizing (leading-shoe) drum brakes used on road cars. Racing drum brakes often were made with a minimum of self-energizing effect, or were even self-de-energizing (some or all shoes trailing), to minimize this instability, at the expense of high pedal effort.

With drum brakes, it was therefore imperative to brake in a straight line, at least most of the time. It was also imperative to use the engine to slow the car, just to have any brakes left at all late in the race.

When disc brakes arrived late in the 1950's, it became possible to use the brakes much harder, and also to trailbrake – carry the braking process into the first part of the turn. Oval track racers have always trailbraked, if they used the brakes at all. It is not uncommon in oval racing for a driver not to lift until the turn-in point, and do all braking in a curved path.

Trailbraking was one of the keys to Stirling Moss's competitive edge. It was further developed by Mark Donohue, who often used cars with spool rear ends to allow him to brake harder while turning, without getting oversteer. Donohue is also credited with being among the first to figure out that it pays to leave braking until really late, apex early, do a lot of braking in a curved path, and give up exit speed, if the turn is preceded by a straightaway and followed by another turn.

Even with discs, it was possible to use up the brakes, especially in sports cars. It still is in many cars, notoriously NASCAR stockers on road courses. And it was still universal to heel-and-toe down through the gears, although some cars now had close enough ratios so that the driver could skip-shift – go two or even three gears down at once.

By the early 1970's, some classes of cars – notably Formula 1 and other large formula cars – had sprouted huge wings and fat, sticky slicks. At speeds above 100 mph, such cars with a high-downforce setup can exceed 1.0g rearward acceleration without using the brakes at all, mainly because they are so draggy. They also have vast amounts of downforce, and when the driver does use the brakes, the car may decelerate at more than 4.0g!

At this point a school of thought developed which held that the brakes and the drag could decelerate the car faster with the engine disengaged, and that the engine's rotating inertia was such that it was actually trying to drive the car even with the throttle closed, and was fighting the brakes rather than assisting them. Advocates of this view held that it was best to disengage the engine while braking, and then try to find the right rpm and gear for the turn from that state right at the end of braking. Among the drivers using this technique was Francois Cevert. When the driver got it right, it worked pretty well. However, there was also a greater risk of not succeeding in snagging that gear at the last moment, in which case the driver was caught out of gear when the car needed power, and the car might easily leave the road, or at best lose a lot of speed.

This school of thought never gained universal acceptance. Older drivers still went down through the gears as they had always done, although the rapidity of the cars' braking made skip-shifting attractive – in which case the "old" method started to resemble the "new" method, for many turns.

Of course, this controversy was only relevant for cars that had big tires and wings, and H-pattern gearchanges. Production cars, NASCAR stockers, and production-based endurance racing sports cars still decelerated at closer to 1.0g or 2.0g than 4.0g, and still needed to conserve their brakes, and their gearboxes. And in the last decade or so, sequential gearchanges have become universal in the faster classes. With a sequential transmission, skip shifting is not possible. Computer control has made rev-matching nearly foolproof, and overrevving due to a premature downshift is prevented as the computer will forbid the shift until the road speed is correct. So we again see, and hear, all the cars rapidly going down through the gears as they approach a turn. In many cars, the driver has the option of having downshifts and upshifts occur either manually or automatically. A popular choice is to downshift manually (with the computer still having veto power) and let the car upshift automatically.

With extremely powerful brakes, the question of whether the engine assists or fights the brakes is mainly relevant to brake balance rather than ultimate braking power.

Regardless of how we get down to the gear we need for the turn, this much is certain: we better catch that gear by the time we need to turn in, except maybe if the turn-in is very gentle, in which case we may do one last shift early in our trail-braking if we feel confident with this. Once we have initiated the turn, we will need to be very smooth with the brakes and steering, rolling out of the brakes as we wind in more steering, keeping the car as close as possible to the perimeter of the traction envelope. This is hard enough without trying to shift at the same time. As soon as we're done braking, we need to smoothly apply power to balance the car and accelerate out of the turn. We can't be pausing to engage the right gear after braking. We need to be in it already. Therefore, we have to do our downshifting before the turn, while braking, whether it's theoretically best for braking distances or not, simply because we can't afford to take a hand off the wheel, and upset the car with a shift, while we're cornering at the limit, nor can we afford to have the car out of gear or in too high a gear when we need power to control it. In short, we must downshift while braking because it's the only time we can.

All of this assumes that we are slowing for a turn, and will need an appropriate gear as soon as we finish braking. Suppose we are braking for a pit stop, or panic-stopping to avoid a wreck or a deer, or doing a brake test?

If we just want to get the car stopped as fast as possible, most of us will simply stand on the brakes, modulating them if we don't have ABS and have sufficient presence of mind, and just try to remember to declutch soon enough to avoid killing the engine. If we're doing a brake test, we may be able to do comparison runs with the clutch engaged and disengaged for most of the stop. In any case, we won't try to downshift in any modern car.

When pitting, we will usually have a pit road speed limit. To avoid violating this limit, we will need to be in a specific gear – usually a lower one, maybe first – and at a specific rpm. Therefore, when entering pit road we are under much the same constraints as when entering a turn. We need to delay braking as late as possible, we need to be at a specific speed in a specific gear right where the speed limit starts, and we need to be at that speed, in that gear, the instant we finish braking. Therefore, we have to downshift while braking. When we reach the pit stall, our object is to brake as late as possible, not overshoot the pit, position the car precisely to help the crew do the stop efficiently, and have the car ready to make a good launch at the conclusion of the stop.

This last concern may impose varying constraints in terms of getting the car in first gear before it comes to rest, depending on the design of the gearbox, and depending on whether we're in first already due to the speed limit. In most cases it will be desirable, though not absolutely essential, to get first gear while still in motion, and declutch during the stop.

To summarize, downshifting while braking may or may not be necessary to get the car slowed down and to allow the brakes to survive. This varies with the car, the event, even the weather. But downshifting while braking remains necessary regardless of this in many cases, to prepare the car for what we need it to do once we're done braking.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

UPCOMING SEMINAR

The Racing Drivers Club in San Francisco has invited me to present a one-day seminar consisting of a lecture on basics of vehicle dynamics, followed by a question-and-answer session. Location is the Sheraton Hotel near the SF International Airport. Time is 8:00am on Saturday, August 10. Fee is $250 for club members, or $300 for non-members. Since dues are only $30, I think the club stands to recruit some members.

For more information please contact Arthur Muncheryan at rose1art@earthlink.net.

STEERING GEOMETRY VARIABLES

In response to numerous requests for information on steering geometry, I am going to attempt a reasonably complete explanation of the various parameters and their effects.

The *steering axis* is a line about which the wheel steers, usually through the two ball joint centers of rotation in an independent suspension, or the kingpin axis in a beam axle. This line can be defined by the point where it intersects the ground and by its angular orientation. These are commonly described in terms of the X and Y coordinates of the ground intercept, with respect to a local origin at the contact patch center, and the transverse and longitudinal angles relative to ground plane horizontal.

The front view distance from ground intercept to contact patch center, or local Y, is called *scrub radius*, or *steering offset*. It would make more sense to call the top view distance from ground intercept to contact patch center the scrub radius, but most people use the term to mean the Y or transverse component of this. This quantity is generally considered positive when the contact patch center is outboard of the ground intercept.

The side view distance from ground intercept to contact patch center, or local X, is called *trail*, or sometimes *caster trail* or *mechanical trail*. It is positive when the ground intercept is forward of the contact patch center.

The front view angle of the steering axis from ground vertical is called *steering axis inclination* (SAI), or sometimes *kingpin inclination* (KPI). It is positive when the steering axis tilts inboard at the top, which is almost always the case.

The side view angle of the steering axis from ground vertical is called *caster*. It is positive when the steering axis slopes rearward at the top.

These parameters are controlled partly by the design and adjustment of the control arms, and partly by the design of the spindle, or spindle/upright assembly, together with the hub and wheel.

The term *spindle* can mean either the stub axle, or pin, that carries the bearings, or the assembly including this pin and the upright, especially when these are one piece.

The spindle or spindle/upright determines two important parameters: *spindle inclination* and *pin lead* or *pin trail*.

Spindle inclination is the front-view inclination of the steering axis, relative to pin or wheel vertical, as opposed to ground vertical. Spindle inclination approximately equals SAI minus camber. Spindle inclination is almost exactly identical to SAI when camber is zero. It is exactly identical when both camber and caster are zero.

The steering axis and the wheel axis do not have to intersect, unless we want the right and left uprights to be identical parts, with bolt-on steering arms and caliper brackets. The steering axis can pass behind the wheel axis, as it does on a bicycle. The perpendicular distance between the two axes is called *pin lead*. This is equivalent to the dimension we call fork rake on a bicycle. If the steering axis passes in front of the wheel axis, that's *pin trail*. So pin trail is negative pin lead, and vice versa.

Effective pin length is the distance, along the wheel axis, in front view, from the steering axis to the wheel centerplane. This distance depends on the wheel and hub as well as the spindle/upright.

We now have sufficient vocabulary to describe and discuss basic steering and spindle geometry. If we can specify all the quantities above, we have enough data to construct a stick model of the basic steering geometry.

We may want to add steering arms. For purposes of spindle/upright design, we can define the position of the outer tie rod end with respect to the pin and the steering axis. We may define a height from pin axis to tie rod end center of rotation. To do this in a manner appropriate for drawing the upright, or inspecting it when removed from the car, this should be the vertical dimension in side view, assuming zero caster and camber -- in other words, we are projecting to the wheel plane, and taking the steering axis in side view as our local vertical.

In such a side view, we may construct a horizontal line from tie rod end to steering axis. This is our *side view steering arm length*.

We may project a top view from the side view, and locate the lateral position of the tie rod end. If we have a longitudinal line corresponding to the side view steering arm described above, we may construct a transverse line from it to the tie rod end, and measure that distance. This we may call *steering arm offset.* It will usually be outboard for a front-steer layout, and inboard for a rear-steer layout. I don't know what sign conventions other people use, but I generally call outboard positive for front steer and inboard positive for rear steer. Thus positive offset is the direction that gives us positive Ackermann.

In terms of coordinates, we are establishing a local origin where the side view steering arm meets the steering axis. The side view steering arm length is our local X, and the steering arm offset is our local Y.

This doesn't mean there's anything wrong with assigning global or front-suspension coordinates to the tie rod ends when doing an overall front end layout. I'm just pointing out that at some point you will have to deal with the spindle/upright/steering arm unit as a sub-assembly, off the car, and it helps to be able to measure and discuss it that way too.

Now we have a fairly complete vocabulary to describe steering geometry, so we can discuss what effects these parameters have.

Trail causes *lateral* forces at the contact patch to produce a torque about the steering axis. This causes the steering to seek a gravitational/inertial center. The driver feels lateral cornering force through the steering. He also feels the lateral force that the tires must generate to make the car run straight on a laterally sloping, or cambered, road surface. It is worth noting that this is only one component of the self-centering forces the driver feels. Another is the tire's own *self-aligning torque*, which is present whenever the tire runs at any slip angle. This will provide some feedback of cornering force even in the total absence of trail. This effect is sometimes described as mimicking trail. The amount of tire self-aligning torque, divided by lateral force, is sometimes called *pneumatic trail*. Note that this is a calculated value which depends on tire properties, and not an actual steering geometry parameter.

One important distinction between the forces from trail and tire self-aligning torque is that tire self-aligning torque is not a linear function of lateral force. It builds at a decreasing rate as lateral force increases, and at a point a bit short of maximum lateral force it actually begins to decrease. This means that if our car has little or no trail, the steering will start to go light a bit *before* the point of tire breakaway. Some argue that this is a good thing, especially for a passenger car, because it gives the driver a signal to ease up short of the point of actual loss of control. In a race car, this type of steering feel requires that the driver be accustomed to driving just a controlled increment beyond the point where the steering wheel tells him/her that the limit of adhesion has been reached. If the driver is used to having more trail, he/she will often find this very difficult.

Trail also causes a small lateral movement of the front of the car with steer, in the direction of steer. We might call this *steer yaw*. It can rationally be argued that this improves turn-in, both by yawing the car promptly and by causing the rear wheels to develop a slip angle promptly.

Scrub radius or steering offset causes *longitudinal* forces at the contact patch to generate a torque about the steering axis. If right and left scrub radii are equal and longitudinal forces at the right and left wheels are equal, no net torque at the steering wheel results. The driver feels the *difference* between the longitudinal forces at the front wheels. The driver feels one-wheel bumps, brake pulsations, and crash impacts where one wheel hits something, in direct proportion to scrub radius.

A car with a lot of scrub radius is sensitive to wheel imbalance and tire and brake imperfections, has a lot of "wheel fight", and has greater tendency to injure the driver's hands in one-wheel crash impacts or curb or pothole impacts. A car with very little scrub radius is less subject to these problems, but the steering will tend to be numb and uncommunicative.

A car with large scrub radius *may* steer more easily at parking speeds, depending on other parameters, provided the brakes are not applied. This is because the wheels can roll as they steer rather than purely scuffing. With the brakes applied and the car stationary, a car with a small scrub radius steers more easily.

Caster causes the front wheels to lean in the direction of steer. With a given spindle/upright geometry, more caster implies more trail.

Caster combined with trail causes s*teer drop* or *steer dive*. The front of the vehicle drops as the wheels steer away from center, if caster is equal on right and left. This tends to cause an anti-centering force at the steering wheel. It is the reason why the front wheels of a dragster at rest tend to flop to one side or the other.

Caster combined with scrub radius causes the car to drop as the wheel steers forward (toes in), and lift as the wheel steers rearward (toes out). When this occurs on the right and left wheels as one steers forward and the other steers rearward, the result is *steer roll.* The car leans away from the direction of steer. The wheel loads also change. The car de-wedges: the inside front and outside rear gain load; the outside front and inside rear lose load. This effect can help the car turn in slow corners, especially with a spool or limited-slip differential. In excess, it can create low-speed oversteer and over-sensitivity to steering angle. In general, cars running on lower-speed tracks need more steer roll, and cars on fast ovals should have very little.

The camber change associated with caster is favorable, particularly for road racing cars, which usually cannot get favorable camber on both front wheels any other way. We can have too much of this good thing, but that's extremely uncommon.

Steering axis inclination (SAI) causes both front wheels to gain positive camber as they steer away from center.

SAI combined with scrub radius causes *steer lift*. The front of the vehicle rises as the wheels steer away from center. This induces a self-centering force in the steering which seeks vehicle center rather than inertial/gravitational center. This is particularly useful in passenger cars because it reduces the car's tendency to follow road camber, and therefore reduces the need for the driver to pay close attention in casual driving on roads with varying slope. The centering force also tends to suppress steering shimmy.

In race cars, the camber change associated with SAI is unfavorable on the outside wheel. The self-centering force increases steering effort, which is a factor for any vehicle without power steering. It also creates what could be considered a false message to the driver about the lateral forces present at the contact patches. There is therefore a rational case for using more caster and less SAI in a race car.

With the packaging constraints we usually face, more SAI generally implies less scrub radius. The main limitation will often be how far outboard we can place the lower ball joint without having it too close to the brake disc. If the wheel has generous negative offset, we may instead be limited by the wheel rim hitting the control arms in some combinations of suspension motion and steer. Either way, we often cannot place the entire steering axis as far outboard as we would theoretically like to. Using SAI allows us to at least get the ground intercept further outboard in such cases. With MacPherson strut front ends, large amounts of SAI are necessary if we are to obtain any camber recovery in roll.

Consequently, in many cars we see SAI used for reasons not directly related to SAI's own dynamic effects.

A full discussion of **Ackermann effect** (increase of toe-out with steer) is beyond our scope here, but we can at least say that in low speed turns with the wheels steered into the turn, the car generally needs toe-out on the front wheels. For high-speed sweepers or ovals, the front wheels generally need toe-in instead. The key determining factor is whether the *turn center* -- the instantaneous center of curvature of the car's path -- is ahead of or behind the front axle line. Other determining factors include the tendency of the loaded wheel to want a larger slip angle than the unloaded one, and what yaw moments we wish to create with the tire drag forces.

The attitude of the front wheels at any given instant depends on both the static toe setting and the change in wheel-to-wheel toe with steer. This means that optimum Ackermann depends on static toe setting.

It should be clear, then, that there is no such thing as perfect Ackermann properties. But we can at least say some definite things about what geometric parameters will do to Ackermann. In particular, increasing steering arm offset increases Ackermann effect.

Ackermann for oval track cars is often asymmetrical. The side view steering arm length is less on the left wheel than on the right. This produces more Ackermann when steering left than when steering right.

We should mention that if we are willing to tolerate a bit of additional complexity, there are ways around some of the tradeoffs in steering geometry. For example, it is possible to create a self-centering force by springing the steering system. This can mimick the self-centering that we get from SAI, without the adverse effects on camber. We can also damp the steering to reduce kick and shimmy.

We can get small SAI and small scrub radius at the same time by using compound control arms (two single links replacing the usual wishbone or A-frame) and dual ball joints. This gives us an instantaneous virtual ball joint outboard of the linkage itself. We can adopt this arrangement at the upper end of the upright, or the lower end, or both.

IMPORTANCE OF STEERING RACK PLACEMENT

I have read that one should either place the steering rack above and behind the wheel axis, or below it and ahead of it, due to deflection steer considerations. How important is this really in race cars, where there is no rubber in the suspension or the steering mounts?

The stiffer everything is, the less this matters. Most formula cars nowadays violate the rule you mention; the steering is ahead of the wheel axis and above it.

Of course it is still true that all cars have some deflection steer, and we would prefer that this be deflection understeer rather than deflection oversteer.

Actually, what determines the critical height for this is not necessarily the height of the wheel axis. Rather, it is the point of zero lateral deflection at the upright, which varies with the combination of loads on the wheel, the geometry of the system, and the distribution of rigidities in the components.

Another consideration is that we may prefer the tie rod on the loaded wheel to be in tension rather than compression, on the logic that the tie rod tends to be the slenderest element in the system. This concern argues for front steer in all cases.

BUMP STEER AS ACKERMANN MODIFICATION

I have a Pantera that has very little bump steer in compression from static position, but has quite a lot of toe-in in droop. It also has a fair amount of Ackermann. Could it be that this is intentional, to provide less Ackermann in hard driving when there is more roll, yet have adequate Ackermann in tight turns at low speed? Have you ever heard of a car being designed that way on purpose?

I haven't ever worked with a Pantera, but I don't really think the combination was thought out on the logic you describe. It is quite common, though, for production cars to exhibit the type of bump steer you describe. The cause is usually that the inboard pivot axes of the lower control arms are splayed out toward the rear in plan view, while the inboard pivot axes of the upper arms are not so much so, or even angle outward toward the front.

This makes for a side view virtual swing arm that shortens dramatically in droop, causing caster to diminish at a rapidly increasing rate in droop. That unavoidably causes some toe-in in droop if the geometry is laid out to avoid major toe change in bump.

The usual reason for designing the lower control arm this way is to get more front view control arm length, while still making room for the engine or the footwell.

THE OLD TORQUE/HORSEPOWER CONTROVERSY

More horsepower makes a car accelerate quicker – true or false?

The above question was recently included in a quiz in a magazine, written by a reputable author who gets this newsletter. To my surprise, he said the correct answer was "false". He explained that, assuming all else identical, including tire size, more acceleration implies more axle torque, not necessarily more horsepower.

What this reasoning misses is that since power is the product of torque and rpm, more torque at the same rpm *is* more power, at the axle.

Now, if we have more axle torque at a given axle rpm, what can we say about conditions at the engine flywheel? Do we have more torque at the flywheel? Not necessarily. The author himself points out that we could have identical torque, and more torque multiplication due to lower (numerically higher as customarily expressed) gearing.

Okay. So if there is identical torque at the flywheel, identical speed at the axle, and shorter gearing, what do we know about engine rpm? It must be higher, in direct proportion to the gear. Again, power is the product of torque and rpm. Therefore, if torque is identical and rpm is greater, power is greater, at the engine.

If the axle has more torque at a given speed, that alone does not tell us anything about torque at the flywheel. But we can say with certainty that both the product of engine torque and gear, and the product of engine torque and engine rpm, are greater. This means we know that both power at the engine and power at the axle are greater.

Looking at it another way, power is the rate of energy flow. To make the engine add kinetic energy to the car at a greater rate, energy must flow from the engine at a greater rate.

Note that this is a somewhat different question from whether to build a “horsepower motor” or a “torque motor”. What counts at a particular instant is horsepower *at that instant*, not at the peak of the power curve. In many cases, especially when running a short oval without shifting, a car will exit a turn at an rpm closer to its torque peak than its power peak. In such a case, making the engine strong at the torque peak may make for better lap times – but we are still talking about a *power* increase (achieved through more torque at comparable rpm – or alternatively through comparable torque and more rpm and gear, or some combination of both approaches), even though the increase is at a point in the rpm range that is below the engine’s power peak.

So a broad power band is worthwhile, and may be worth more than maximum peak horsepower. But at a particular instant, at a particular speed, more axle torque definitely implies more engine horsepower, and by itself implies nothing about engine torque.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

PANTERA FRONT END REVISITED

The person who submitted last month's question about his Pantera's steering geometry informs me that the car does not have increasing anti-dive in droop, contrary to my speculation. The car has inner control arm axes that are very close to perfectly horizontal and longitudinal, meaning it has little or no anti-dive at any ride height.

In such a front end, the presence of substantial toe-in in droop, with minimal toe change in bump, suggests that (assuming front steer) the rack length is a little too long for the control arm layout – at least at the height the rack is mounted. This makes the tie rods a little too short, and produces a bump steer curve that we could call C-shaped, as opposed to J-shaped. But we can minimize bump steer in bump, and therefore on the loaded wheel when cornering, by raising the rack, or lowering the outer tie rod ends, relative to the setting that would give least instantaneous rate of toe change at static position. In other words, we can tilt the C-shaped curve. We can make the wheel toe out a little, and then in a little, through the bump range, at the expense of having it toe in a lot in extreme droop.

This is not an ideal situation, and it is doubtful that the builder is trying to get any benefits this way. Rather, this looks like a way to use a readily available rack that's a little longer than ideal, without getting roll oversteer.

HORSEPOWER AND TORQUE REVISITED

The author of the article on horsepower and torque on which I offered some comments last month points out that he was not necessarily assuming equal road speed in the cases he was comparing. I agree that it is possible for a car with lower horsepower to accelerate faster than one producing more horsepower, and to correspondingly have higher axle torque, if we are supposing that the car with less power is at a lower road speed. This is an unusual way to compare two cars, but the point does stand.

A corollary is that a given car accelerates slower and slower the faster it goes. Most of us know this already, but perhaps not everyone is aware that this would be true even if aerodynamic drag were absent or did not vary with speed. As road speed rises, a car has to use a numerically lower overall gear ratio to get a given engine rpm, and consequently torque multiplication diminishes.

Be this as it may, at any given road speed, other things being equal, more torque at the axle still implies more power at the engine, and does not necessarily imply more torque at the engine.

FRONT AND REAR CENTER OF GRAVITY?

I write to you in the hope that you may be able to help me with a few calculations regarding center of gravity. The situation is that I am chasing to restore the balance of our car (Australian Sports Sedan – 20B/RX7) once had. [The writer is referring to an early Mazda RX-7 with a later twin-rotor engine installed.] *The car's balance was disturbed once we fitted the 20B engine, adding about 200 lbs of weight. After chasing our tails for 4 years I decided to do something about it. I collected all of Carroll Smith's books, and Allan Staniforth to mention another, in order to gain a better understanding of racecar dynamics.*

Of interest at this time is a chapter in Allan Staniforth's book written by David Gould. The chapter has some nice formulae to determine lateral weight transfer and roll resistance. However, I am stuck at determining a few things that the formulae require.

The first problem I have is that he asks for the height of the CG of the front or rear unsprung mass (to calculate front and rear unsprung weight transfer). The only clue given to work this out is that the figure is usually similar to the radius of the tire. Is there another way of calculating these figures?

The second problem is similar: the equation to determine the mean center of gravity of the sprung mass. The author has mentioned here that to find the center of gravity at the front and rear axle lines can be an extremely difficult process and the preferred process is deduction from the known location of major items that comprise the sprung weight. Once again, would you know of any other techniques that can be used to determine the height of the center of gravity at the front and rear axle lines?

It is customary to model a car as a single sprung mass, flexibly connected to front and rear axles. Even if the suspension is independent, we approximate the independent systems at the front or rear to an equivalent single axle for modeling purposes. We thus have three bodies, each of which has a center of gravity (or center of mass). We may also regard the whole vehicle as a single mass, and this will also have a center of mass, which can be taken as a weighted average of the centers of mass of the three masses that comprise it. All three of these masses are assemblies of components, all of which have their own centers of mass, but we treat the car as three masses because these three are movable with respect to each other, for chassis analysis purposes.

For most designs, it works quite well to assume that the CG heights of the front and rear unsprung masses are at their respective hub heights, or tire loaded radii. This approximation is very accurate for the wheels and brake rotors. Other components may be higher or lower, and you can weigh them and estimate their individual CG locations, then calculate a CG location for the entire unsprung assembly, but for most cars this will come out very close to hub height.

I am amazed at the number of books that talk about front and rear **sprung** mass centers of gravity. The sprung mass is a relatively rigid single assembly, and it has one center of gravity (or center of mass), not a front and a rear one, not a string of them. The sprung mass is not two men in a horse suit, nor is it a train of flexibly connected bodies. Like any other body, it has one and only one point at which (or through which) a linear force can be applied from any angle without causing the mass to rotate in any direction, and that point is the center of mass. This point lies between the axles, at a height somewhere between 12 and 20 inches for most cars.

The best method to estimate the location of the center of gravity for the sprung mass, or the whole car, depends on whether we are dealing with a design or an actual car. If we are trying to model a car not yet built, we have to weigh as many purchased components as we can, estimate where their individual centers of mass are, estimate the weight and individual centers of mass for components we're designing, and then take moments in three planes about a convenient origin point to calculate the overall CG. One common choice for this origin point is a point centered between the front contact patch centers, at ground level. Another is a point centered laterally, and midway along the wheelbase, at ground level. Any point at all can work mathematically, but these are convenient. To estimate the overall vehicle CG location, we include the unsprung components. To estimate the sprung mass CG location, we omit the unsprung components.

For a car that's an assembled vehicle, sitting in our shop, we directly measure the whole car rather than taking it apart and weighing components. To find the lateral and longitudinal coordinates of the overall CG, we simply scale the car as we usually would when setting up for a race, and calculate the CG position in plan view from the left and rear percentages. If the car has, for example, 55% rear, then the overall CG is 55% of the wheelbase aft of the front axle line. If it has 51% left, then we find the overall CG's lateral position as follows:

1) In a top or plan view, construct lines connecting the contact patch centers on each side of the car.
2) Construct a line perpendicular to the vehicle centerline, through a point 55% of the wheelbase aft of the front axle.
3) Find the length of the line from step 2 between the lines from step 1.
4) Find the point on this line 51% of its length from the right end. This is the plan view location of the overall CG.

That's two thirds of our answer. We still need to find the height of the overall CG. This is the tricky part. The method usually described in chassis books involves elevating one end of the car, with the other end on scales, and noting the weight change at the scales. The distance of the CG above the unraised axle, the one on the scales, is given by: $h = [(\Delta W)(L)\sqrt{(L^2 - x^2)}] / (W)(x)$, where:

h = CG height above the axle on the scales, at static condition
W = total vehicle weight
ΔW = weight increase on the scales when opposite axle is raised
L = wheelbase length
x = height that axle not on scales is raised

Units of length for h, L, and x can be inches, feet, meters, or whatever, but must all match. Likewise, units of weight for W and ΔW must match.

One drawback to this method is that the overall CG of many race cars is not much above the axle lines, meaning that we are measuring a very small h. This means that accuracy becomes a problem. The smallest shift in fluids or other masses can noticeably affect the measurement, and the car may have a significantly different CG height if we drain all the oil and fuel than it has with a normal load of fluids. This may require a calculated correction. The end that's raised has to be raised a significant distance, generally at least 3 feet, and must be supported with the wheels free-rolling on perfectly level surfaces. This requires special stands and the means to hoist one end of the car that far. The suspension should be immobilized at normal ride height. The car may be rather precariously balanced when raised, and in many cars ground clearance of the front and rear overhangs will limit how far we can lift either end, unless we support the lower end on pedestals too.

A client of mine came up with an alternative method. It likewise requires immobilizing the suspension and draining the fuel and oil. He raises one side of the car, rather than one end. He has made large "shoes" of angle stock that cradle the two tires that remain on the ground. He uses a wrecking truck to lift the other side of the car. He raises it until the car balances on the edges of the "shoes". He leaves the hoist attached, so the car can't fall over, and gives the car just a little slack. He can then stand beside the car, move it gently by hand, and feel for the balance point. Once he finds the balance point, he measures the car's tilt.

Once that point is found, he knows that the overall CG lies in the vertical plane passing through the "shoe" edges. This we might call the balance plane. When the car is sitting normally on its wheels, this plane assumes an angle from vertical equal to the tilt of the car when the balance plane was found. He knows the plan-view position of the CG from scaling the car during setup, as described above. So the CG lies on a vertical line that appears as a point in plan view. He just finds where the balance plane intersects that vertical line, and that's the overall CG.

Once we have the overall CG location, we can calculate sprung mass CG if we have weights and estimated CG heights for the unsprung masses. For cars with independent suspension all around, the sprung mass CG will be at approximately the same lateral and longitudinal position as the overall CG, and a little higher. For cars with live rear axles, the sprung mass CG will be a bit forward of the overall CG, and a little higher.

If all this seems intimidating, be aware that you can do calculation exercises, and learn a lot about vehicle dynamics, using assumed values for the CG locations. You can also determine the longitudinal and transverse location of the overall CG by measurement, and assume only the height. For your RX-7, assumed heights of 16 inches for the overall CG, 17.5 inches for the sprung mass CG, and tire loaded radius for the front and rear unsprung mass CG's are probably pretty realistic.

As a general rule, to re-balance the cornering of a car that has received a larger engine, you need to lighten whatever you can in the front, move whatever you can to the rear, increase the share of the roll resistance at the rear, and adjust the brake bias for more front.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MOVE THE REAR WHEELS BACK TO IMPROVE TRACTION?

Regular readers will be aware that this newsletter is used as the basis of a column in *Racecar Engineering* called “The Consultant”. In the October 2002 issue, Simon McBeath has an article about the DJ Firehawk hillclimb car, in which he remarks: *Creating that slightly longer wheelbase was achieved by sweeping back the rear wishbones, which also had the benefit of minimising chassis overhang, effectively shifting some weight forward and keeping the major masses well within the wheelbase. The result is almost exactly 40:60 weight distribution front to rear. The swept-back rear suspension has, apparently, also been found to be an aid to traction in US buggy racing. Perhaps this is one for The Consultant’s column?*

First, I agree that 40/60 weight distribution is a good general-purpose target for a car where the rear tires can be larger than the front ones. I also agree that moving the differential forward of the axle line a bit is advantageous in terms of yaw inertia, although there is a penalty in halfshaft joint longevity and frictional losses.

I don’t work much with off-road cars, but I can speak to the physics of traction. In any rear-wheel-drive vehicle, more rear percentage – static or dynamic – translates to more traction, in straight-line forward acceleration or hill climbing. From this standpoint, moving the rear wheels aft (or the front ones) hurts traction, since it reduces both static and dynamic rear percentage. So the simple answer is that moving the rear wheels aft for better traction is wrong. For better traction, you need to move them forward, or move major masses back.

However, in the real world we are not always dealing with pure forward acceleration, and forward acceleration is not always limited by the rear wheels’ capability to generate forward force.

In the real world, we not only have to put power down – we also need to steer. We also have to avoid flipping the vehicle over backwards. These considerations may lead us not to go for maximum static rear percentage in all cases.

The first case I can recall of somebody moving the rear wheels back on an existing vehicle occurred in the 1970's when Porsche moved the rear wheels on the 911 back a couple of inches. This was not done to improve forward traction. It was done to make the car a little less tail-heavy and thereby reduce limit oversteer without making the rear tires drastically larger than the fronts.

It is commonplace in hillclimb and drag racing motorcycles to fit a long rear swingarm and move the rear wheel aft as much as a foot. Why is this done? Because the vehicle is wheelstand-limited rather than traction-limited. The center of gravity is high. The wheelbase is short. When grip is good or the grade is steep, the rider is limited by the need to keep the bike from flipping rearward on top of him, rather than wheelspin.

Drag cars, especially ones that are required to resemble road cars, can be wheelstand-limited up to surprisingly high static front percentages, because of the extraordinary grip of modern drag tires. When drag slicks were less evolved, everybody tried to move the weight back as far as possible and raise the CG as far as possible. This is the way to go, up to the point where the front wheels come off the ground. Once that point is reached, we are faced with a delicate balancing act: we must maximize rear wheel loading but still be able to steer with the fronts, within the range of forward accelerations we can anticipate. It turns out that the best design is one with a very long wheelbase, a low CG, and ample static rear percentage – a rear-engined dragster. But a funny car, with a much shorter wheelbase and much greater static front percentage, can turn nearly as good a time – provided that the traction characteristics of the track and tires are what we normally expect. Pit the same cars against each other on a slippery surface, or on street tires, and the funny car has no chance.

When we are dealing with forward acceleration while cornering, things get even more complex. The front tires not only have to afford us some measure of directional control, they also have to provide sufficient cornering force to keep the car from developing a power push.

For a given track and tire, limiting lateral acceleration at the front or rear of the car depends on the relationship between the dynamic normal (vertical) force on the tires and the centrifugal force they are required to resist, plus the fact that a tire's ability to generate lateral force diminishes when we ask it to generate longitudinal force at the same time.

When we have a car at its limit in lateral acceleration, and we then ask it to accelerate forward as well, we know that the car will have to reduce its lateral acceleration and widen its arc. Whether it gets looser or tighter depends on the balance between two conflicting effects. The first effect is rearward load transfer. The normal force on the rear tires increases, and the normal force on the fronts decreases, while the proportion of centrifugal force at each end of the car remains essentially unchanged since its constituent masses do not shift significantly. This tends to tighten the car (increase understeer). The other effect is that with rear wheel drive the rear tires must give up some lateral force capability in order to make forward force, while the front tires are not required to do this. This effect tends to loosen the car (increase oversteer).

Many factors other than wheelbase and CG location influence cornering balance under power, including differential type, toe settings, tire design and inflation, and suspension design and settings. In oval track cars, we have additional effects from suspension asymmetries, tire stagger, and static left percentage. But let's consider the influence of wheelbase, CG height, and front/rear CG position, assuming other factors are held constant.

Almost any rear-drive car will be tightened by very moderate power application, and loosened if we apply enough power to create obvious wheelspin. In between these extremes, the car has a range of throttle position where power tightens the car compared to steady-state, and at some point a transition to power oversteer. If that transition to power oversteer occurs only with obvious wheelspin, and the car just gets tighter with power up to that point, we say it has power understeer, or a power push. If anything more than minimal power application loosens the car, we say it's loose on power or prone to power oversteer.

Desirable behavior lies somewhere between these extremes. For best exit speed, we would like the car to remain reasonably balanced as power is added, and just use up more road as we add more power. The driver may like some power oversteer so he can steer the car with the throttle, but as a rule this will exact some penalty in exit speed.

The limiting lateral acceleration at the front or rear wheel pair depends in part on the relationship between normal force on the wheel pair and the percentage of the car's mass that the wheel pair must control. Let's look at how normal force varies at the front and rear on some hypothetical cars, at 0.5g forward acceleration. The load transferred from the front wheel pair to the rear wheel pair due to forward acceleration is given by: $\Delta F_n = Wa_x h_{cg} / L_{wb}$ where:

ΔF_n = absolute change in front or rear wheel pair normal force due to longitudinal acceleration
W = total vehicle weight
a_x = longitudinal (x axis) acceleration
h_{cg} = height of vehicle overall center of gravity
L_{wb} = length of wheelbase

Working in English units, we use pounds for W, express a_x in g's, and obtain ΔF_n in lbf. With metric units, we would classically use the vehicle's mass in kg for W, express a_x in m/sec^2, and obtain ΔF_n in newtons. However, since our wheel scales will probably read in kg, we may find it more convenient to use g's for a_x and get ΔF_n in kg.

Suppose our vehicle has a CG height equal to 1/5 the wheelbase. That's a fairly short, high car, perhaps a sprint car or a midget. We now have a value of 0.2 for h_{cg} / L_{wb}. If $a_x = 0.5g$, then $\Delta F_n = W(0.5)(0.2) = 0.1W$.

If our car weighs 1000 lb and has 50% static rear weight, then assuming no turn banking and neglecting aerodynamic effects, our front normal force is 500 lb in steady-state cornering. If we accelerate forward at 0.5g, ΔF_n = 100 lb and total dynamic normal force on the front wheel pair is 400 lb. This means that the front end has a limiting lateral acceleration a bit greater than 80% of what it had in steady-state cornering. I say a bit greater than 80% because a tire's coefficient of friction usually increases a bit as normal force diminishes.

Now let's suppose we have the same situation, except the car has 60% static rear. ΔF_n is still 100 lb, so we now have front normal force going from 400 lb to 300 lb. Therefore, the front end's limiting lateral acceleration is now a bit greater than 75% of its steady-state capability – less than in the previous example.

If the car has 70% static rear, front normal force goes from 300 lb to 200 lb, and front lateral acceleration capability is a bit greater than 67% of steady-state.

If we look at these three cars at a_x = 1.0g, we have front end lateral acceleration capabilities slightly greater than 60%, 50%, and 33% of steady-state. At a_x = 1.5g, the first car has front end lateral acceleration capability that is >40% of steady-state, the second car has >25%, and the third has **zero**; it's at the point of impending wheelstand.

What happens at the rear wheels is somewhat more complex. With more static rear, the percentage increase in normal force is less for a given a_x, but the forward force required to produce that a_x is also a smaller portion of the tire's overall vector force capability. In general, it is harder to induce power oversteer in tail-heavy cars.

My point here is that tail-heavy cars put power down better, both in a straight line and when cornering, but as the car becomes more tail-heavy the unloading of the front wheel pair becomes greater percentagewise for a given a_x, and the car becomes more prone to power understeer. At some point, power understeer will limit power application on exit before wheelspin will. When we have such a condition, the car may actually achieve better a_x on exit with less static rear percentage.

It will be apparent that as h_{cg} decreases, or L_{wb} increases, ΔF_n decreases for a given a_x. This means that a lower or longer car can have more static rear percentage without encountering wheelstanding or power push. It also means that a lower, longer, more tail-heavy car will perform more consistently as grip varies.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANOTHER TIRE CARE TIP

In the past I have written in these pages about bagging tires and keeping them in cool, dry, thermally stable places. I recently read that it's also a good idea to keep them away from electric motors or generators that use brush contacts, as these generate ozone. I'm not sure just how significant this is, but it sounds logical.

PULL RODS, PUSH RODS, OR BOTH?

Do push rods or pull rods have an adverse effect on handling? If so, does each selection of rod induce oversteer or understeer in the chassis? I know a combination of pull and push rods was used by McLaren in the early 1990's, taking Senna to numerous victories. Has anyone tested using a combination of pull rods on one side of the car and push rods on the other (left and right)?

Push rod and pull rod suspensions are primarily packaging solutions. They get the shock or coilover in out of the airstream. They also allow convenient tailoring of rising-rate effects via rod and rocker geometry, and facilitate interconnections such as third springs and anti-roll bars.

However, it is also possible to provide most of these effects with outboard or direct-acting springs and dampers. The Z-bar on the rear of a Formula Vee, the camber compensator on a late Porsche 356 or '64 Corvair, the swing spring on a late Triumph Spitfire, or the coil spring above the differential on a Mercedes swing axle all do essentially the same thing as the third spring in a modern race car suspension.

It is a fairly common misconception that push rods or pull rods affect the dynamic load transfer (weight transfer). Actually, all the added forces are reacted within the car, and the only thing the tires "feel" is the wheel rate – the effective spring rate, at the wheel, in the modes of suspension movement being encountered at a given moment.

Another misconception is that inboard coilovers reduce unsprung weight, rather like inboard brakes. That is also erroneous. Anything that moves when the wheel moves is unsprung weight, whether it moves horizontally, vertically, or obliquely. So adding the rod and the rocker to the system is a disadvantage in terms of unsprung weight, and the part of the damper and spring that move with the rocker are still unsprung.

Ordinarily, we try to make suspension layouts on road racing cars symmetrical in their behavior, and the easiest way to do this is to make them symmetrical in their layout. It is possible to obtain symmetrical dynamics with push rods on one side and pull rods on the other, but that's doing it the hard way unless we are up against some very unusual packaging constraints. There might potentially be a reason in an oval-track car, perhaps a supermodified. I'm seeing a lot of pavement supers, midgets, and sprint cars these days with right side coilovers hung way out from the frame on outrigger structures. Rockers, or rockers with push rods or pull rods, could bring these inside the body, improve the effectiveness of a low front wing, and move the weight of the right side dampers lower and further left.

TRAIL AND STEERING FEEL

I'm on a Formula SAE team. In your August 2002 newsletter, you mention increased steering effort and also better ability to sense the amount of traction at the front wheels, if the steering geometry includes more mechanical trail. You also went into the difference in feel between pneumatic and mechanical trail near the point of breakaway. I am hearing conflicting advice from other sources on this question. Carroll Smith told us FSAE participants that ample trail is good. On the other hand, the first Milliken book suggests that mechanical trail masks the feel from the tire's self-aligning torque. We are confronted with driver complaints about steering effort, but we don't want a numb race car. Could you elaborate on the pros and cons of mechanical trail as it relates to steering effort and feel?

As I mentioned briefly in August, this is to some extent a question of personal preference, and drivers differ on whether they like the steering to lighten before true breakaway, or at the point of true breakaway. My own preference is for the steering not to lighten unless lateral force is actually going away.

This may be partly just my own conditioning, but I do think there is a rational or objective case for not having the steering lighten below the actual point of traction loss. Even if the driver can get used to the feel of the steering going light short of the limit, how is even the best driver to distinguish between lightening due to pneumatic trail reduction and lightening due to reduced grip? To some extent the driver can "interpret from context" but basically less force is less force, and your hands shouldn't be expected to know if a lightening in the steering is the tire doing its normal thing or the tire hitting pavement with less grip. With ample trail, you don't have to be psychic. Less steering force means less cornering force, period.

I haven't seen your car, but a common problem in FSAE cars I have seen is excessive U-joint angularity in the steering shaft. In the quest for a short turning circle via short wheelbase, many FSAE cars position the front wheels a lot further aft than they really ought to be from a weight distribution standpoint. This not only adds load to the front wheels, but often requires outrageous bends in the steering shaft. If your last car had this problem, consider moving the front wheels forward, while shortening the steering arm length to get more lock with similar rack travel. The reduced arm will add steering effort, but for a given turning circle you approximately get the difference in effort back again from the reduced front wheel loading, and the steering shaft can be much straighter.

KART SCALING

As part of my Mechanical Engineering degree, I have been tasked with redesigning a kart chassis to make it easier to manufacture. I want to understand some of the dynamics of the current kart so I can model the new one more effectively:

a) *How can I measure corner weights without resorting to expensive electronic scales?*
b) *How can I measure the loads the chassis is exposed to as the kart corners? I want to understand how the chassis is stressed in the dynamic state so I could load up my model in a similar way and alter the model to achieve the same kind of handling characteristics.*

That's some assignment. Karts have been around for over 40 years, and they were originally conceived as a simple, low-cost, easily manufactured item. Improving manufacturability of something like that, or simplifying it at all, is a real challenge.

There are many kinds of karts, but all of them are about as simple in terms of chassis design as a four-wheeled vehicle can be. You can simplify some of the more elaborate ones by making them more like the lower-class ones, but that's obvious.

I can at least offer some help with the scaling. You can't get away from the need to use four scales of some sort, but you can use bathroom scales, which are nowhere near as costly as electronic racing scales. Try to get four that read as nearly the same with you standing on them as possible, and calibrate from there as needed.

You will find that your diagonal percentage is greatly affected by whether all the scales are in a common plane, and whether the steering is centered. You will want to level the scales carefully and devise a repeatable regime for positioning the steering. You will also find that tire pressures affect wheel loads, so you need to have accurate inflation for repeatable wheel load measurements.

Most karts have adjustable front ride height. The simplest system consists of washers on the kingpins. You will find that this allows you to obtain almost any desired diagonal percentage.

If you are setting up for road racing, and you don't want wedge in the chassis, you can get fairly decent results by setting front ride heights so both front wheels come off the ground simultaneously when you jack up or lift the front end at a central point. Again, you will find that even this crude test is sensitive to tire pressures and steering centering. This doesn't tell you wheel loads, of course.

With a vehicle that has suspension, it is highly desirable to have platforms that allow you to roll the vehicle back a few feet and then forward again onto the scales, to settle the suspension. With a kart, you don't need to do that. You can just set the kart on the scales.

You will find that left and rear percentages do not depend much at all on steering centering and tire pressures. You can therefore easily calculate what the wheel loads would be in an unwedged condition (equal left percentage at both ends, equal rear percentage on both sides).

Driver weight is generally at least half of the total, so things change dramatically depending on who's driving.

Kart frames are deliberately made torsionally flexible, and somewhat flexible in beam as well. This design objective is a happy mate with low cost and ease of manufacturing; the simplest perimeter frame imaginable fills the bill. The frames receive torsional loads not only from bumps, but also from caster jacking as the wheels steer. Spindle lengths are sometimes adjustable with shims to vary the magnitude of this effect by changing scrub radius (steering offset). Caster jacking is used to unload the inside rear wheel in tight turns and help the vehicle rotate in yaw without a differential.

I expect you can determine stress levels in an existing kart frame while running if you can instrument it with strain gauges. I don't know if your school has the necessary equipment to do this. I would think buying it yourself would be tough if affording scales is an obstacle. Stresses in any vehicle depend greatly on operating conditions, and this is especially true for a vehicle with no suspension except tire and frame deflection.

It is a long-established tradition in motorsports to do laboratory tests on frames which do not necessarily faithfully reproduce running stresses but are easy to do. For a kart, relevant tests would be overall torsional stiffness and maybe yield strength, with load points where the axle and spindles attach, and beam stiffness and yield strength with load applied at the seat and support points at axle and spindle mounts.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
December 2002

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ACKERMANN RECOMMENDATION

I am modifying a road racing Formula Ford for SCCA Solo 2 [American autocross]. I am considering adding more Ackermann effect to make the car work better in tight turns. Is this a sound idea, and if so, what do you suggest for geometry?

Without writing a really long piece on Ackermann, yes you will probably help the car. I don't know what geometry you have now, but as a general rule a car needs more Ackermann for events with tight turns, e.g. autocross or hillclimbs.

There isn't a universally agreed way to express how much Ackermann (toe-out increase with steer) a car has. The closest thing we have is to take the plan-view (top-view) distance from from the front axle line to the convergence point of the steering arm lines, divide the wheelbase by that number, and express the quotient as a percentage. If the steering arms converge to a point on the rear axle line, that's said to be 100% Ackermann. If they converge to a point twice the wheelbase back, that's said to be 50%. If they converge to a point 2/3 of the wheelbase back, that's said to be 150%. If they are parallel, that's zero Ackermann. If they converge to a point twice the wheelbase ahead of the front axle, that's said to be –50%.

Supposedly, with 100% Ackermann, the front wheels will track without scuffing in a low-speed turn, where the turn center (center of curvature of the car's motion path) lies on the rear axle line in plan view. This is actually not strictly true, even for the simplest steering linkage, which would be a beam axle system with a single, one-piece tie rod. With either a rack-and-pinion steering system or a pitman arm, idler arm, and relay rod or center link, we can't fully predict what the Ackermann properties will be at all, merely by looking at the plan view geometry of the steering arms. The whole mechanism affects toe change with steer.

Even knowing what instantaneous toe we want in a specified dynamic situation is not simple. We don't necessarily want equal slip angles on both front tires. For any given steer angle, the turn center might be anywhere, depending on the situation. All the infinitely numerous possible situations will have different optimum toe conditions. Therefore, there is no relationship between steer and toe that is right for all situations.

The toe we have at any particular instant results not only from Ackermann effect, but also from static toe setting and toe change with suspension movement (roll and ride Ackermann).

Because of these complexities, there is no single obvious way to define what constitutes theoretically correct Ackermann. It is possible to come up with a rationally defensible definition for your own purposes, but there is no standard rule, and it is unlikely that there ever will be.

Having entered these abundant caveats, I will now make some general-purpose recommendations for autocross and hillclimb applications:

1. In plan view, at zero steer (straight-ahead position) the steering arms should converge to a point somewhere between the rear axle line and the midpoint of the wheelbase. In traditional parlance, that’s somewhere between 100% and 200% Ackermann. The tighter the turns, the higher the percentage.
2. At all steering positions, the rack or relay rod should be either slightly behind the outer tie rod ends or even with them. This applies to both front steer and rear steer cars. With rack and pinion steering, it means that at zero steer, the rack should be a bit behind the outer tie rod ends on a front steer car, and about even with the outer tie rod ends on a rear steer car. Purpose of this is to assure that tie rod angularity adds Ackermann at large steer angles, rather than subtracting.
3. Angle between any arm and any link in the system should never be less than 30 degrees or greater than 150 degrees. This helps to assure that the mechanism cannot snap over-center due to deflections of the components. Alternatively, over-centering can also be prevented by provision of stops at appropriate points in the mechanism.

TORQUE, RPM, AND POWER DISTRIBUTION IN DIFFERENTIALS

I would like some clarification on the issue of torque distribution between the front and rear axles on 4wd vehicles. I find the matter fairly easy to understand when you have wheels spinning, and a limited-slip differential, but I find it more confusing when I read statements that a vehicle has a permanent torque distribution of, say, 32% front and 68% rear.

To me, torque and revolutions go hand in hand: reduce rpm and you increase torque, as in a ring and pinion. Doesn’t that mean that if you want different torque at the front and rear axles, they have to turn at different speeds?

I know that in vehicles with viscous coupling drive to one axle, one can have a different overall drive ratio at each end, and this is often deliberately employed just to load the system in normal driving, and make it respond quicker to traction loss. But how does a rigid system, with a planetary differential for example, split torque unequally?

When we are dealing with one input torque, from one gear or shaft, and one output torque on a single shaft or other member, the relationship you describe between torque and speed does hold. Neglecting friction, power in equals power out. If rpm is changed, torque must change too, in inverse proportion, for the product of the two (power) to remain constant.

However, when the output power is divided between two shafts by a differential, things change a bit. Total power in still equals total power out (again neglecting friction), but power at each of the two output shafts is not necessarily equal to power at the other shaft. Any non-locking differential maintains a fixed distribution of torque between the two output shafts, while letting their relative speeds vary freely. In a conventional differential, the torque split is 50/50. In a planetary differential with one planetary gearset, the torque split is unequal but still fixed, while the shafts can turn at different speeds.

Usually the differential carrier or planet carrier is driven by a gear, which receives power from another gear driven by the input shaft. At the carrier, the simple inverse relationship between speed and torque applies. Torque at the carrier is input torque times rpm reduction factor. The sum of the output torques equals the carrier torque. The average of the output speeds equals the carrier speed. Power at each individual output shaft can be any value at all. It is even possible to have negative power (retardation) at one output shaft if that shaft is being forced to turn backward (opposite to torque). But the sum of the two power outputs must equal the power input. (That's the sum of their signed values, not their absolute values.)

It is helpful to think of each spider or planet gear as being similar to a beam, with a load applied at its midpoint, and reaction or support forces at two points equidistant from the load. The load is the drive force applied at the spider or planet gear's shaft. The reaction forces are the output shaft resistances to vehicle motion, acting at the points of mesh between spider and side gears, or between planet and sun and planet and annulus. Since the spider or planet shaft is always at the gear's center, the forces at the mesh points are always equal. This is true regardless of the rotational speeds of the various elements.

In a conventional differential, the side gears are equal diameter, so the equal forces at the mesh points act on equal moment arms, and produce equal torques. In a planetary, the annulus is larger than the sun, so the output torque at the annulus is greater than the output torque at the sun. The ratio of the output torques is the ratio of the pitch diameters of the annulus and sun. So the bigger the planet gears are in comparison to the sun, the more unequal the torque split becomes. Usually, the annulus drives the rear axle and the sun drives the front axle.

We can, in fact, regard the conventional differential as a unique version of the planetary, cleverly reconfigured by the use of bevel gears to allow the sun and annulus to be the same size.

All of this determines the torques at the front and rear drive shafts. Usually, the main rpm reduction and torque multiplication (after the transmission) happens at the axle, not at the transfer case. It is possible to use different ring and pinion ratios at the front and rear axles, and/or different tire sizes front and rear, and further alter the drive force distribution at the tire contact patches. At the axles, the usual rpm/torque inverse proportionality applies. To get more front torque and less rear by using dissimilar axle ratios, the front drive shaft must turn faster than the rear. That will increase wear at the center diff, rather like traveling a long distance with unequal size tires on an axle. Actually, the least wear at the center diff comes with slightly less torque multiplication at the front axle than at the rear – say a 4.10:1 ring and pinion at the front and a 4.11 at the rear. This is because even on a straight road, the car doesn't quite go perfectly straight, and in most turns the front wheels will track outside the rears. Consequently, the front wheels travel a few more revolutions per mile more than the rears, even if the effective radii of the tires are equal.

A spool or completely locked differential drives both output shafts at the same rpm, and does not split the torque in any fixed proportion. This is opposite to an open differential, which controls relative torque at the output shafts but not relative speed. With a spool, torque distribution depends on relative resistance at the two output shafts. It is quite possible for one output shaft to have negative resistance (wheel dragging and trying to drive the axle), while the other output shaft has a torque greater than the sum of the two (wheel driving the car plus overcoming drag from the other wheel). The former condition exists on the outside wheel, and the latter on the inside wheel, when making a turn with a spool and no tire stagger.

A partially locking or limited-slip differential is midway between. It allows some difference in speed, but adds torque to the slower output shaft and takes that torque from the faster output shaft.

A viscous coupling transmits torque according to the amount of slippage at the coupling. The faster the input shaft turns relative to the output, the greater the torque at the output shaft. Unlike a gear set, however, the relationship is usually not a simple linear function of the rpm ratio.

Note that none of these alternatives split power equally. No known passive mechanical device does that.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SHOCKS WANTED FOR RESEARCH

In previous newsletters, notably August 2000 and December 2001, I have discussed acceleration sensitivity of shock absorbers. (Acceleration is the rate and direction of change in velocity.) I noted that just looking at the difference between the extended end of the stroke and the compressed end of the stroke in a standard sinusoidal shock dyno test will give you a crude indication of a damper's acceleration sensitivity. If the two ends of the stroke look substantially different, that suggests a high degree of acceleration sensitivity.

It is reasonable to suppose that differences in acceleration sensitivity are a big part of the reason why shocks that generate similar traces in the most common dyno test (sinusoidal motion produced by a crank, 2" stroke, 100 rpm) can act so different on a car.

I have for some time been interested in investigating the matter more systematically. I would like to come up with test procedures that will give us a way to measure and quantify sensitivity to acceleration, and also investigate the importance of jerk sensitivity. (Jerk is the rate and direction of change in acceleration.)

It now looks like I'll get the chance. At least one, and quite possibly two, shock dyno manufacturers are interested in working with me on this. At least one up and coming damper manufacturer has expressed interest in building shocks for test. At least one racing team has expressed interest in working with us. That's enough to get started.

What I'm looking for now is additional teams actually running cars, who are interested in collaborating on this. As things stand, teams won't pay us, and we won't pay them. Teams will furnish shocks they already run, and/or experimental ones, for dyno testing, and provide us with feedback regarding how the various shocks affect their car. We are particularly interested in obtaining shocks that are reported to dyno similarly but act different. The team gets free shock dyno testing and a better understanding of how shocks work. We get free shocks to test, and a better understanding of how shocks work. At some point, perhaps I will write a feature article in *Racecar Engineering*, and get some publicity for the consulting business. Teams interested in being on the inside of this cutting-edge effort are invited to contact me at the address, e-mail, or phone at the top of this page.

TIRE WARMERS FOUND

Back when the forum on www.RacecarTech.com was running, somebody asked me where to get tire warmers for their dirt Late Model. At that time, I pursued several leads, but they all ultimately came up dry.

Finally, at the PRI show last month in Indianapolis, I found a US source. They are Chicken Hawk Racing, at 249 Hapeman Hill Rd., Red Hook, NY 12571. Their phone is 866-HOT-TIRE (468-8473). They have a website at www.chickenhawkracing.com.

For those unfamiliar with tire warmers, they are basically high-temperature electric blankets that wrap around a tire's tread and heat it up. Apart from the obvious advantage of giving you sticky rubber right from the green, tire warmers also allow you to heat and cool the tires gently, and keep them hot between runs. This reduces the effects of heat cycling, keeping the rubber soft longer. Additionally, they allow you to set your "cold" pressure at a controlled temperature, rather than ambient. This temperature can be high enough to assure that the tire won't have any significant liquid water in it. Regular readers may recall that water in a tire does not cause any unusual pressure rise if it's in the vapor state when you set the pressure.

Why wouldn't you use them? First of all, many sanctioning bodies and tracks have outlawed them as a cost-containment measure. And they aren't cheap. Chicken Hawk sells two models, one for around $1500 and one around $2000. That's each, and you need at least four for a car (they make them for motorcycles too). The less expensive model has a pre-set thermostat, ordinarily 175deg F (80deg C). The more expensive model has an adjustable thermostat, and a digital thermometer so you can see if the tire's up to temperature yet.

Whether the performance gain is worth the money depends on your personal situation, but the performance gain is real.

SPRING PLACEMENT ON TRIANGULATED 4-LINK

I have a question on rear spring placement on a stock metric 4-link suspension. I have built several chassis and have been mounting the rear spring centerline forward 2½ inches of the centerline of the axle. I've started on a new chassis and thought I would go back to mounting the spring directly on the axle centerline. Since the housing does not rotate under power I don't feel I'm gaining anything. Does mounting the spring forward of centerline affect the static rear percentage or in any way change the motion ratio of the spring?

US oval track racers will need no introduction to this type of suspension. For readers unfamiliar with it, this is what is sometimes called a triangulated 4-link, or Chevelle-style 4-link. It has been used on various GM cars, including the "metric" series referred to here, and also recent Mustangs. It is illustrated on p.648 of Milliken and on p.260 of Gillespie. It uses four angled trailing links to locate a beam axle both longitudinally and laterally, with no Panhard bar, Watt linkage, or other purely lateral locating device.

In most such layouts, the side-view geometry gives a substantial amount of anti-squat. The axle does rotate with ride motion, nose-down in bump and nose-up in droop. However, the only rotational *compliance* with drive torque comes from flexure of the parts, mainly the bushings. When the questioner here says the axle does not rotate, he means that there is no highly compliant torque absorbing device such as a torque arm or pull bar incorporating a spring.

In roll, there is little or no axle housing rotation.

The location of the springs has no effect on static rear percentage, except that the mass of the springs is positioned slightly further forward or back. Spring location fore-and-aft does affect motion ratio a little bit in ride. Moving the spring forward makes the spring-to-wheel motion ratio slightly less than 1:1 in ride. In roll, the motion ratio is the same as it would otherwise be, assuming the lateral spring spacing is unchanged. Note that this motion ratio in roll is always less than 1:1 for any beam axle, which means that any beam axle without an anti-roll bar has a substantially softer wheel rate in roll than in ride.

So on a stock metric suspension, moving the springs forward softens the wheel rate in ride somewhat, without softening it appreciably in roll. This makes the ride and roll wheel rates less unequal. However, if the spring is moved forward only 2½ inches, that will have only a small effect.

Note that we are speaking here of springs (on buckets, on coilovers, or on sliders) mounting directly to the axle, not to a link or a birdcage.

Even in cars with compliant torque arms or pull bars, mounting the springs forward of the axle does not add a lot of rear jacking, and rear jacking only adds total rear wheel loading due to the overall vehicle CG being slightly higher when accelerating forward. Such effects tend to be small.

Remember that jacking up both rear corners does not increase rear percentage, in and of itself. Remember also that jacking one rear corner up more than the other also doesn't significantly change rear percentage, but it does change diagonal percentage.

Correspondingly, fairly significant effects in torque-compliant axles can result when the fore-and-aft spring offset differs on the right and left, as when the left spring is ahead and the right spring is behind. Then there can be a meaningful change in instantaneous diagonal percentage as power is applied. This in turn will affect the car's cornering balance under power.

ROLL CENTER WITH A J-BAR

Many books, forum posts, and websites go into great detail on on the front roll center and only touch on the rear. I run an IMCA modified with a j-bar [short, off-center Panhard bar, bent into a J shape to clear the pinion snout – usually mounted to the left side of the frame and the right side of the pinion snout, with the left pivot somewhat higher than the right]. I would like to determine where my rear roll center is.

This is actually a fairly complex question. First let's discuss **what** a roll center is, and isn't.

A roll center isn't a real thing. It's a modeling construct – an invented idea that helps us think and talk about the suspension's behavior. It's a way of representing the geometric roll resistance of a front or rear wheel-pair suspension system, to simplify prediction of wheel loads when cornering. In the simplest method of modeling wheel load changes due to lateral acceleration, the suspension is imagined as a beam axle (which yours actually is), and the roll center describes a **height** at which **lateral** force is transmitted between the axle and the sprung mass.

It is vital to recognize that we are not talking about a point the car actually rotates around, or a point whose lateral location determines how vertical forces react. The roll center is best thought of as a point in a side view of the car, that has no defined lateral location at all, or perhaps as a point in the same longitudinal plane as the sprung mass CG. In other words, we should imagine the roll center as the height of a pin in a vertical slot, or the height of a horizontal Panhard bar, not as a pin joint. It is a notional device that transmits **horizontal force only**.

Okay, now with a Panhard bar that's curved, offset, and inclined, how do we assign that imaginary point to get the best wheel load prediction? There are two answers to this, depending on how much work we want to do, and how accurate we want our model to be. In both methods, we disregard the bend in the bar, and think of it as a straight link connecting its two pivot points.

In the simpler method, we find the point where the centerline of this imaginary straight bar intersects the longitudinal CG plane, and take that point's height as the roll center height. With this method, we disregard effects due to the off-center, inclined Panhard bar jacking the rear of the car up or down.

In the more rigorous method, we take the midpoint height of the imaginary straight bar as the roll center height. We then must also take into account the vertical forces resulting from bar inclination. We likewise consider these as acting at the bar's midpoint. When the left pivot is higher than the right pivot, in a left turn the jacking force tries to raise the sprung mass. When the bar centerpoint is left of the sprung mass CG, this effect tries to roll the car rightward, reducing effective roll resistance at the rear. So we have a higher roll center, suggesting more roll resistance, but also a pro-roll moment from the jacking. Net result will be similar to the load transfer predicted from the lower roll center in the simpler method, though not exactly the same. (As described here, both methods have some inaccuracy due to the bar being forward of the axle. Correction for this is possible, but beyond our scope here.)

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SHOCK RESEARCH UPDATE

Last month I announced that I am undertaking a project to explore sensitivity of suspension dampers to acceleration, jerk, and perhaps other factors. I am still looking for shocks to test, particularly groups of two or more shocks that produce similar results in typical crank dyno tests but act different on-track.

I do have some preliminary feedback. One shock company tested one of their dampers at the usual 2" stroke and 100 cpm (5/3 Hz), then at 1" stroke and 200 cpm (10/3 Hz). These two tests produce equal ranges of velocity, but at any velocity the second test produces 2 times the acceleration and 4 times the jerk (change of acceleration). In this case, the forces were identical in both tests, within the window of accuracy attainable. This doesn't mean the test was a failure. It means that the shock tested is insensitive to acceleration and jerk, at least at the values tested.

I think it is quite possible that many shocks are acceleration-insensitive. My object is to devise ways of systematically testing to find out, and also to find out whether acceleration sensitivity can be a performance advantage if used correctly.

The shock dyno company I'm working with at this point, Performance Data Systems, also tested two shocks that were provided to them by a different shock company, which were identical except for gas pressure. PDS has a unique dyno design that allows unusually precise motion control, and will follow almost any desired motion pattern, since it uses a linear motor rather than a crank. One test this dyno can do is a step test: the shock is rapidly accelerated to one velocity, held at that velocity for a given distance, then accelerated abruptly to a higher constant velocity, held at that velocity for a specified distance, and so on.

The acceleration zones between one step and the next can be programmed to have defined limiting values for acceleration and jerk. The machine can also be programmed to reach a particular peak acceleration with either maximum jerk at the ends of the acceleration zone, or minimum jerk for a desired mid-zone acceleration, within a specified acceleration time between velocity steps. If jerk is set at maximum, then jerk is zero in the middle of the acceleration zone. If jerk is set at minimum, then peak jerk is much less, but there is still a non-zero jerk value in the middle of the acceleration zone. This means that this test can produce points where both velocity and acceleration are the same, but jerk is either zero or some known value. This allows isolation of jerk effects from acceleration effects, which is not possible in sinusoidal testing. Alternatively, a shock can be tested at different known accelerations, with identical velocities, and zero jerk. This allows isolation of acceleration sensitivity from both velocity and jerk effects.

In the test of the similar shocks with differing gas pressures, PDS reports that varying accelerations did not produce different forces at mid-acceleration, but varying jerk values did produce differing forces. And the difference was greater in one shock than in the other. In other words, the shocks appeared to be jerk-sensitive without being acceleration-sensitive, and the jerk sensitivity appeared to vary with gas pressure.

Stock car teams are reporting that shocks with a given piston and shim package definitely feel softer to the driver when gas pressure is reduced.

Some caveats here: I was not present at the tests I am describing. I am relying on the accounts of others. Also, we are not looking at results of an exhaustive, systematic testing program. What we do have is preliminary, anecdotal evidence that suggests there are effects worth measuring and exploring through unconventional damper testing.

MORE ON REAR WHEEL PLACEMENT AND TRACTION

Simon McBeath, whose comments regarding rear wheel placement and its effects on traction prompted my remarks in the October 2002 newsletter, writes:

I've just been catching up on some overdue reading and noticed in your October newsletter that you picked up my suggestion for a discussion on the above. Many thanks a) for doing that and b) for reading the feature (on the DJ Firehawk hillclimber) where the suggestion was placed!

I read what you had to say in your newletter with great interest. But is there also another mechanism at work with swung back rear suspension? The Firehawk's designer mentioned to me something I was very unclear about, and hence did not go into in the article, but it involved the suggestion of gyroscopic effects aiding traction, and it was in reference to buggy racing. Have you heard of this effect being exploited this way? I couldn't figure how that would work, to be honest.

I tried an experiment in the workshop with a hand held grinder, angled back, as it were, as if the grinder's disc was a wheel swung back on its suspension, and as you move such a tool around up and down you can feel gyroscopic forces, but when the tool is held still (but powered up) there are no sensations or reactive forces.

But when you first power the tool up there is, obviously, a reaction force. I wondered if this instantaneous response could be usefully exploited for improved traction - to add to the weight transfer under acceleration and make the tyre dig in harder, initially at any rate. I have a feeling as I type this that what you might gain on one side of the car you'd lose on the other, but I can't figure it out in the middle of a Sunday afternoon! Any thoughts would help still my curiosity and soothe my confused brain!

What you're feeling when you turn the grinder on is mainly the grinding wheel acting as a flywheel, not a gyro. The body of the grinder is more analogous to an axle housing than a semi-trailing arm in a buggy rear suspension, because drive torque reacts through the grinder body. The arm on the buggy only reacts thrust under power. Drive torque reacts through the powertrain mounts, and does not act through the suspension.

Wheels on a car do produce gyroscopic forces, but only when their toe or steer angle changes, or their camber angle changes. Rotational acceleration or velocity about the wheel's main axis (axle axis) does not produce gyroscopic forces. When we steer the wheel to the left, it tries to lean to the right. When we lean the wheel to the left, it tries to steer to the left. These effects are called gyroscopic precession.

The precession force depends on the wheel's angular velocity in the plane perpendicular to the force. That is, when the wheel steers left, the magnitude of the rightward camber-wise or roll torque about the wheel-longitudinal axis depends on the wheel's velocity (not acceleration, not position) about the vertical axis. The wheel's rotational speed on its axle also matters. More rpm, more precession force; wheel not rotating, no precession. Lastly, the wheel's moment of inertia about the axle axis matters. More flywheel effect, more precession force.

In a motorcycle or bicycle, precession forces are an important factor in vehicle behavior. We use them to hold the vehicle upright, and to steer it. But in a tricycle or a car, we just live with these forces; we don't harness them. If anything, they're a problem, because they are part of the reason for shimmy in steering systems.

With the grinder, you are holding the device by the body, which is not quite in the same plane as the disc. Consequently, the grinder may try to move in a complex manner when you power it up. It may try to tilt the disc as well as rotate the body about the spindle axis. If the disc tilts, then there will be some gyroscopic precession.

In any case, gyroscopic precession does not increase traction.

As for transient (short-lived) forces that try to lift the car momentarily increasing traction, that's possible. However, the brief traction improvement is followed by a corresponding unloading of the wheel a fraction of a second later. What counts for this is the **vertical acceleration** (not position, not velocity) of the sprung mass ($F = ma$). The sprung mass is lifted a bit, but only to a point. So its velocity upward increases to some value, and then decreases to zero again. That means its acceleration is first upward, and then downward. When the sprung mass acceleration is upward, there is a wheel load increase. When the sprung mass acceleration is downward, there is a wheel load decrease.

It's probably better for traction not to have such an effect. In certain instances, the driver may be able to time the momentary traction increase to occur when it's most needed, but in general the car is limited by its instants of poorest traction, rather than its instants of greatest traction. Therefore, we would like the wheel loads to vary as little as possible.

There is also another effect when the car is being carried in a lifted position: the center of gravity is higher, and that increases rearward load transfer. So anti-squat does improve traction, but not as much as many people imagine.

Now, if you move the rear wheels back on a buggy, what happens to the anti-squat, and other properties? The answer depends on what type of rear suspension the car has, and exactly what you change to move the wheels back. Traditionally, buggies have semi-trailing-arm rear suspension, derived from the design on late VW beetles. However, this is not always the case any more.

Assuming we have semi-trailing arms, there are a number of ways the wheelbase could be lengthened, and the various methods have different effects on the rear suspension geometry. Probably the simplest method, on an existing car, would be to merely fit longer arms, without modifying the frame. If we do this, we get the following effects:

1. The static rear percentage decreases. As previously noted, this hurts traction.
2. The static anti-squat diminishes, assuming the trailing arm slopes up toward the front.
3. Changes in anti-squat with suspension motion are reduced, because of the longer side-view swing arm.
4. Changes in camber over bumps are reduced, due to the longer front-view (rear-view, end-view) swing arm. Also, there is less bump steer.
5. The rear roll center is lower.
6. In all likelihood, the suspension will be softer with a given spring and shock package, due to a decreased spring-to-wheel motion ratio.

The last five of these effects could all improve traction, especially while cornering, and on bumpy surfaces. This might account for perceived or reported improvements. Note, however, that all of these effects could also be achieved by moving the pickup points forward, and leaving the wheel location unchanged. That would probably involve redesigning the frame, of course. And a better approach yet is to forget about using semi-trailing arms altogether, and build a proper five-link, or short-and-long-arm, suspension.

If we do that, we can have any rear geometry we want, with any wheel location, and we can have much less variation in anti-squat than with any semi-trailing arm system. And any arguments for moving the wheels back that might apply with semi-trailing arms become irrelevant.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

FREE SEMINAR MARCH 17

For readers in the Charlotte area, I will be presenting a one-hour talk at UNC Charlotte, in Fretwell 126, Tuesday, March 17, at 12:00 noon. The title is "Minding Your Anti: Understanding Roll Centers, Jacking Forces, and Other Factors in Weight Transfer". This will not be a purely standard treatment of the subject. I will include discussion of "lateral anti" and shed important light on the much-discussed topic of lateral roll center location and migration. This is a free presentation of the UNCC student SAE chapter.

TIRE WARMERS LESS EXPENSIVE THAN I THOUGHT

In the January issue, I mentioned I'd found tire warmers, supplied by Chicken Hawk Racing , 866-HOT-TIRE or www.chickenhawkracing.com. I said they had a standard model for around $1500 and an adjustable one with temperature readout for around $2000. I was under the impression that those prices were for a single warmer, but actually they're for a set of four.

INDEPENDENT REAR SUSPENSION FOR DIRT?

Would there be any advantage to running an independent suspension on the rear of a dirt car? This refers primarily to a modified, but would it help on a dirt Late Model, also? We were wondering if a design similar to a Corvette would work.

There is no doubt that independent rear suspension can work very well on dirt. This is provable not only by theory, but by example. Independent rear suspension is used with great success in off-road buggies, rally cars, and Unlimited hill climb cars at Pikes Peak. The only place IRS isn't used on dirt is in oval track racing.

The biggest single reason for this is that in most classes, and in most sanctioning bodies, independent rears are illegal, presumably for cost containment. I'm not sure about the high-dollar mods that

D.I.R.T. runs in the northeastern US, but in IMCA, UMP, NASCAR, and WISSOTA, there are specific rules against independent or “sports car” rear ends. They don’t even allow quick-changes. Down here where I live, we have the Carolina Modified Tour, which runs similar cars, but with quick-changes permitted.

For Late Models, the rules vary. I haven’t checked all the sanctioning bodies that run these cars, but WISSOTA abolished all suspension rules in the Late Model class a few years back, after previously prohibiting independent rears. To my knowledge, everybody still runs live axles, partly so they can go on the road and race in other series, and partly because they are mostly car buyers, not builders, and no Late Model builders offer independent rears.

Far as I know, all sprint car and midget sanctioning bodies, including World of Outlaws, now require beam axle suspension front and rear.

So the first obstacle to overcome is to find a sanctioning body that will let you run IRS. You have to think about not only what the current rules are, but also how the organization is likely to react if you are successful with an independent rear, and everybody else faces the prospect of having their cars obsoleted. You will have to invest a lot of time and money in building your own car and developing it. If it’s outlawed as soon as it starts winning, you take a big loss.

You will face another problem that besets all innovative owner-builders: when you tear up equipment, you can’t just order replacement cars or parts; you have to make them. If you have a heavy schedule and are running for points, or you’re on tour, this is a major concern.

Twenty years ago, there were some attempts to build independently suspended sprint cars. These efforts were mainly the work of backyard builders, who had little formal training. They attempted to build systems that looked like what they’d seen on road racing cars of the era, with little real understanding of what they were doing. I recall one case where the builders didn’t realize they’d still need tire stagger, and blamed the suspension when the car went straight into the wall the first time they ran it.

The lesson here is that a mediocre concept, executed and set up well, will beat a superior concept, executed or set up poorly. Independent suspension has the potential to win races on dirt, but only if it’s done right. Since you’ll be pioneering a new idea, you won’t be able to rely on conventional wisdom; you’ll have to study sufficiently to understand the principles of the system. I will be happy to help you as a consultant, but those actually doing the project will need considerable knowledge as well.

Okay, assuming you aren’t daunted by the practical aspects of trying something radical, and assuming you’ve found a class where IRS is legal, what are the pros and cons of IRS, and what sort of design would be best?

Independent rear suspension is good, but it is a mixed blessing in some respects. In general, overall weight is greater for independent suspensions than for beam axles. However, unsprung weight is much less for an independent suspension, especially if the brakes are inboard, and most Late Models run to a minimum weight rule that requires them to add ballast. So in terms of weight, the only drawback to IRS is that you have somewhat less ballast to move as desired. There is a big benefit in roadholding, meaning ability to keep the tires in contact with the track, and minimize tire load variation, on bumpy surfaces – and dirt tracks are often bumpy, though not always.

Anti-squat in independent systems is different than in live axles. In a live axle system, we can separate rear jacking forces under power into thrust anti-squat and torque anti-squat. In a typical Late Model, torque anti-squat is the lift we get from the torque arm, and thrust anti-squat is the lift we get from the geometry of the linkages at the ends of the axle, which most commonly attach to birdcages (brackets that can rotate on the axle). With independent suspension, we only have thrust anti-squat to work with, because axle torque reacts through the differential mounts and does not act through the suspension.

This leads some people to suppose that overall anti-squat is necessarily less with independent suspension, and that therefore independent suspension would be at an inherent disadvantage compared to current state-of-the-art dirt Late Model live axles. I question this myself, although I do agree that in theory at least, a live axle can probably be made to lift more under power than an independent system. As I have mentioned at various times in the past, the advantages of anti-squat are often over-estimated, and it is possible to get ample lift from an independent system.

It is safe to say that the live axle has some edge in terms of anti-squat properties, particularly as regards the potential to manage variation in anti-squat properties as grip varies. However, current systems do not exploit the possibilities in this area as fully as they could, so this potential advantage of the live axle is hypothetical until somebody decides to exploit it. These possibilities might be a future newsletter topic.

Compared to current live axles, an independent system could have similar, or at least adequate, anti-squat, and much better adhesion over bumps. The independent system might reasonably be expected to compare most favorably on a bumpy track, and least favorably on a smooth and slippery one.

For Late Models, there are rules about transmissions, at least in WISSOTA. They have to be mounted to the engine, so transaxles are out. That means the diff would be an IRS quick-change, with either a spool or a Gleason.

You mention Corvette rear suspensions. There are two basic styles of independent suspension used on Corvettes. The C2 and C3 used the halfshafts as upper lateral or camber-control links, a lateral link sometimes called a strut rod below the halfshaft to complete the camber-control linkage, and a trailing arm for toe location, longitudinal location, and brake torque reaction. A variation of this system, with a third lateral link for improved toe control near the front of the trailing arm, was used in second-generation Corvairs.

The C4 and C5 Corvettes have a 5-link system. There are three transverse links to control camber and toe, and two longitudinal links to provide longitudinal location and react brake torque. On the C4, the halfshaft is still used as the upper camber control link. On the C5, the model currently in production, the halfshaft is only used to transmit power, and the five links are all purely suspension parts. Similar 5-link systems are used on the Viper and most purpose-built race cars. On some current race cars, two pairs of links are combined into upper and lower a-arms, with a toe-control link. The system then visually resembles a front suspension.

If I were designing an independent rear for any form of racing, including dirt oval-track, I would use a five-link system, or the a-arm and toe-link variation of the 5-link. Using the halfshafts as camber control links saves a little weight and cost, but it compromises geometry. Specifically, it forces you to choose between a high roll center or meager camber recovery in roll. Also, the consequences if you break a shaft or U-joint are particularly nasty, though of course they aren't pleasant regardless.

One key decision is whether to use inboard or outboard rear brakes. The advantage of inboard brakes is that you reduce unsprung mass, and thereby maximize the system's roadholding advantage on bumps. The advantage of outboard brakes is that you can have ample anti-squat under power, without having excessive anti-lift under braking. A lot of anti-lift in braking tends to cause wheel hop when used with generous rear brake bias, and many dirt drivers like to use a lot of rear brake to get the car to turn in. Typical 4-bar Late Model rears have more than 100% anti-squat and zero or negative anti-lift.

To get such properties with an independent system, you need geometry that makes the hub travel rearward approximately .15" to .20" per inch of suspension compression, and makes the upright rotate rearward approximately 0.6 to 1.0 deg per inch of suspension compression. In terms of side view geometry, this means a side view instant center something like 80" behind the rear wheel, and at or slightly above ground level.

That's with outboard brakes. With inboard brakes, you'd want the hub to move rearward no more than .10" per inch of suspension compression, unless the driver never uses a lot of rear brake. Upright rotation doesn't matter with inboard brakes. Probably the simplest approach would be to make the whole upright move along a line inclined about 5 deg rearward, and not rotate at all.

In either case, I'd consider having a bit more anti-squat on the left than on the right, to make the car gain wedge under power.

For lateral location, I'd try instant centers between 70 and 100 inches from the wheel and try to keep the force line slopes between zero and 10 degrees, upward toward the center of the car, in all combinations of ride and roll. This would correspond to a static roll center height of 3 inches, give or take an inch. I would try to make the lower control arms as long as possible – all the way in to the center of the car if possible – and have the upper arms shorter than the lowers by as much as needed to achieve least possible force line slope changes in both ride and roll, with perhaps a bit more emphasis on roll than ride.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April/May/June 2003

My apologies to anybody inconvenienced by the lateness of the April and May issues,
and my thanks to those who have kept me so busy lately,
especially those who have done so on a paid basis!

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SHOCK AND SPRING FORCES

With the increased emphasis on tuning using shocks, could you explain how the shock absorber forces are fed through the chassis to affect the tire loads? How do these forces differ from the forces that are transmitted through the springs and sway bars?

To discuss any subject, we need a vocabulary. So first, let’s define some terms.

Car wheels move in three dimensions, but we can simplify and think of the suspension as mainly just letting the wheel go up and down. Viewed this way, the suspension for a particular wheel can move in two directions: compression and extension. We also sometimes use the terms bump and droop or rebound for these. As with any motion, the system can be said to have a position, a velocity, an acceleration, and a jerk at any instant we may choose to look at.

Position, or displacement, is the inches (or millimeters) of compression or extension from some designated reference or zero point. Usually, we take the static position as this zero – the position the suspension is in when we set the car up on the scales.

Note that suspension position or displacement (for one wheel) can be expressed as a single number. (Position or displacement of the sprung mass requires six numbers to completely express it: three for linear position along three axes, and three for angular position about those axes. Rotation about a longitudinal axis is roll; rotation about a transverse axis is pitch; rotation about a vertical axis is yaw.)

The suspension’s position may be one-dimensional, but it is still a *vector* quantity: it has a magnitude, and a direction – so many inches or millimeters, compression or extension.

The position or displacement can change over time. This change of position with respect to time is called *velocity*. It likewise has a magnitude and a direction – so many inches or millimeters per second, compression or extension.

The velocity can change over time. The change of velocity with respect to time is called *acceleration*. Acceleration has a magnitude and a direction – so many inches or millimeters per second per second (in/sec^2 or mm/sec^2), compression or extension.

The acceleration can change over time. The change of acceleration with respect to time is called *jerk*. Once again, this is a vector quantity – so many inches or millimeters per second per second per second (in/sec^3 or mm/sec^3), compression or extension.

Readers who've had calculus will recognize that we are taking a series of derivatives here. Velocity is the first derivative of displacement, with respect to time. Acceleration is the second derivative of displacement and the first derivative of velocity. Jerk is the third derivative of displacement, the second derivative of velocity, and the first derivative of acceleration. We could go on taking derivatives indefinitely (the next one is called *quirk* – no, I'm not making this up), but the usefulness of doing so is doubtful.

It definitely is useful to look at longitudinal and lateral acceleration and jerk. As we will see, the car's accelerations determine the suspension's displacements, and the changes in the car's accelerations (its jerk values) determine the individual suspensions' velocities. And since the suspensions' velocities are not constant, the suspensions likewise have non-zero acceleration and jerk values.

Those with engineering backgrounds may feel I'm belaboring the obvious in this discussion. However, I have recently been recruiting volunteers to collaborate in a project to test sensitivity of dampers to suspension acceleration and jerk, and I have found that many readers and clients have not understood what I meant by acceleration or jerk as applied to a shock or a suspension system.

We can express the direction of any of these quantities with a mathematical sign – positive or negative. Which way should we do this? Should positive be compression or extension? It is customary to use positive for compression displacement in data acquisition, so maybe that's the way to go. Calling compression positive also more or less agrees with the conventional tire axis system, in which normal or road-vertical force is considered positive. In general, we think of increased suspension compression displacement and increased normal force as going together, although this is not always so. Then again, compression displacement generally implies extension force within the suspension, so if extension is positive, then positive force in the suspension corresponds to positive force at the tire, at least as a crude generalization.

In shock dynamometer testing, the dyno manufacturers have to establish sign conventions within their own software. In the software for Roehrig dynos, the most popular make in the US, compression strokes have positive force (that's extension force, resisting the compression motion) and negative velocity. This agrees with the extension-positive reasoning above. However, it disagrees with the conventions generally used in data acquisition. No matter what we do, we will either disagree with the shock testing convention or the data acquisition one, since they disagree with each other.

If this were not confusing enough, even these sign conventions are not universal, as we will see shortly.

One might think we could avoid all confusion by dispensing with signs and simply stating direction with the word “compression” or “extension”. This works fairly well when we are talking about displacement, velocity, acceleration, and jerk direction. However, some confusion arises when we discuss damping force. It is customary to speak of compression (or bump) damping force as the force occurring when the shock is compressing, and extension (or rebound) damping force as the force occurring when the shock is extending. But in fact these forces ordinarily are opposite in direction to the shock’s velocity: compression damping force acts in the extension direction; extension damping force acts in the compression direction. So substituting words for mathematical symbols is no refuge. Either way, we have to keep in mind the actual physical phenomena we’re trying to describe, and apply some common sense, to avoid confusion.

When a number or quantity has a positive or negative sign, we may speak of its *absolute value*. A quantity’s absolute value is the greater of the quantity and its opposite. The absolute value of 4 is 4. The absolute value of -4 is 4. ($|4| = 4$; $|-4| = 4$)

Correspondingly, when we have a unidimensional vector quantity such as suspension displacement, velocity, acceleration, or jerk, whose direction can be expressed by a positive or negative sign, we may speak of the quantity’s absolute value. This means the quantity’s magnitude, irrespective of direction. So, for example, when we speak of large **absolute velocities**, that means large compression or extension velocities. When we speak of large **velocities**, on the other hand, that means large velocities in whatever direction we call positive.

In casual conversation, these distinctions are often disregarded, so again we face the need to apply common sense, and understand people’s words (and also their math symbols) in context.

A familiar synonym for absolute velocity is *speed*.

The most common type of shock dyno plot is force (vertical axis of the graph, forces resisting compression positive, forces resisting extension negative) versus absolute velocity (horizontal axis, all values positive). Also available is force versus velocity. Here it is customary to show velocity as negative in compression, and force resisting compression as positive. The trace is generally S-shaped, and lies mainly in the second and fourth quadrants. These sign conventions are opposite to the compression-positive convention used in data acquisition, but they do show a realization that the velocities and the most common forces should have opposite signs.

That’s with most dyno software I’ve seen. A correspondent in Australia recently sent me force-versus-velocity traces from an SPA dyno, in which the compression stroke has both velocity and force positive, and the extension stroke has both velocity and force negative. The trace then lies mainly in the first and third quadrants of the graph. With this choice of sign convention, velocity agrees with the data acquisition, but forces acting opposite to velocity are shown with the same sign as velocity.

Note that I refer to the **most common** forces associated with a particular velocity direction, rather than all the forces. A true damping force acts in opposition to motion – otherwise it wouldn't suppress motion. However, not all the forces our dampers generate are actually damping forces in this sense.

If you examine shock dyno plots, you will see that sometimes shocks generate forces in the same direction as velocity. There are at least three known phenomena at work here, and perhaps additional ones. The first known phenomenon is gas spring effect. In gas pressure shocks, the gas compartment acts as a rising-rate spring. The smaller the gas volume, and the higher the pressure, the greater the gas spring rate. The gas spring force always acts in the extension direction. So when the shock is moving very slowly in extension, it exerts a net extension force.

If we look at a force vs. absolute velocity plot of the full stroke from a crank dyno, there will be two *noses* or points at zero velocity, representing the extended and compressed ends of the stroke. In a gas-pressure shock, the compressed end of the stroke will show a higher force reading (meaning more extension force) than the extended end. If the dyno is cycled very slowly, and the shock has very soft low-speed valving – especially if it has bidirectional bleed – the difference between the two noses will be almost entirely from gas spring effect.

It is customary to zero the dyno, and omit gas spring force from the force reading, at some point in the cycle – typically the extended end of the stroke, although mid-stroke and full-compression zeroing are also common. Even when this is done, the force reading will be higher at full compression than at full extension. Thus, the shock will either show an extension force early in the extension stroke, or a compression force early in the compression stroke, or both, just from gas spring effect.

There is a second known effect that will cause the noses to spread further apart as low-speed valving is stiffened, and as the shock is cycled at higher frequencies or longer strokes. This effect is fluid compressibility.

Suppose we have a shock with the body sprung, mounted body-up. As the shock nears the end of the extension stroke, the fluid below the piston is under substantially greater pressure than the fluid above the piston. If it is not allowed to bleed off very rapidly, it will still be under pressure as the piston comes to rest and starts to move upward. Consequently, the shock will not resist compression until it is some distance into the compression stroke. This effect is sometimes called *lag*. If the fluid were perfectly incompressible, this couldn't happen. Pressure would equalize instantaneously as soon as velocity reached zero. But shock fluid has substantial compressibility, despite our efforts to reduce this.

Thus, the fluid itself will act partially as a spring rather than a damping medium. And until there is greater pressure above the piston than below it, there will be no flow downward through the piston, and therefore there will be no extension force damping the compression. Indeed, as long as pressure is greater below the piston, fluid will try to flow upward through the piston.

Lag occurs at both ends of the stroke.

Lag is somewhat distinct from acceleration sensitivity, but it does relate to accelerations, especially when the velocity is changing sign.

A third known effect comes from the masses and inertias of the valving elements. With deflective-disc valving, these effects are generally small. Where a valving element of considerable mass acts against a spring, the effects can be large. Shocks made by Ricor and sold under the Edelbrock name make deliberate use of this effect and advertise it as a selling point. These shocks use a weighted element on top of the piston, which softens the valving when acceleration is in the extension direction, i.e. during the more compressed half of the stroke, or the compression closing/rebound opening (cc/ro) portion. I have also seen a patent description for a shock with a similar weighted element under the piston, to soften the rebound closing/compression opening (rc/co) portion.

Intentional acceleration sensitivity usually reduces the forces generated by the shock – although it could be made to increase them – when acceleration is in a particular direction, and sometimes only when a particular combination of acceleration and velocity directions is present. Since lag is related to valving stiffness when the velocity is changing sign, acceleration sensitivity affects lag.

Acceleration sensitivity affects force whenever relevant accelerations are present, not just near velocity reversals. Thus, although acceleration sensitivity affects lag, it is a distinct phenomenon.

Acceleration sensitivity is not necessarily bad for car behavior, and may in fact be beneficial when intelligently applied, but it’s a complication in terms of modeling or understanding.

So shock forces are complex. Sometimes our dampers create spring forces. Sometimes we can’t predict their behavior just by knowing their velocity.

Spring forces can also be complex. In leaf springs especially, there is damping in the spring, mainly from inter-leaf friction. In some large vehicles, with many leaves in the springs, this effect provides all the damping; there are no shocks. Even coils and torsion bars have some internal hysteresis. They will heat up as they flex, and they will come to rest after a number of oscillations, even in the absence of external damping.

There is also friction in all the pivots in the suspension and steering, and there is friction in the sliding contacts in the shocks.

So we get some damping forces from our springs, and from other components in the system.

I am not trying to confuse matters here. I merely wish to point out that the remarks which follow are based on simplifying assumptions, rather than comprehensive models of spring and shock behavior.

In understanding how shocks and springs affect wheel loads, we think of springs as being exclusively sensitive to position or displacement. We estimate their forces on the basis of their displacement. We think of shocks as being entirely velocity-sensitive. We assume that they always make compression forces when they are extending, and extension forces when they are compressing. We assume that if the shaft speed is greater, the absolute force is greater, in some predictable relationship, though usually not a linear one.

An anti-roll bar is an interconnective spring. It generates forces based on its displacement, but its displacement depends on the relative displacement of the two wheels it connects, rather than their individual displacements. It generates equal and opposite forces in the two suspensions it connects.

It is useful to divide suspension forces affecting wheel loads into the forces present at static condition (as the car stands on the scales at the conclusion of static setup) and the forces that add to or subtract from these static forces as the car runs. At static condition, all suspension displacements from static are zero, suspension velocities are all zero, suspension accelerations are zero, and suspension jerks are zero. The anti-roll bar or bars may have preload or may not. In a road racing car they usually will not.

The springs, shocks, and anti-roll bars act on the suspension **in parallel**. Although these elements may act through different motion ratios, each of them can be thought of as exerting a particular force at the wheel at any given instant, and these forces can simply be summed (with proper attention to sign) to arrive at the resultant effect. A 400 pound extension force from the spring, with no force from the a/r bar and the shock, is equivalent to a 600 pound extension force from the spring, countered by a 100 pound compression force from the a/r bar and a 100 pound compression force from the shock (all as measured at the wheel). Or, either of these is equivalent to a 200 pound extension force from the spring, and a 200 pound extension force from the shock, and no a/r bar present. The tire doesn’t know the difference. It only responds to the total force spreading the wheel away from the sprung mass.

With springs and a/r bars, we have a *wheel rate*. That’s the rate of the spring, or the bar at the lever arm end, times the **square** of the spring-to-wheel, or arm-end-to-wheel, motion ratio. The wheel rate defines a simple relationship between force and displacement. Using the wheel rate, we can calculate the spring and bar forces at the wheel when we know displacement at the wheel.

A shock doesn’t have a wheel rate in the sense that a spring does, because it is not a displacement-sensitive device. To find shock force at the wheel, we need to know either velocity at the wheel, or velocity at the shock. If we are working from data acquisition outputs, often the sensor will be set up to read shock motion one-to-one, or as nearly so as practicable. If we are calculating from an assumed or predicted suspension motion, or from photographic data, we may be working from wheel motion. To calculate shock force at the wheel from velocity at the shock, we first estimate the force the shock generates at that velocity, based on dyno testing, then multiply this by the **first power** – not the square – of the shock-to-wheel motion ratio. To calculate shock force at the wheel from velocity at the wheel, we first find shock velocity by multiplying wheel velocity by the shock-to-wheel motion ratio, then proceed as before: estimate shock force, multiply by motion ratio. So we do multiply by the motion ratio two times in this process, but in between, we estimate the shock force.

A complete dissertation on all possible wheel loading effects from springs and dampers is beyond our scope here, but let's consider a simple case: a turn on a level, smooth surface. We will assume that the road has no small-scale or large-scale irregularities – billiard-table flat, no hills, no crests, no dips, no banking. We will also ignore aerodynamic effects. This means that the sum of our four wheel loads is the same as we'd see in the shop while doing our setup on the scales. It also means that any change in the distribution of those wheel loads is entirely the result of the way the suspension transfers weight or wheel loading in response to horizontal forces generated by the tires. This lets us isolate these effects and look at them.

We will also assume that the suspension generates no extension or compression forces due to linkage geometry: no anti-roll or pro-roll, no anti-dive or pro-dive, no anti-squat or pro-squat, no anti-lift or pro-lift. This is actually impossible to achieve for all conditions of suspension motion, and it wouldn't be desirable, but we can imagine it, and it is not too far from the actual properties of current four-wheel-independent chassis. This simplifying assumption lets us focus on wheel load changes from the springs and dampers.

As when the car is on the scales, **an increase in positive (meaning extension) force at one corner of the car adds wheel loading at that corner and the diagonally opposite one, and correspondingly reduces loading on the other two corners of the car. And a negative (compression) force reduces loading at that corner and the diagonally opposite one, and adds load at the other two corners.** This is true regardless of whether the force is generated by the damper, the spring, or the a/r bar. The tire doesn't know which part does what. It only behaves according to the resultant loading generated by the suspension elements acting together.

Unlike the static condition, front, rear, left, and right percentages do change. However, the suspension does not control these changes in this simplified case; the wheelbase, track width, and CG height – not the suspension – control how much load transfers at a particular longitudinal and lateral acceleration. The springs and shocks control how the diagonal percentage varies as all this is going on, and thereby influence the car's cornering balance. More diagonal percentage (meaning outside front wheel load plus inside rear, as a percentage of total) at any point in the cornering process adds understeer, or tightens the car. Less diagonal percentage adds oversteer, or loosens the car.

The July 2001 newsletter contained a troubleshooting chart based on five parts of a turn, with complete explanations of what the five parts were. For the convenience of readers receiving the newsletter by e-mail, I am sending that back issue with this one, as a reference. Readers seeing this issue as hardcopy can order back issues from me.

<u>Part One, or early entry – braking increasing while turning in:</u> This may or may not happen at all. In oval track racing, it is very common. In road racing, braking force more commonly reaches its maximum while the car is still running straight.

The car as a whole is accelerating rearward at an increasing rate, and accelerating laterally in the direction of the turn at an increasing rate. Angularly, it is pitching forward and rolling out of the turn. Its roll displacement is outward. Its pitch displacement is forward. Its roll velocity is outward. Its pitch velocity is forward.

Therefore, the outside front suspension has a compression displacement, and a compression velocity. The inside rear has an extension displacement, and an extension velocity. Without more information, it is hard to say exactly what the displacements and velocities at the inside front and outside rear are, but they are relatively small, because the effects of roll and pitch are subtractive at those corners.

Consequently, spring and damper changes at the outside front and inside rear will have the greatest and most certain effects on the car.

Taking springs first, the important principle is that a stiffer spring creates more load **change** with displacement change – not necessarily more load. So a stiffer outside front spring increases load at that corner (negative displacement, positive load change), and at the inside rear, and correspondingly unloads the inside front and outside rear. This adds diagonal percentage, which tightens the car, or adds understeer.

A stiffer spring on the inside rear creates a bigger load **decrease** with displacement change. That translates to less diagonal percentage, and a looser car (more oversteer or less understeer).

A stiffer front anti-roll bar creates a positive (extension) force at the outside front, and an equal and opposite negative (compression) force at the inside front. This also creates equal and opposite load changes at the rear – more load at the inside rear, less at the outside rear. Result: more diagonal percentage, tighter car (more understeer). A stiffer rear bar does the opposite, and loosens the car.

As for the dampers, if we stiffen the outside front low-speed compression valving, that adds a positive (extension) force at the outside front, adding diagonal percentage and tightening the car (adding understeer). If we stiffen the inside rear low-speed extension valving, we are creating a negative (compression) force at the inside rear. This reduces diagonal percentage and loosens the car.

Important things to note regarding the role of the dampers:
1) When the suspension velocity and the suspension displacement are in the same direction, stiffening the damper and stiffening the spring have qualitatively similar effects on oversteer/understeer balance.
2) Contrary to a very common misconception, stiffening the dampers does not slow down or momentarily reduce the load changes at the outside front or inside rear – these load changes are sped up, or are momentarily **increased**. Spring loads are momentarily decreased at the outside

front and increased at the inside rear – in other words, spring load changes are momentarily decreased by the shocks – but the effect on tire loads is the opposite.

3) If the low-speed valving is soft and the velocities are small, the damper forces may be relatively insignificant.

Also, note that:

1) We are assuming that the road is smooth. As long as this is true, the shock movements will be low-speed (less than 2 in/sec) and will be caused by sprung mass motion. When the surface is bumpy, bumps become the main factor in shock motion and none of what we've been saying about load transfer effects from the dampers applies. There still are sprung-mass-motion components to the shock motion, superimposed on the motions from the bumps. When looking at track data we can, at least to some degree, separate these components, but the shocks can't do that. They only know their actual motion at a particular instant.
2) We are assuming that the brake bias is such that the front wheels do at least half of the braking. If the car is slowed primarily by the rear wheels, the effects of diagonal percentage may reverse. This is due to the distribution of rearward force at the rear tires, and not to any fundamental difference in tire properties during entry.
3) Contrary to the contentions of some writers, tire load sensitivity (the decrease in coefficient of friction with increasing load, which is responsible for the car getting tighter with increasing diagonal percentage) does not reverse or work backwards during entry or with cold tires.

Part two, or late entry – braking decreasing, cornering force increasing: This may be the first or only phase of entry if the driver reaches peak braking while traveling straight. It is also possible for a period of "semi-steady-state" braking and cornering to exist between early (increasing) and late (decreasing) braking, particularly on ovals. More on this later.

In terms of spring and damper behavior, the difference between part one and part two is the direction of pitch velocity. In part one, the car has a forward pitch displacement and a forward pitch velocity. In part two, the car has a forward pitch displacement and rearward pitch velocity. We may say it's de-pitching; it has a forward tilt, but a decreasing one.

Roll displacement and velocity are both outward, same as in part one. However, roll displacement is increasing at a decreasing rate, which reaches zero at the conclusion of part two. So roll velocity is outward and decreasing, and roll acceleration is inward.

At the beginning of part two, the car has a combination of outward roll displacement and forward pitch displacement. At the conclusion of part two, the car has near-zero pitch displacement and increased outward roll displacement.

The biggest individual suspension displacement changes from the conclusion of part one, and the greatest individual suspension velocities, are at the inside front and outside rear. These are the wheels where the effects of rearward pitch velocity and outward roll velocity are additive.

At the beginning of part two, the displacements of the inside front and outside rear will be small. At the conclusion of part two, the inside front will have an extension displacement, and the outside rear will have a compression displacement. The velocities, therefore, are in the extension direction at the inside front, and in the compression direction at the outside rear.

Consequently, the low-speed extension (rebound) valving on the inside front and the low-speed compression valving on the outside rear are in a position to influence wheel loads. Stiffening inside front extension introduces a negative (compression) force and increases diagonal percentage, tightening the car (adding understeer). Stiffening outside rear compression introduces a positive (extension) force, which decreases diagonal percentage and loosens the car (reduces understeer).

Inside front and outside rear spring rates will be of little importance at the beginning of part two, but will be as significant as outside front and inside rear rates at the conclusion of part two. Stiffening the inside front spring will tighten the car (add understeer), and stiffening the outside rear will loosen the car (reduce understeer). As in part one, velocities and displacements at the most influential wheels have the same sign, and stiffer springing and stiffer damping have qualitatively similar effects on balance.

<u>Part three, or mid-turn – steady-state cornering:</u> Most turns, with most drivers, will include some interval of approximately steady-state cornering. This means that the driver applies just enough power to maintain or slightly increase speed, and most of the tires' traction is used in the car-lateral direction. The car will be traveling in a nearly constant-radius path. In a street-intersection turn on a street circuit, this phase may be so brief as to be negligible. In a carrousel-type turn or long sweeper, or on a high-speed oval, the car may experience approximately steady-state cornering for as long as five seconds.

In this situation, the car has pitch and roll velocities very close to zero. Pitch displacement is also close to zero. Roll displacement is substantial, and outward. Suspension displacements are compression on the outside wheels, and extension on the inside wheels. Suspension velocities are close to zero. Therefore, damper forces will likewise be close to zero.

This means that the car will be sensitive to springs and anti-roll bars, and insensitive to dampers. Stiffening either front spring, or the front anti-roll bar, will tighten the car (add understeer). Stiffening either rear spring, or the rear anti-roll bar, will loosen the car (add oversteer).

Remember we are assuming that the turn is completely unbanked. In banked turns, the inside suspensions may compress. With soft springs and stiff bars, this may happen at surprisingly shallow banking angles. In such cases, effect of outside spring and anti-roll bar changes are the same as in a flat turn, but effects of inside spring changes reverse. A stiffer inside front spring will loosen the car (add oversteer). A stiffer inside rear spring will tighten the car (add understeer). With a beam axle, we may have moderate compression of the inside spring even though we have moderate extension at the outside tire, because the spring will be inboard of the tire. In this situation, there will be a *node*, or a point where there is neither compression nor extension, somewhere between the inside tire and the inside spring.

I am digressing from our simplified flat-turn example here to remind the reader that our example is simple, but the real world is complex. As we proceed through our hypothetical flat turn, it is important for the reader to pay attention to how things work and why, rather than treating these simplified dynamics as a universally applicable troubleshooting guide.

Part four, or early exit – car-forward acceleration present and increasing, but less important than car-lateral acceleration: The driver now begins to apply greater throttle than required to merely sustain constant speed, and begins to release the car in terms of cornering. The car's lateral acceleration is diminishing, and its forward acceleration is increasing. In this phase, lateral acceleration still dominates the car's behavior in terms of suspension displacements. The car has an outward roll displacement, but this is decreasing, so the car has an inward roll velocity. The car has a rearward pitch displacement, and a rearward pitch velocity.

The effects of roll and pitch velocities are additive at the outside front and inside rear corners, so these will be the most influential wheels in terms of damper tuning. The outside front suspension will have a compressed displacement, but this will be diminishing, so the velocity will be in the extension direction. The inside rear suspension will have an extended displacement, but this will again be diminishing, so the velocity will be in the compression direction.

The effects of roll and pitch are subtractive at the inside front and outside rear. We cannot generalize about the net velocities at these corners, except to say that they will be smaller than at the outside front and inside rear. Therefore, the car will be relatively insensitive to damping changes at these corners.

Note that we now have at least two corners where the displacement and the velocity are opposite in direction. This means that stiffening the spring and stiffening the damper have opposite effects on wheel load, diagonal percentage, and oversteer/understeer balance. Spring and anti-roll bar effects are as in earlier parts of the turn: stiffer front tightens (adds understeer); stiffer rear loosens (adds oversteer). As part four progresses, the outside front and inside rear suspensions approach their static positions, and the influence of spring rates at these corners correspondingly diminishes. So, especially toward the end of part four, the corners where the shocks matter most are the corners where the springs matter least.

At the outside front, stiffening the low-speed extension damping adds a negative, or compressive force. This reduces wheel loading, reduces diagonal percentage, and loosens the car (adds oversteer) – an opposite effect from stiffening the spring or the bar, as long as the spring is compressed compared to static. At the inside rear, stiffening the low-speed compression damping adds a positive, or extension force. This increases wheel loading, increases diagonal percentage, and tightens the car (adds understeer). Again, this is opposite to the effect of stiffening the spring or bar, as long as the spring is extended compared to static.

Note that I am contradicting the much-repeated advice to soften inside rear compression damping to hasten loading of the inside rear and tighten exit. In fact, softening the inside rear compression damping momentarily diminishes total extension force, and therefore momentarily diminishes inside rear tire loading and diagonal percentage, compared to stiffer inside rear compression damping.

Part five, or late exit – combined forward and lateral acceleration as in part four, but with forward acceleration dominant: The difference between part five and part four is the displacement direction at the outside front and inside rear. Forward acceleration is now large enough, and lateral acceleration is small enough, so that the outside front is extended relative to static (though less than the inside front), and the inside rear is compressed relative to static (though less than the outside rear). Suspension velocities are similar to part four: greatest at the outside front and inside rear; extension on the outside front; compression on the inside rear.

So for shock tuning purposes, exit can be treated as a single phase of the cornering process, and does not need to be broken down into two parts. However, part five is distinct from part four for spring tuning, because the car's response to spring changes at the outside front and inside rear reverses. Stiffening the outside front spring loosens the car (adds oversteer), and stiffening the inside rear tightens the car (adds understeer).

It may be worth clarifying what basis of comparison I'm using when I speak of a change tightening or loosening the car (adding oversteer or understeer). In the above remarks, we are referring to the car's behavior compared to the same part of the turn, before the change in question. It is also possible to consider how a change affects a given part of the turn relative to the previous part of the turn, or some other part, with the same change, as opposed to the same part of the turn, before the change. Both of these modes of comparison are useful. We do have to be mindful of which mode we are using, however.

For example, in a flat turn, stiffening the inside rear spring loosens the car during part four of the turn – but less than it does in part three or part two, especially toward the end of part four. So we might reasonably say that a stiffer inside rear spring tightens the car in part four, **relative to its condition in the preceding portions of the turn**. Changing the choice of baseline for a comparison can change the outcome of the comparison.

We have given much attention to the distinctions between inside and outside wheels. It will of course be obvious that when the car has to turn both right and left, any given wheel will be an inside wheel in some turns and an outside wheel in others. For most road racing applications, we can condense spring, bar, and shock tuning to a surprisingly simple set of rules:

1) To tighten the car (add understeer) overall, add spring and/or bar to the front and/or take spring and/or bar out of the rear. To loosen the car (add oversteer), do the opposite: add spring and/or bar to the rear, and/or take spring and/or bar out of the front.
2) To loosen the car (add oversteer) on entry and tighten it (add understeer) on exit, add rear damping and/or take out front damping. For opposite effect, do the reverse.

On ovals, suspension tuning in general is considerably more complex, because we can use asymmetries of many kinds, in addition to everything we use in road racing.

One other nuance mainly relating to ovals is that there may exist what might be called semi-steady-state cornering conditions (my own terminology) between parts one and two and parts three and four. In steady-state cornering, longitudinal acceleration is zero, or near zero. In semi-steady-state cornering, longitudinal acceleration is substantial, but not changing.

Such a state can occur during entry if the drivers applies the brakes and then holds braking force roughly constant for a time before releasing the brakes. Assuming the driver is using the tires' full capability, lateral acceleration will also be close to constant. This will place the car on a path whose instantaneous radius is steadily decreasing, even though the car's vector-sum acceleration is not changing. The car will have an outward roll displacement and a forward pitch displacement. These will be substantially constant, and therefore all suspension displacements will be nearly constant and suspension velocities will be close to zero. That means that damping forces will be negligible, and the car will be unresponsive to damper tuning.

Semi-steady-state cornering can also occur during exit if the driver holds forward acceleration roughly constant with the throttle, while using the tires' full capability. Again, neither speed nor instantaneous path radius is constant, but vector-sum acceleration of the car is constant. And again, the suspension velocities will be close to zero, damping forces will be negligible, and the car will not be affected by damper tuning.

Regarding whether to add or reduce damping on compression or extension, and at high velocities or low, some widely repeated advice would have us set compression damping to control sprung mass motion, and set extension damping to control unsprung mass motion. In my opinion this is incorrect. At some time it may have served as simple advice to racers faced with setting the earliest double-adjustable shocks, but now we have revalveable and four-way adjustable shocks, and reasonably good shock dynos. My advice nowadays is:

1) Use low-speed damping, in both extension and compression, to manage transient weight transfer and sprung mass motion. Do not expect this to work unless the surface is smooth enough so that sprung mass motion is the main cause of suspension movement. Use the springs and bars as your main means of managing weight transfer.
2) Use damping properties at velocities above 2 in/sec to manage sprung and unsprung mass behavior over road irregularities. Again, both compression and extension matter.
3) Keep compression and extension damping in reasonable proportion to each other. At most absolute velocities, extension damping should be at least a little stiffer than compression damping, but usually not more than twice as stiff and never more than three times as stiff unless you are deliberately trying to make the car jack down.

This relatively brief discussion will inevitably not have covered all possible situations, but hopefully it has covered the main principles, and illustrated a useful way to think systematically about springs, anti-roll bars, and shocks. Evaluate effects of a change in terms of whether it adds to or diminishes extension force at the corner you're changing, then imagine the car on the scales and imagine you are adjusting the extension force the same direction with the spring seat or jacking screw, and you can predict the change's effect on car behavior.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

RIGHT REAR SPRING STIFFEST?

I frequently watch the pre-race shows for NASCAR events and listen carefully when they are speaking of setups. Two times now in the past year they have spoken about the stiffest spring on the whole car being the right rear. One crew chief asked a former driver if he ever thought he would see the day that was the case. The driver said he had won many a race with the right rear being the softest spring, and that has been the case with my experience as well. They spoke further to say that the trend now was for the younger former open wheel drivers to run off the right rear tire. Physics to my knowledge haven't changed. Are they running extremely heavy front bars and super light springs, or what gives? I thought the jounce bumper trend had been made illegal. Are their ways of thinking something we local drivers could apply to our cars for an outside-the-box thinking advantage?

At the risk of disillusioning my legions of admirers, first let me confess that I do not have throngs of Winston Cup crew chiefs and engineers climbing over each other to tell me their setups. I've worked with just a few, and a few get this newsletter. However, maybe I can be of some help, and if I say anything wrong, maybe I'll get straightened out.

Traditionally, the most common approach to spring splits has been to run the front end right-stiff and the rear end left-stiff. This usually makes the right front the stiffest spring and the right rear the softest. I never really understood that approach, and I have generally advised either running right-stiff or left-stiff at both ends, except perhaps in cases where the rear end lifts under power – a condition normally only seen in dirt cars.

Two main considerations determine whether we run the car right-stiff or left-stiff: cornering balance when forward and rearward accelerations accompany cornering, and adapting the car to the banking angle of the track. Taking the former of these first, when a car has different amounts of pitch resistance on the right and left sides, and the tires are accelerating the car forward or rearward, the diagonal percentage changes. If the car is right-stiff, that tends to make the diagonal percentage increase in rearward acceleration (braking) and decrease in forward acceleration. That tightens entry and loosens exit. If the car is left-stiff, that does the opposite: frees up entry, tightens exit.

Regardless of banking angle, stiffer right springs add roll resistance. However, if the track is banked enough so that the left springs compress (I don't mean more than the right ones, just compress rather than extend), we actually increase the roll resistance with softer springs, on the left side only. This means that with soft left springs, we can make the left wheels more compliant over bumps, and improve camber control by reducing roll, at the same time. This is in contrast to our usual dilemma of having to stiffen up the suspension to control roll and camber change, at the expense of roadholding.

On a short track, with tight turns, especially on dirt, it is most common to have trouble getting the car to turn readily enough on entry, and to have trouble hooking up the rear tires on exit. If this is a problem, that argues against right-stiff springing. But on high-speed tracks, it is common to want to tighten entry. The driver is going in as deep as possible, and braking hard while cornering, from speeds as high as 190 or even 200 mph. It may take as long as ten seconds to complete the turn, and the car is following a very large-radius path. We don't want to achieve large yaw accelerations then. The need is for a car that doesn't want to come around when braking and cornering at the same time.

On top of that, the front ends on stock cars generally lose anti-dive rapidly as the suspension compresses, because the side view projected control arm is shorter for the lower control arm than for the upper. This means that when the car is rolled to the right, anti-dive asymmetry de-wedges the car when braking, unless the right front has a lot more anti-dive than the left front at static position. Right-stiff springing helps compensate for this.

It is possible to tighten entry by using lots of front brake. However, if the turns require significant braking, using lots of front brake overworks the front brakes. As the front brakes go away, the brake bias shifts toward the rear again. Also, we really would like to set the brake bias for shortest stops, which means we want the rears to do around 30% of the work. If we run more than 70% front brake, we hurt the car's ability to stop quickly for pit stops, and to brake well when avoiding accidents on the track. Therefore, it's best to get the desired entry balance with the suspension.

What about exit balance? On a short track, we are often fighting to get the car tight enough, to control wheelspin. However, even on a short track it is possible to get exit too tight, and have a power push. At high speed and with steep banking, it is harder to get wheelspin. Cup cars do have stout motors, but they also have ample tire loading from the banking and the aero, and a lot less torque multiplication from the gears than they would on a short track. It is therefore not uncommon for the driver to report a "push-loose" condition. That means the car is basically tight power-on, so the driver feeds in more power trying to get it loose, but this doesn't happen until power is sufficient to really cause wheelspin, whereupon the car goes wheelspin-loose. The driver then has a hard time finding a stable point where the car has good balance. In this situation, if the car is freer power-on, the driver can find good balance at moderate throttle.

We see an analogous situation in road racing or short-track racing in classes like Formula Ford or pavement mini-stock, where the car has good grip and modest power. Such cars often want a freer setup than more powerful cars would, to get proper exit balance.

Are soft front springs and stiff anti-roll bars still favored since bump rubbers have been outlawed? Yes. The springs can't be as soft as when bump rubbers were legal, but it is common to run the front end as soft in ride as is possible without the rubbers. This soft ride rate makes the front end run lower through the turns. That adds aero downforce. This will work on a short track too, though the effect will be less pronounced. One might think we could do the same thing by running a lower static ride height, but stock car racing rules usually include a minimum static ground clearance.

In the August 2001 newsletter, I addressed the subject of things that make spring changes work backwards. I introduced the term *critical angle* to describe the track banking angle at which the left spring neither compresses nor extends. This angle is usually not identical for both ends of the car. At angles steeper than critical, effects of left spring changes reverse, for the end of the car in question.

Running soft springs and a stiff bar reduces critical angle. The left spring will compress rather than extend at surprisingly small banking angles. (Putting it another way, with the soft ride rate, the entire front end will drop more on the banking.) That means a softer left front spring tightens the car. This is neither a disadvantage nor an advantage, just something to be aware of when running such setups.

Regarding the suggestion that the young drivers come from open-wheel racing and therefore like to run a stock car with the right rear tire heavily loaded, I question that.

First of all, it may be true that some drivers are getting seat time in sprints and midgets – Jeff Gordon and Tony Stewart for example – but the most common road to Cup is through lower pavement stock car divisions: NASCAR Weekly Racing Series, USAR Hooters Pro Cup, ARCA, ASA, Busch. It is normal for young drivers to spend time in these series before transitioning to Cup. And not all drivers who do well in sprints and midgets are able to run well with a stock car. Some do, some don't.

As you correctly note, no laws of physics have been repealed. A stock car doesn't have 60% rear and a huge right rear tire, or a big rear wing. If you make the right rear carry a lot of load when cornering in a stock car, you get a loose car. Driver preferences on car balance vary, but only within a narrow range. Nobody likes a car that's way loose at 150+ mph. Also, the car has to be fairly neutral to avoid having a tire wear problem.

But you can't necessarily conclude that the right rear is more heavily loaded just because the spring is stiffer. Other things being equal it would be, but other things don't have to be equal. If you combine a stiffer right rear spring with more static diagonal, a bigger front anti-roll bar, or a lower Panhard bar, you can compensate for the effect of the spring and load the tire about like you were loading it before.

Of these different possibilities, the one that looks most appealing to me is running more static diagonal. If we do that, then when grip is poor and lateral acceleration is less, the springs have less effect on wheel loading and the static loadings have more influence. That means the car loads the right rear less, and runs tighter. When grip is good and lateral acceleration is greater, the spring affects wheel loadings more, so right rear loading is greater and the car runs looser. This means that with more right rear spring and more static diagonal, the car will have less tendency to go loose on slick, as stock cars have traditionally done. With enough rear roll resistance and static diagonal, a stock car can even be made to go tight on slick. Between these extremes, we can find a setup whose balance varies relatively little with changes in the condition of the track or the tires.

I have clients successfully applying this reasoning on dirt and pavement short tracks, although generally the added rear roll resistance is achieved without a right-stiff spring split, except in cars that lift the rear under power.

INFLUENCE OF PUSHROD ANGLE ON WHEEL RATE

I would like to know how to calculate vertical stiffness and roll stiffness taking into account the angle of the pushrods. I have seen lots of formulas but none of them take into account the pushrods. Also, how do I calculate the total roll of the car starting with the difference in movement between the wheels?

Taking the second question first, take the difference between the right and left wheels, and divide by the track width. The quotient is the tangent of the roll angle, so you find the angle that has that tangent. As an equation:

$$\theta_r = \tan^{-1}[(h_l - h_r)/T]$$

where:

θ_r = roll angle
h_l = left ride height (average of front and rear)
h_r = right ride height (average of front and rear)
T = track width (average of front and rear)

Now, as to the angle of the push rods, you need that if you want to calculate stresses in the pushrods. But to calculate wheel rate in ride or roll, you just need to know the motion ratio from the spring to the wheel. Everything in between – the pushrod, the rocker, the control arm – only matters to the extent that it influences the motion ratio. The pushrod angle does affect the motion ratio, but it is just one factor. For an existing car, the easiest method is to simply measure how much the spring shortens or lengthens for an inch or centimeter of wheel motion. For a car that’s in the design process, if you’re designing on a computer, you can move the wheel and measure the spring length on the computer. If you’re drawing manually, you basically do the same thing by hand and estimate the motion ratio.

Once you have the motion ratio, square it and multiply by the spring rate, and you have the wheel rate in ride. Then do the same for the anti-roll bar, which may have a different motion ratio. Add the rate from the anti-roll bar to the wheel rate in ride, and that’s the wheel rate in roll. Be sure that when figuring the rate of the anti-roll bar, you calculate the pounds or Newtons per inch or millimeter **per wheel** – that is, the force change per unit of opposite motion on both wheels, not just one.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

LOAD TRANSFER BASICS

I am a bit confused regarding the subject of load transfer. I have always thought that the load transfer in a curve should be maximized (if tire wear is not an issue), and that one will achieve that by stiffening the suspension.

Now I have met people saying that load transfer should be minimized and the reason is that you want as near the same load at all tires as possible – claiming decreasing tire friction with increasing load. Okay, a softer suspension gives more grip, I know, but how come you run faster with a stiffer setup (to a certain limit).

I have also met people saying the opposite is the truth. By softening the suspension, the CG moves more towards the outer side in roll, and therefore more load will be transferred to the outer side. Sounds reasonable, but isn't it so that the shocks have a role here, absorbing energy?

Another point of view from a well renowned man is that you will always have a certain amount of load transfer at a constant speed in a constant-radius turn, no matter the chassis setup – the thing is to find the balance front to rear. Okay, I fully accept that, but what is the importance of the general stiffness of the setup then? Why are we bothering at all about roll centers, cambers, springs, and anti-roll bars?

Can you please enlighten me (and probably many more) what is happening when load transfers, and what it is we want to achieve? Can you please also discuss how this applies to different types of cars and the differences in thinking regarding suspension design and setup?

This is not exactly the first time the subject of load transfer has been addressed, even by me. In fact, the April/May/June issue dealt with the influence of springs and dampers (shocks) on load transfer at considerable length. So by addressing the basics of the subject now, I'm doing things backwards. But the question above came in recently, reminding me that the whole world hasn't been reading my newsletter like a textbook, and the newsletter cannot have the logical progression of information from issue to issue that a textbook has from chapter to chapter, and still respond to questions.

So we will tackle the basics of load transfer this month. Hopefully, those readers who found the April/May/June discussion a bit over their heads can read this, and then read April/May/June and have it make more sense.

For simplicity, I will ignore aerodynamic factors, even though these are often highly significant. Let’s also consider pure steady-state cornering only, for now, and assume that the turn is unbanked.

I share the preference for calling what we’re discussing load transfer rather than weight transfer. Actually, if we are considering an unbanked, level turn, and we’re ignoring aero, then the total tire loading really is due to gravity acting on the car’s mass, and therefore is weight. Any transfer of this load can then legitimately be called weight transfer. However, since the terms *weight* and *mass* tend to be confused, and used interchangeably in informal discourse, “weight transfer” suggests that the change in tire loads is primarily due to movement of the center of gravity, or center of mass, relative to the tire contact patches – and that is not mainly what’s going on, although it is a small factor.

What is mainly happening when load transfers is this: The tires are generating a horizontal force (toward the center of the turn, called centripetal force) at ground level. This produces an acceleration of the car (toward the center of the turn, called centripetal acceleration). The car resists this with an inertia force (centrifugal force) which acts horizontally at the center of mass, in the opposite direction to the tire force.

These two forces are equal and opposite, but their lines of action are offset. The sum of the tire forces acts at ground level; the inertia force acts at CG height. We therefore have a *couple*: a torque or rotational force due to the offset lines of action – equal to the force times the offset: the centrifugal force times the CG height. This torque acts about an axis lengthwise to the car, so it tries to move the car in the angular mode we call roll. It is therefore called a *roll couple*. It acts opposite to the direction of the turn: in a left turn, it tries to roll the car to the right.

Short of the point of rollover or wheel lift, this couple is resisted by an *antiroll couple* which takes the form of a load increase on the two outside wheels and an equal load decrease on the two inside wheels. Essentially, the ground pushes up harder on the outside wheels, and less hard on the inside wheels, in response to the way the tires are pushing down on it, and exerts a counter-torque on the car that keeps it right side up. The load change on the inside or outside wheel pair, times the track width, has to equal the centrifugal force times the CG height.

So the person who told the questioner that suspension design and setup have no effect on total amount of load transfer is basically correct. The important factors are CG height, track width, and the amount of centrifugal force. It is true that when the vehicle rolls, the CG moves a little bit toward the outside wheels. Therefore, softening the suspension increases load transfer slightly, and stiffening the suspension reduces load transfer slightly. However, for cars these effects are small. For tall vehicles such as trucks, CG movement is somewhat more important. For trucks with cargoes that can shift or slosh, CG movement can become highly significant.

So the questioner's idea that stiffening the suspension increases load transfer is incorrect – **as it relates to TOTAL load transfer**. But it's correct as it relates to the front or rear wheel pair's **share** of the total.

The anti-roll couple we mentioned is actually the sum of two anti-roll couples – one from the front wheels, one from the rears. These two couples are not necessarily equal. The front and rear tire pairs split the job of resisting roll according to their **relative** roll resistance: the stiff end sees a greater share. In considering the roll resistance, we need to include both the springs (anti-roll bars are interconnective springs) and the suspension geometry.

In steady-state cornering, on a smooth surface, the suspension is not in motion, so the shocks, at least theoretically, have no effect on load transfer. For a more detailed discussion of the effects of springs and dampers on load transfer, see the April/May/June 2003 newsletter.

Now, is lateral load transfer beneficial or harmful to cornering (centripetal) force capability? The questioner correctly notes that the answer to this is directly dependent on how friction varies with load.

Does friction decrease with load? No, it increases. But the **coefficient** of friction decreases with load. The coefficient of friction is the ratio of friction force to normal (road vertical) force – how many pounds of cornering force we can get from the tire, per pound of loading.

According to the classical model of *Coulomb friction*, which models sliding friction between hard, dry, clean, smooth materials, the coefficient of friction is a constant for any given pair of materials. The friction force is directly proportional to load. For many situations, the Coulomb model is fairly accurate. However, tires and roads do not conform to this model. This should not be too shocking, since neither the tire nor the road is smooth and hard. The road is generally hard compared to the tire, but it isn't generally smooth. It is gritty, like sandpaper. The tire may be smooth, but it isn't hard. It is soft enough to conform to the road's gritty surface, and interlock with it. It is tough, meaning it resists being torn apart. If it is a racing slick, it may actually be tacky, meaning it sticks to a smooth surface even in the absence of sustained normal force.

Moreover, a tire is not a rigid structure; it is a flexible – but not totally limp – bladder, carrying an inflation gas under pressure. This means that the macroscopic, or visible-scale, size of the contact patch increases with load.

These factors make it very difficult to mathematically model a tire's behavior. That hasn't kept people from trying, but any equation that seeks to express the relationship between tire loading and limit cornering force will not only be complex, but must necessarily include values for the particular tire that can only be determined by test, because whatever the relationship may be, it isn't the same for all tires. Therefore, it is impossible to accurately know the relationship between loading and cornering force without testing. And any equations seeking to describe tire behavior must be developed by fitting curves to experimental data, rather than predicting behavior from abstract theory. And if this weren't enough, tire properties vary with temperature, tire age, camber, inflation pressure, and of course the road surface – which also has properties that vary with the weather, the age of the surface, how much rubber and oil have been deposited, and so on. To avoid the effects of real-world road surfaces and weather, tires are now tested indoors, under fairly careful climate control, against belts or drums rather than roads. This is a valid necessity if we want to meaningfully test one tire against another, but it is important to remember that tire behavior in the real world is much more variable than tire behavior in a controlled environment.

Fortunately, it is much easier to look at trends in tire behavior than to predict precise numbers. And one trend that all tires exhibit is that the coefficient of friction decreases with load. It may decrease rapidly, or almost imperceptibly. It may decrease very little at light loads, and more rapidly at greater loads. But it always decreases with load. This property is termed *load sensitivity of the coefficient of friction,* or just *load sensitivity* for short.

This means that a pair of identical tires has the greatest cornering capability when loaded equally, and as we load the outside one more and the inside one less, we lose cornering capability. It's not that the outside tire loses friction force. The outside tire gains grip – but the inside one loses grip at a greater rate, so the total decreases.

Some writers contend that there are times or conditions where load sensitivity reverses. Some say it reverses on dirt, some say just on dry-slick dirt, some say on snow or ice, some say when the tire is cold, some say during corner entry. I don't believe any of these theories myself. If load sensitivity worked backwards, the car's responses to roll stiffness adjustments and static corner weight adjustments would also reverse, and I have not encountered this – except in situations where there was clearly some explanation other than reversal of load sensitivity.

Now, if load transfer in cornering does not depend significantly on overall suspension stiffness, why does overall stiffness of our setup matter at all? It matters for other reasons than load transfer: control of camber and aerodynamics, ability to absorb road irregularities, how high we have to set the static ride height.

A stiff setup reduces camber changes, particularly with independent suspension. It reduces changes in the car's roll and pitch angles, relative to the airflow and relative to the road. It allows the car to run with a lower static ride height, without bottoming excessively. It makes the car more responsive to driver inputs.

A soft setup allows the wheels to follow bumps better. Loads on the tires vary less, and they spend less time airborne. The tires often heat up more slowly, and wear less. The car and the driver don't get beaten up so badly by the bumps, improving endurance of both car and driver.

So overall stiffness has to be a compromise that balances these conflicting factors in a manner appropriate to the conditions, while relative stiffness of the front and rear suspensions manages load transfer distribution.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September/October 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SHOCK RESEARCH UPDATE

Regular readers will recall that I am coordinating an effort to investigate sensitivity of dampers to factors other than velocity, including acceleration and jerk. I still have a long way to go on this, and am still seeking persons to contribute shocks for test, especially shocks that dyno similarly in standard crank dyno testing but behave differently on the track.

I have been able to learn a little already, however, and I would like to relate these findings.

When I look at dyno output for a shock, I like to look at a force-versus-absolute-velocity trace for the full stroke: both the rebound-closing/compression-opening (extended) end and the compression-closing/rebound-opening (compressed) end. Such a trace will have two "noses" or points at the left side of the graph, each having its vertex at the zero-velocity point. Another way of displaying the information is a force-versus-velocity plot. This will generally take the form of an S-shaped loop, which will cross the zero-velocity line at two distinct points. The force-versus-absolute-velocity trace is just the force-versus-velocity trace, with the negative-velocity part of the trace folded over to the positive side because the velocities are being expressed as their absolute rather than signed values.

Almost invariably, the zero-velocity point for the cc/ro trace will be higher (indicating greater extension force) than the rc/co trace.

When the shock is cycled very slowly, if it has any bleed or leakage at all past the piston, and especially if we hold the shock stationary at the ends of the stroke and let the force readings stabilize, we will get a reading of the gas spring effect. We should expect the two nose points to be separated by this amount even in the absence of acceleration sensitivity or other effects.

The appearance of a loop-shaped force-versus-velocity trace has led some writers to call the effect *hysteresis*. This may or may not be strictly correct, depending on exactly how we define the term, and how meticulous we want to be. My Webster's dictionary defines hysteresis as "a retardation of the effect when the forces acting upon a body are changed (as if from viscosity or internal friction)".

My dictionary traces the etymology of the word to the Greek verb *hysterein*, meaning to be late or fall short.

How does this relate to what dampers do, or to what we'd like them to do? I think it's pretty safe to say that we want a damper to generate a force opposing **motion** of the suspension at all times. How big this force should be, in what circumstances, is less certain. But we can at least say that if the damper is generating a force in the same direction as the system is moving, the damper is exaggerating motion rather than damping it, for as long as this state of affairs prevails. A suspension system in such a condition is sometimes said to be *self-exciting*, which is the opposite of damped.

I am not saying that damper opposition to motion is always optimal for roadholding – merely that this is what damping means. And the shock, being a passive mechanical device, cannot be expected to know when damping is desirable. All we can reasonably expect a damper to do is damp motion, and do so in some consistent manner with respect to velocity, acceleration, and jerk of its sprung and unsprung elements.

Note that motion, force, and force change are different from each other. Actually, viscosity does not always retard response to a change in force as the dictionary's language would suggest. Sometimes it exaggerates response to a change in force – so even my dictionary doesn't make perfect sense. For example, suppose we have a system damped by an ideal viscous fluid (totally incompressible and non-volatile). Suppose the system is in motion, and the force moving it is diminishing or reversing, and the motion is slowing. The damping force will oppose the motion, and therefore hasten the slowing – meaning the viscosity is actually hastening response to the force change, not retarding response.

So when we discuss the meaning of "hysteresis", we confront a situation where even those charged with defining the term have an imperfect grasp of the phenomena the term attempts to describe. Does hysteresis mean damping? If we're talking about rubber, it does. Does it mean a response lag in damping? Maybe, as some people apply it to shocks. But clearly then it does not mean damping. It may even imply temporary absence of damping, or anti-damping – behavior more like a spring than a damper. But would we speak of an ideal spring (no self-damping or non-linearity) as having hysteresis? Not ordinarily.

Engineers tend to think of hysteresis as anything that produces a loop-shaped data trace when a system is subjected to forced oscillation. But of course the meaning of this will depend on what variables we're plotting. The earliest shock dynos generated a force-versus-displacement plot. For a damper that's working reasonably well, this will always be a loop, typically shaped something like an egg, and sometimes called an egg plot. More modern dynos, including crank dynos with continuous computerized data acquisition, can also produce an egg plot.

Looking at such a plot, we may say we're seeing a hysteresis loop. A shock is supposed to have hysteresis in this sense. It won't do its job if it doesn't.

An ideal spring should produce a force-versus-displacement plot that's a straight line – or actually two straight lines, overlaid.

A spring is not a velocity-sensitive device, but in a sinusoidal test, we have a fixed relationship between velocity and displacement. Therefore, we know the velocity at any given displacement, and we can create a force-versus-absolute-velocity trace or a force-versus-velocity trace based on that. For an ideal spring, the force-versus-absolute-velocity trace will be half of a sine curve, stood on its side, or an arcsine curve for half a period – or, more precisely, two such traces, overlaid. That trace may be regarded as having two "noses" at the zero-velocity line or vertical axis, rather like what we see with many shocks. The force-versus-velocity trace will be the same, only with the negative-velocity trace folded over to the left side of the vertical axis instead of overlaid on the positive-velocity trace. That gives a loop-shaped trace, resembling an ellipse, although mathematically not a true ellipse.

So an ideal damper, with nothing in it that acts like a spring, produces a loop-shaped force-versus-displacement trace, and produces a force-versus-velocity trace that is not a loop but rather two curves overlaid. Its force-versus-absolute-velocity plot is likewise two traces overlaid. The force-versus-velocity and force-versus-absolute-velocity plots have only one zero-velocity point.

Conversely, an ideal spring produces no loop in its force-versus-displacement trace, but produces a big loop in its force-versus-velocity trace, and two widely-spaced noses in its force-versus-absolute-velocity trace, which is two traces overlaid.

A device that acts as both a spring and a damper will produce loop-shaped traces for both force-versus-displacement and force-versus-velocity. Other effects may also produce a loop-shaped force-versus-velocity trace.

In an automotive suspension damper, we can get spring effects from both the gas spring and other compliances, primarily compressibility of the fluid.

What we see when testing some – but not all – shocks is that the noses are separated by a greater amount than gas spring effect alone can account for. We also sometimes see that the nose separation tends to increase as the valving gets stiffer. This is particularly easy to see when dynoing certain adjustable shocks at a series of settings: as you stiffen the damping, the noses spread further apart.

Also, in some cases, the noses spread further apart when the shock is cycled faster, with stroke unchanged. In such a case, we are looking at an unchanged shock, but greater velocities, accelerations, and jerks.

What my collaborators have done to date is to test single-tube, deflective-disc shocks at twice the usual frequency and half the usual stroke (1" stroke, 3.2 Hz rather than 2" and 1.6 Hz), and also test them upside down. So far, these tests have not been done on shocks with other types of valving, or on low-pressure, twin-tube gas shocks.

The results have more or less confirmed what the appearance of a deflective-disc valve suggests: this type of valving is not highly acceleration-sensitive. The discs do have some inertia, of course, so this was not entirely a foregone conclusion.

When a typical stock car shock, with reasonably soft valving and some bleed, is tested upside down, the forces it generates do not change noticeably. When the shock is body-up, as it's usually installed, the piston and discs are subjected to accelerations. When the shock is body-down, the body moves instead, and the piston and discs have a constant velocity of zero – therefore no acceleration or jerk. When the shock dynos the same both ways, that implies that, at least within the range of accelerations present in the test, the shock is not acceleration-sensitive.

Also, when typical stock car shocks are tested at half the stroke and double the frequency, that generally does not have much effect on the forces. Compared to the standard test, this test produces identical velocity at any given crank angle or point in the cycle, but twice the acceleration and four times the jerk. The lack of much effect on forces in this test suggests that, at least for relatively moderate acceleration and jerk values, the shock is a predominantly velocity-sensitive device.

The finding that many dampers are not truly acceleration-sensitive is not a setback. The important thing is that we have a test for this, and therefore we can separate true acceleration sensitivity from other effects that may make a damper act different at the compressed and extended end of its stroke, and act different when cycled at differing speeds or adjustment settings. I would like to explore some such effects here.

I am adding some six shock dyno graphs to this newsletter. Three of them are on pages 6, 7 and 8. The other three are separate attachments. These last three are Adobe pdf files. You will need Adobe Acrobat to open them. An Adobe Acrobat reader is available free at www.adobe.com.

The first two plots are from a mountain bike shock used on a Formula SAE car. This is a deflective-disc shock. It has relatively little bleed, judging by the lack of a soft, progressive region at low absolute velocities. This shock is adjustable, and the plots are for a soft setting on both compression and extension, and a stiff setting on both compression and extension. The test is the same in both cases: 1" stroke, 200 rpm. This shock is so small that it can't be tested at 2" stroke and 100 rpm. It doesn't have 2" of stroke. Consequently, the only way to get the absolute velocity up to the 10 in/sec customary in race car shock testing, using a crank dyno, is to turn the crank faster than customary.

The first two pdf files are from Bilstein stock car shocks, with valving codes 5030 and 7030. Plots were furnished by Bilstein's Mooresville, North Carolina facility. These are full-size, non-adjustable shocks, also with deflective-disc valving. I chose these for comparison with the mountain bike shock because they generate roughly similar forces around 4 to 5 in/sec, but they have much more piston area (hence lower working pressures) and more bleed, as indicated by the relatively soft and progressive behavior at low absolute velocities.

Now consider the zero-velocity points on these four plots. The ones on the soft mountain bike and soft stock car shock are separated by very similar amounts: about 20 lb. However, the zero-velocity points on the stiffer stock car shock are separated by about 30 lb., while those on the stiff mountain bike trace differ by at least 130 lb. The spread on the small shock with high working pressure and little bleed grew much more as the damping was stiffened.

The third pdf file is from an 8060 Bilstein. It is stiffer at 4-5 in/sec than either of the two mountain bike calibrations shown. Yet it also shows relatively little spread at the zero-velocity points. Big piston; substantial bleed; still not much spring-like behavior.

The plot on page 8 is from a correspondent at Bilstein's Australia headquarters. It shows a shock of specifications unknown to me, tested two different ways with no changes to the shock. The Bilstein 8060 appears to be similar, although perhaps a bit softer at high speeds. The 8060 is considered a stiff shock by stock car standards, so the Australian shock is definitely stiff by stock car standards. Yet the progressive character at low speeds suggests it has similar bleed and preload to the 8060.

The plot from Australia is different from what most Americans will be used to looking at, in two obvious ways. First, the sign conventions for the velocity are reversed. (This is because the dyno is in the southern hemisphere and therefore upside down – no, not really; it's because these sign conventions are arbitrary and the choice is up to the dyno manufacturer.) Second, the units are metric. Handy conversion factors:

1" = 25.4 mm
1 mm = approx. .04 in
100 mm/sec = approx. 4 in/sec
200 mm/sec = approx. 8 in/sec
300 mm/sec = approx. 12 in/sec
400 mm/sec = approx. 16 in/sec
500 mm/sec = approx. 20 in/sec

1 Newton = approx. .225 lbf
1 lbf = approx. 4.45 N
1000 N = approx. 225 lbf
2000 N = approx. 450 lbf
3000 N = approx. 675 lbf
4000 N = approx. 900 lbf
5000 N = approx. 1125 lbf

I am told that the Australian shock was tested at the same stroke, at two different rpm's. The smaller trace appears to have been done at 30% the speed of the larger one. Applying the above conversion factors, if the stroke was 50 mm, which is close to the 2" common in shock testing, the rpm's would have been about 60 and 200. So compared to the standard 2" stroke, 100 rpm test, we are looking at a considerably lower-speed test, and one about twice as fast.

In the low-speed test, the trace shows almost no spring effect. We very nearly have two identical traces overlaid – hardly a loop at all. The medium-speed 8060 test, shown in force-versus-velocity format, would be only slightly more of a loop. But double that speed, with more high-speed damping thrown in, and we see a really fat loop. This indicates that even a shock with generous piston size and significant bleed starts to act like a spring if the speeds and forces get high enough.

Okay – what conclusions can we draw? One would be that when we set out to investigate a particular phenomenon, we may stumble upon others we weren't looking for. Nothing new here; many of the greatest discoveries in science were accidentally made this way.

Another is that end-of-stroke phenomena include not only acceleration sensitivity but effects related to entrapment and compressibility of the damping fluid – and possibly some deflections of other components. These effects are not readily separable from acceleration sensitivity, although testing the shock upside down can tell us quite a bit. Of course, not all shocks can be tested upside down.

Finally, it appears that it is highly desirable to be able to create tests that vary acceleration and jerk in mid-stroke, without altering velocity. This will involve exploiting the capabilities of recently introduced linear-motor dynos.

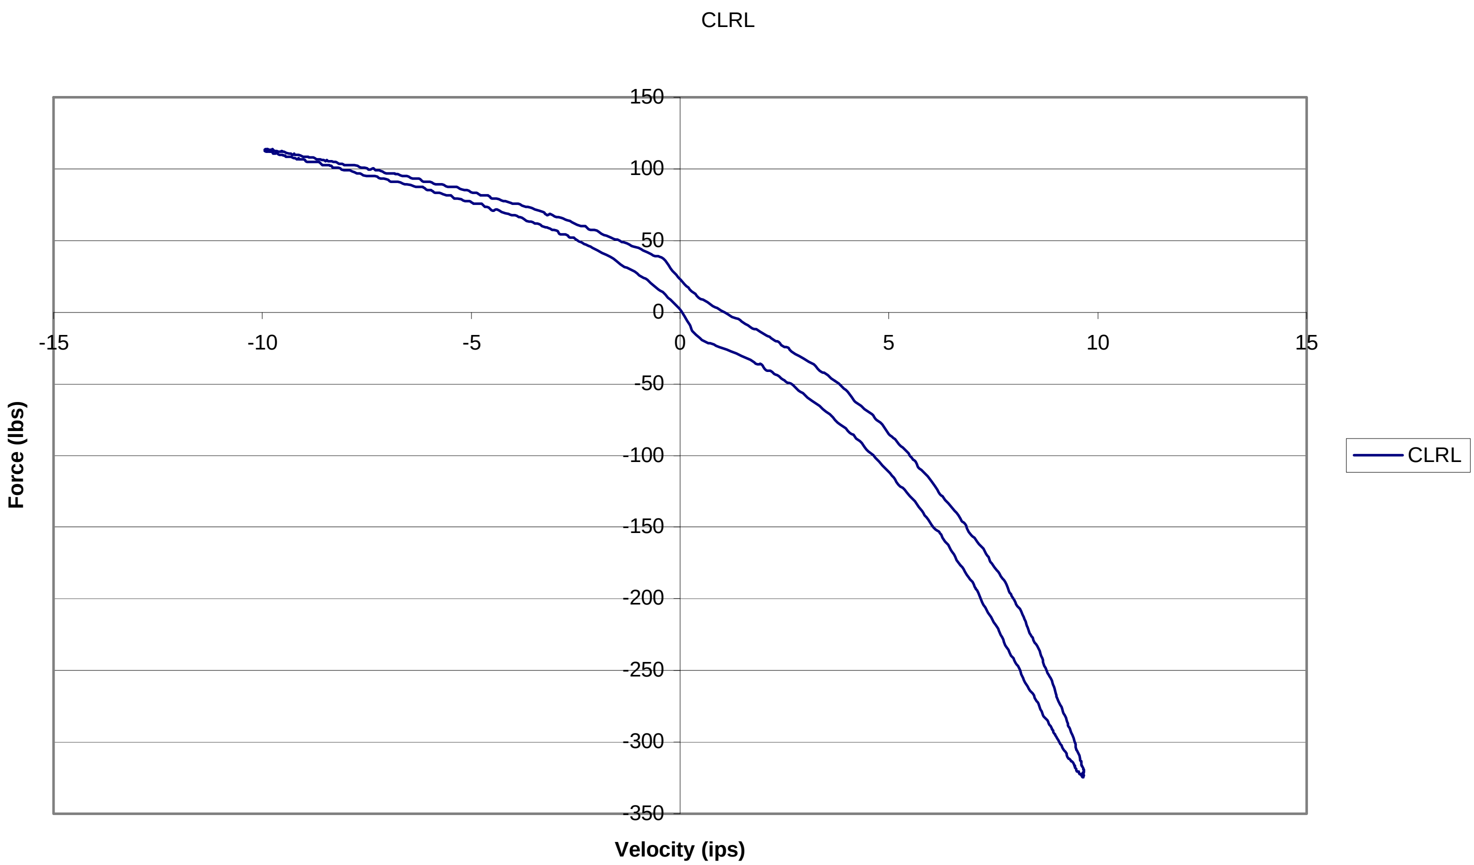
CLRL
150
100
50
0
-50
-100
-150
-200
-250
-300
-350
-15
-10
-5
0
5
10
15
Force (lbs)
Velocity (ips)
CLRL

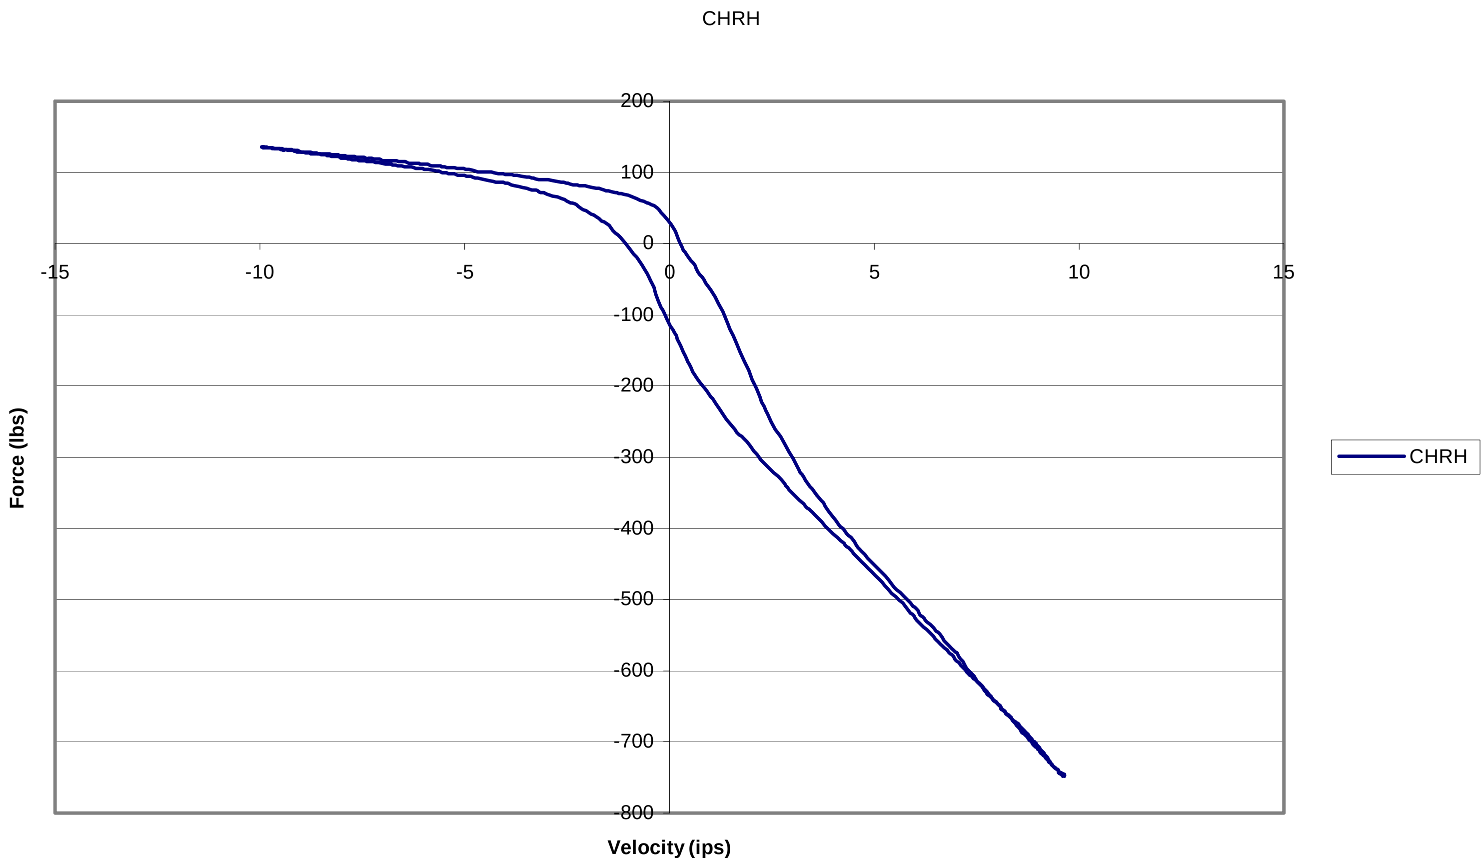
CHRH
200
100
0
-100
-200
-300
-400
-500
-600
-700
-800
-15
-10
-5
0
5
10
15
Force (lbs)
Velocity (ips)
CHRH

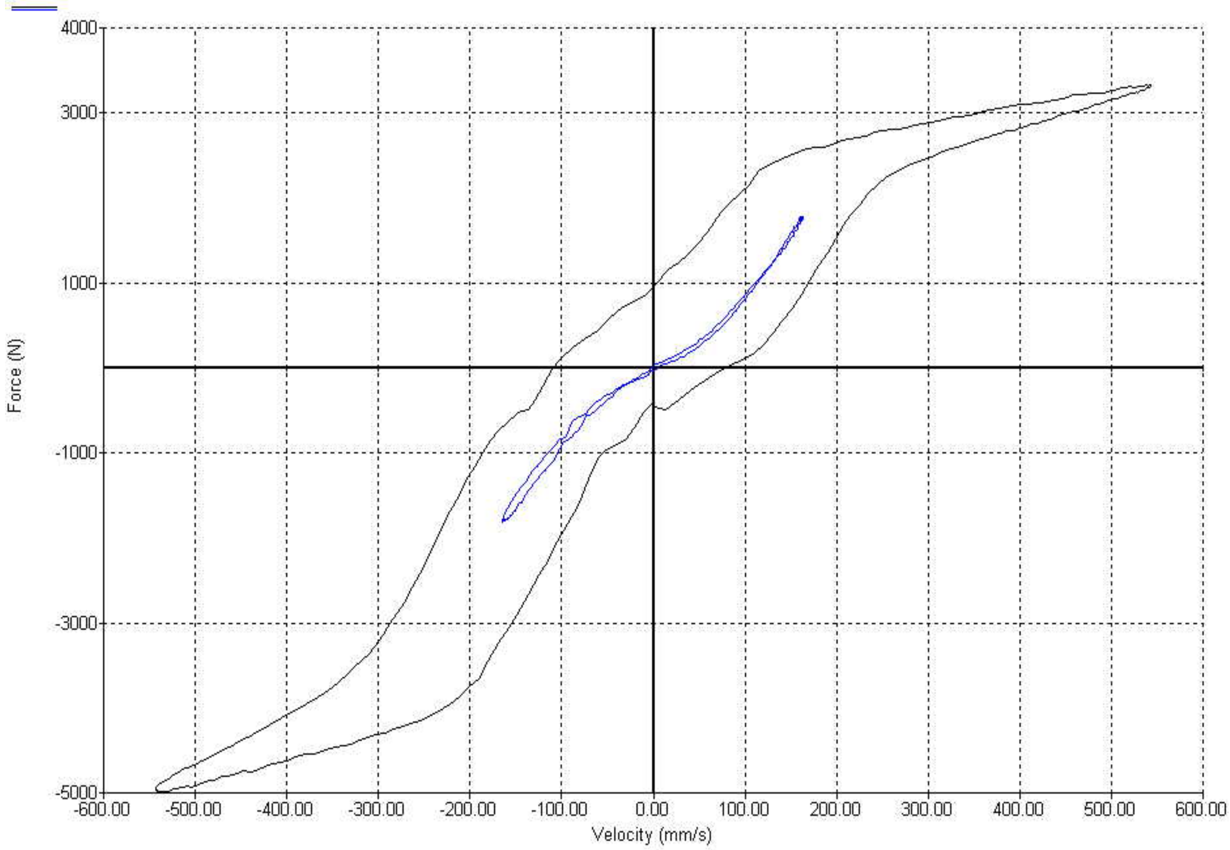
4000
3000
1000
-1000
-3000
-5000
Force (N)
-600.00
-500.00
-400.00
-300.00
-200.00
-100.00
0.00
100.00
200.00
300.00
400.00
500.00
600.00
Velocity (mm/s)

Force vs. Absolute Velocity

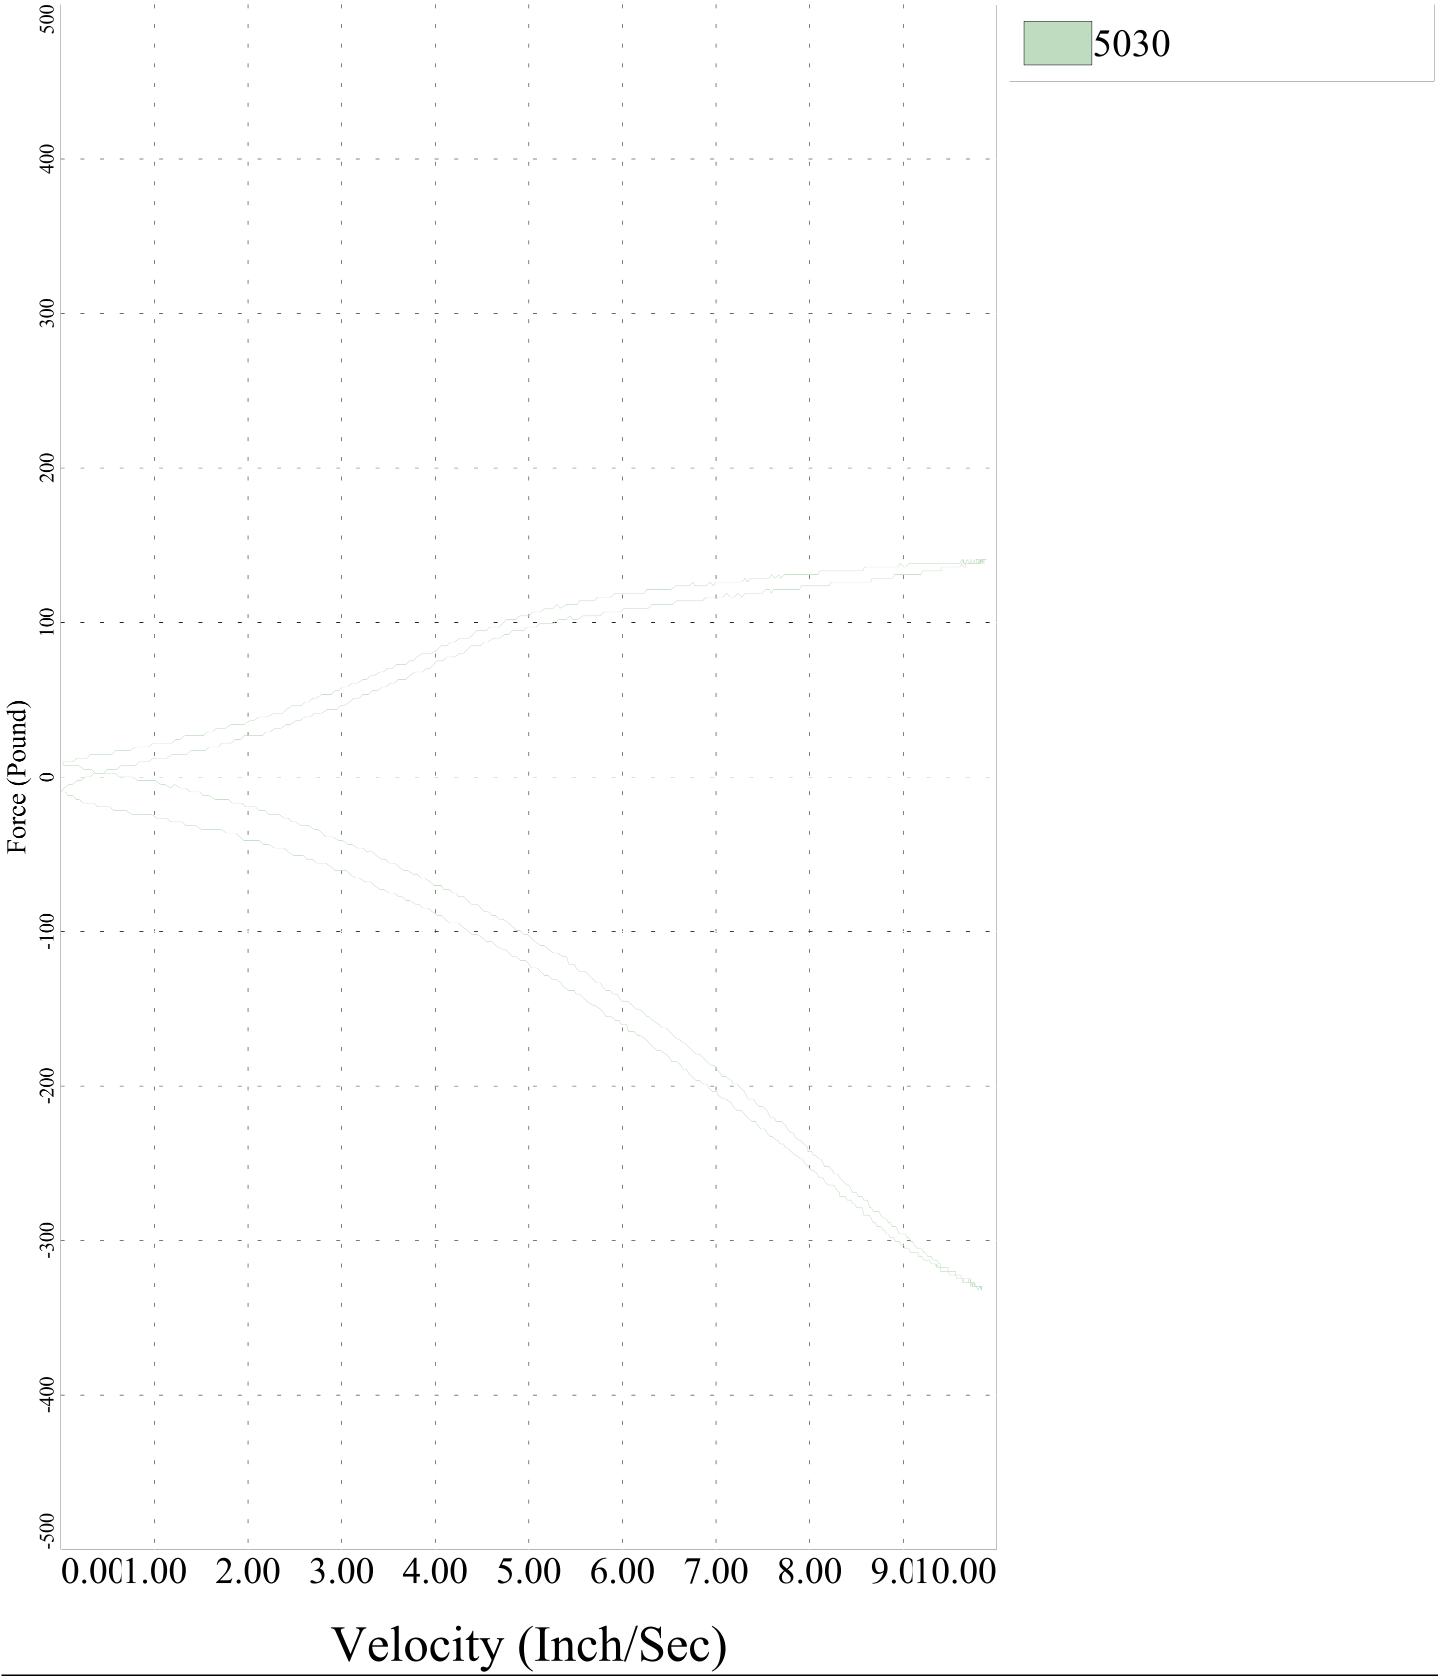

Force vs. Absolute Velocity

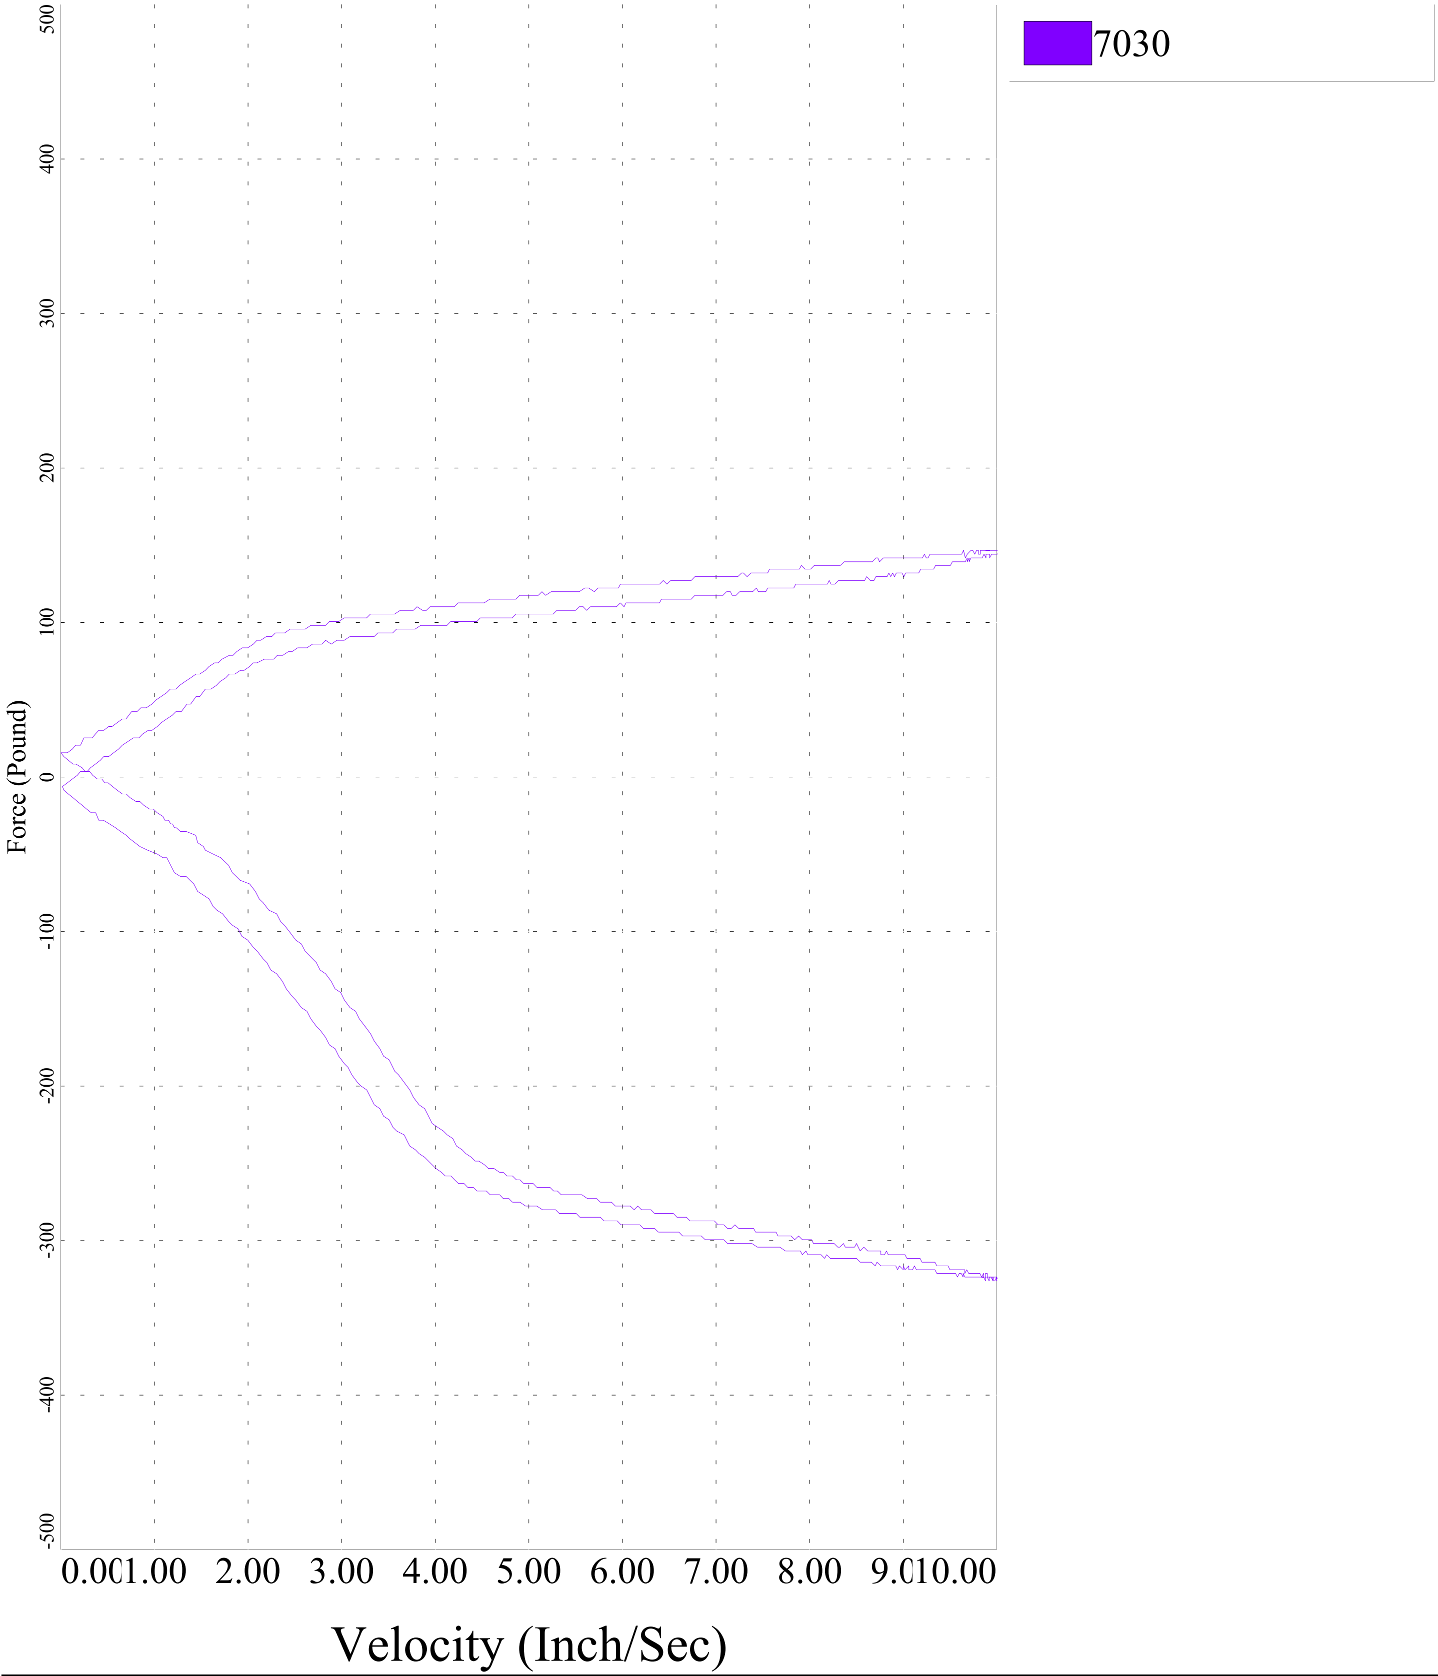

ROEHRIG ENGINEERING INC.

Force vs. Absolute Velocity

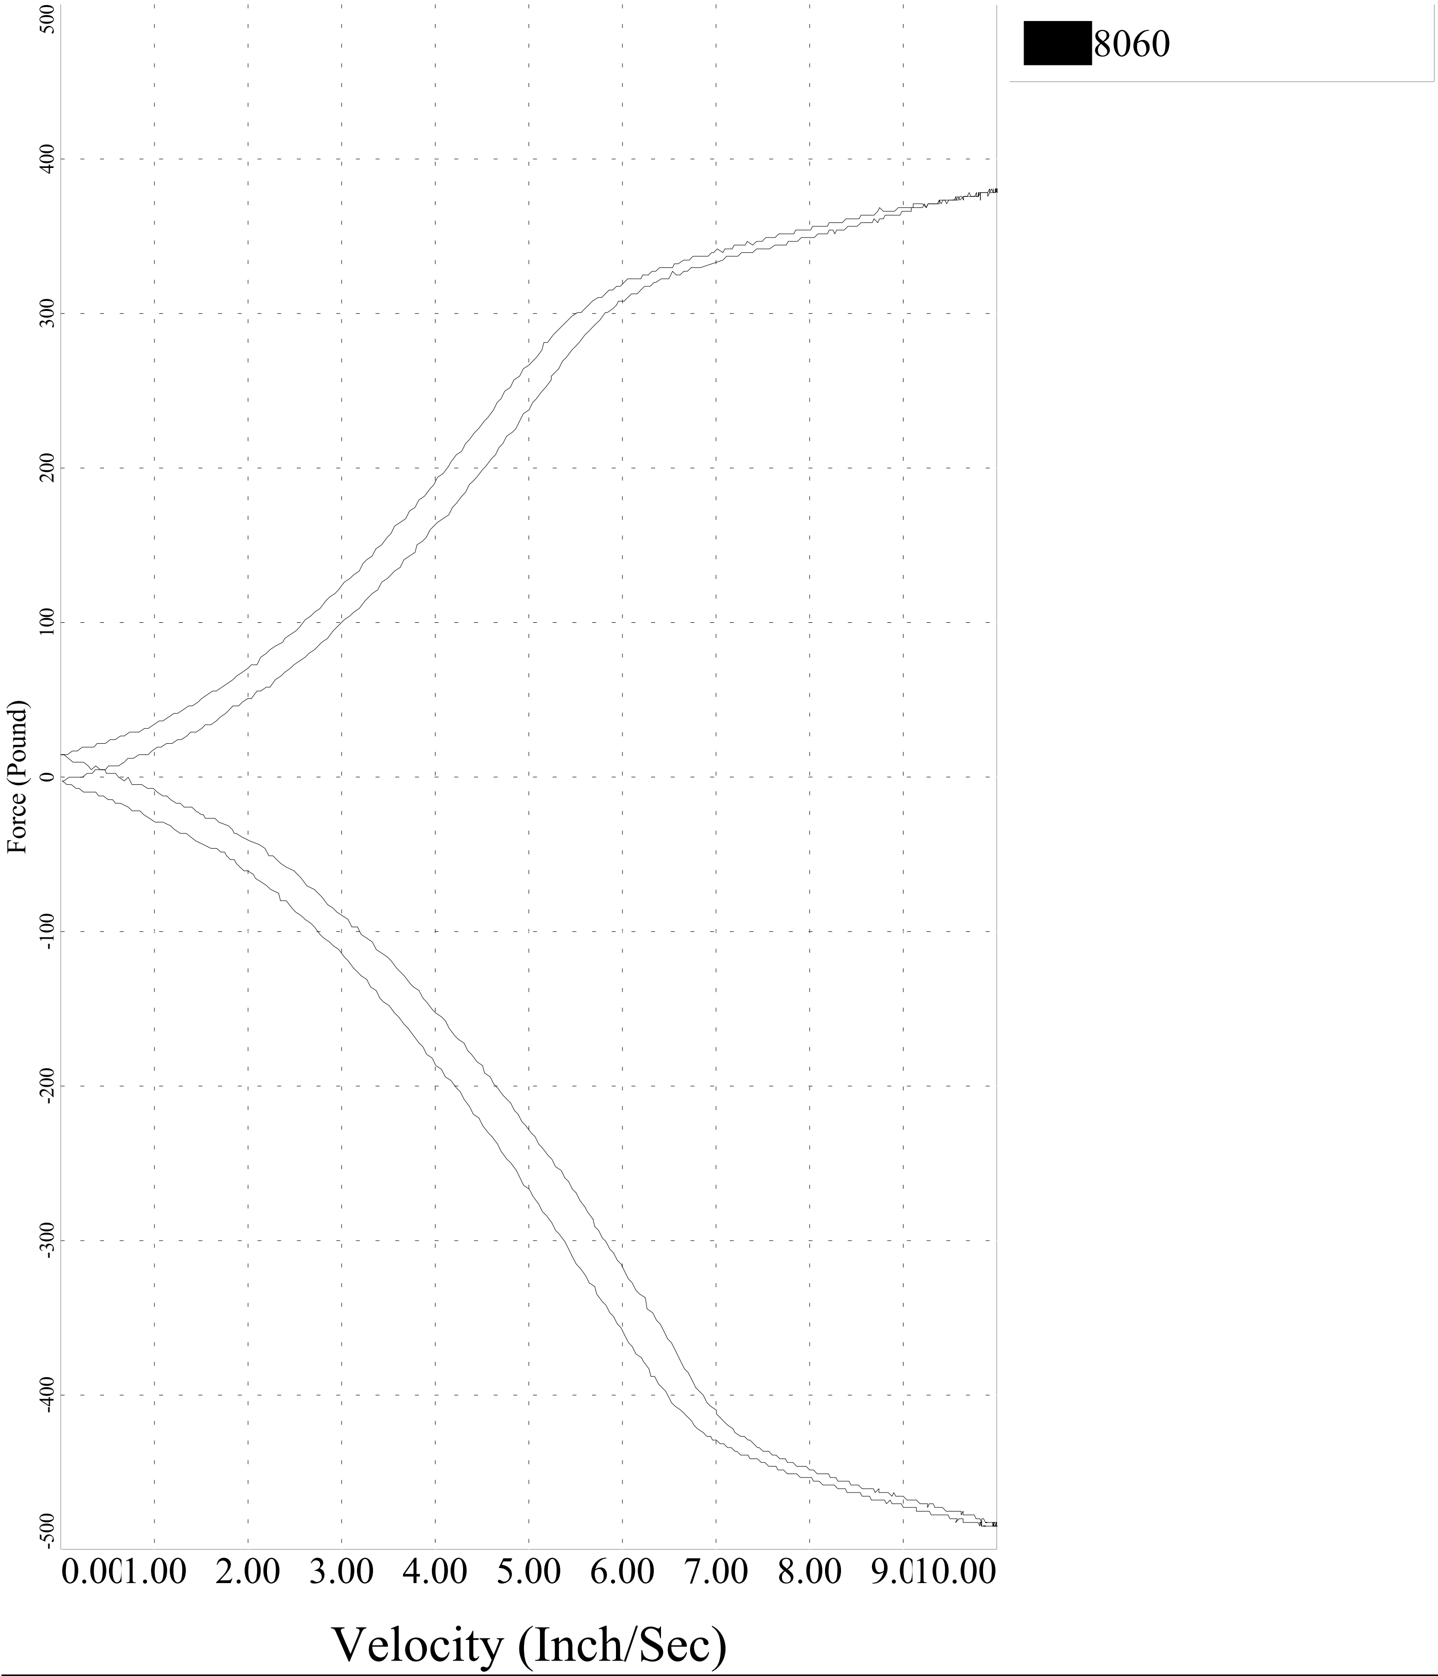

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

NEW VIDEO

I am pleased to announce that I now have videos available of the presentation I gave at UNC Charlotte this past March. The one-hour lecture is entitled "Minding Your Anti: Understanding Factors in Load Transfer". It deals with the origins of load transfer and presents a "force-based" or "lateral anti" approach to the notion of roll centers. This is original and very current thinking on the matter, and not to be found elsewhere. Videos are single VHS cassettes, and sell for US$50.00. This price includes shipping to any destination, worldwide. North Carolina residents please add 7½ % ($3.75) sales tax.

WEIGHT DISTRIBUTION AND TIRE SIZE

At what point is it worthwhile to install wider wheels and/or tires on the back of a rear-wheel-drive car – specifically if the car has close to 50/50 front/rear weight distribution? Directly related to this point, I often wonder: why is it that an F1 car can brake so much more rapidly than it can accelerate, even though the front tire contact patches are smaller?

The basic rule of thumb is that tire size should be roughly proportional to tire loading, assuming we are talking about a car that has to corner well. So if the rear percentage is close to 50%, the tires should be equal size, and if the weight distribution is 40/60, the rear tires should be half again as wide as the fronts. Ordinarily, we go by tread width for this, rather than overall width at the sidewalls. Of course, the rule is only an approximation in any case.

The rule gives us roughly optimal values for steady-state cornering. That may not be the only thing we're trying to get the car to do well, but it's certainly important in most cases.

In many cases, we do not have a free choice of tires or tire sizes. Often, our task is to optimize the car for the tires, rather than the other way around. With race cars, our tire sizes are usually limited by the rules. For street cars and some race cars, we may be constrained by the fenders. In many production cars, if we just put the biggest tire at each corner that will fit without hitting the fenders, we end up with the rears bigger than the fronts. This is partly because the rears don't have to steer, and partly because it is usual to allow for snow chains on the rears, with equal size tires.

Other practical constraints may intrude as well. I have an old Chevy Impala station wagon that serves as both my transportation and my bedroom. When I first got the car, I set it up with 8" wheels with ½" offset all around, a rear anti-roll bar, and a much stiffer front anti-roll bar than stock. The car was fun to drive, cornered quite flat, and was well-balanced. Only trouble was that it kept breaking front a/r bar links and other a/r bar hardware, and I couldn't keep front wheel bearings in it for more than 10,000 miles. I could have put a full racing front end in, but this is an old, beat-up car that I don't want to put huge amounts of money and time into. So I went to 7" wheels with zero offset in front, a much softer front a/r bar from a sedan, and 8 ½" rear wheels with the same stiff rear bar as before. The car is balanced this way too. It rolls more, but less than stock. The ultimate lateral grip is less, but acceptable. And I don't have to fix the front end all the time.

I help out a Formula SAE team. Sometimes they take my advice, and sometimes they don't. They have had a policy of doing their car in a single year, and making minimal changes from the previous year's car. One consequence of this is that many design elements get adopted simply as carryovers. The team has been using wider tires in back, and the 2003 car had the same feature, despite my urging the team to use the widest and biggest tires possible all around. The car has only 52% rear. The rear tires have 58% of the tread width. We still managed to get balanced handling, by using stiffer springs and anti-roll bar at the rear.

My point here is that in many cases tire sizes are chosen based on factors other than vehicle dynamics theory – sometimes rationally, sometimes irrationally. And because cornering balance depends on suspension as well as tires, a surprisingly wide range of tire size combinations can be made to work acceptably on any given car.

Well, okay – but limiting ourselves to considerations of vehicle dynamics, why might we want bigger tires on the rear, when the car is not markedly tail-heavy?

Depending on aerodynamic balance, higher speeds may argue for bigger rear tires, or alternatively for more nose-heavy weight distribution. We know that a tire has limited capability for combined lateral and longitudinal force. To get more longitudinal force from a tire, we sacrifice some ability to generate lateral force. We speak of the traction circle, traction ellipse, traction perimeter, or traction envelope, as a representation of the limiting values for the vector sum of lateral and longitudinal force.

When somebody mentions steady-state cornering, we may think of a typical skidpad test, with speeds somewhere in the 60 mph (100 kph) range. Throttle application to maintain this speed will probably be fairly moderate, meaning the rear tires have a large percentage of their traction envelope available for cornering. But we can also have steady-state cornering at, say, 150 mph (250 kph). Just to run that fast in a straight line requires a fair amount of power. Add the drag of four tires operating near peak slip angle, and the car may need full power, or something close, just to maintain constant speed. And powerful cars can sometimes spin the wheels in top gear, in a straight line. So in this situation, how much of our rear tire traction envelope do we have left for cornering? Not a lot, unless the traction envelope was generous to start with (big tire). Or maybe quite a lot, if the car generates sufficient rear downforce at high speed to compensate for the other effects. In a case such as a NASCAR Cup car, both the tires and the aerodynamics (except for details) are dictated by the rules, and we pretty much tune the suspension and the ballast placement around the tires and aero package. The tires are required to be equal size at both ends, and for medium to high speed tracks, the car likes around 52% front. To run more rear percentage, wider rear tires would be helpful. We could run more rear, with equal tires, but we would be making less use of the left front tire, and mid-turn speeds would be lower.

We may want to run larger rear tires in search of greater forward acceleration. Any tire has an optimum inflation pressure for making lateral force, and another, lower, optimum pressure for longitudinal force. Consequently, if we have a car that's balanced with equal size tires front and rear, and then we install larger rears but run them somewhat underinflated for cornering, we still have a balanced car, but it puts power down better.

We can take this a step further, and add roll resistance at the rear, reduce roll resistance at the front, and further increase the tire size disparity. If we take this to an extreme, we have a car that is optimized for drag racing, but also has acceptable cornering balance – although it isn't really optimized for cornering. The inside rear tire will be very lightly loaded when cornering, but with a limited-slip diff, this may be acceptable. With a live axle rear, we improve the car's launch at the drag strip if we provide a very stiff wheel rate in roll at the rear and a very soft wheel rate in roll at the front. This helps because driveshaft torque produces less change in diagonal percentage when the car is stiff at the rear and soft at the front. We may disregard cornering completely in a car we only race in a straight line, but even if we are concerned with balanced cornering behavior on the street, a drag racing suspension setup will often call for larger tires in back. Note that this reasoning regarding reduction of driveshaft torque effects does not apply with independent rear suspension, where these effects are absent regardless.

Looking at the opposite end of the spectrum, IMCA-style modifieds, as raced in the US, may have as much as 59% rear, or even 60% with a full fuel load, and they are required to run equal size tires front and rear. In this case, we have a car that often runs on very slick dirt tracks, with tires that don't give much grip, and has lots of power. It needs the rear percentage to put power down. Even in the turns, a large percentage of the cornering force is car-longitudinal drive force from the rear tires, applied at an angle to the car's direction of travel because the car is powersliding. To get decent cornering balance in such a car, we have to make it corner on three wheels, or very nearly so. We under-utilize the left front tire, but we accept this to get forward traction.

If we run the same car on pavement, we need to move ballast, and perhaps even the engine, forward in the car.

In answer to your question of why F1 cars achieve greater accelerations rearward when braking than forward under power, even though the front tires are smaller than the rears, actually almost all vehicles exhibit this property. The main reason is that we have brakes on all four wheels, but propulsion on only two. A secondary reason is that drag, aerodynamic and mechanical, acts rearward, so it assists braking but opposes propulsion. About the only way we could produce a vehicle that accelerates faster forward than rearward would be to have no front brakes.

WIDTH VERSUS DRAG

I'm preparing an SCCA GT2 car (production car shape fiberglass body, tube frame). I have the opportunity to increase the car's width up to 4 inches, either by splitting the body or extending the flares. I have always assumed that going with the widest width is best chassis-wise, to minimize load transfer and maximize grip. However, I would think that at some point you lose more due to drag than you gain from improved cornering. These cars don't make much downforce, so I don't think there's much gain there from building the car wider. Is there any way to calculate an ideal width?

As long as cornering is an important part of the game, the best approach will usually be to go for greatest allowable width. This is particularly true when power is ample. Certainly, the drag does cost you some speed on a long straight. But if you come out of the preceding turn faster, and you also enter the following turn faster, you are faster not only through those turns, but also over the first and last portions of the straight.

Width is a mixed blessing though, even for cornering. A narrow vehicle lets you take a better line, especially if the turns are tight and the road is narrow. This is one of the main reasons motorcycles are as fast as they are. Narrowness can be very important in autocross, especially for slaloms. A narrow vehicle also can have an easier time passing other cars, which was the reason the FIA narrowed the width limit for F1 cars a few years ago. But on most road courses, the cornering power gained from a wide track width is worth more than the improved line and reduced drag with a narrow track width.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
December 2003

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

RAISING MY RATES

For the first time in three years, I have decided to raise my hourly rate for consulting. The new rate will be $50/hour, which is still reasonable compared to what other consultants have told me they ask. Retainer rates will likewise go up proportionately. A month will be $300; a year will be $1500.

I was considering having the new rate take effect at the turn of the year, but what I'm going to do instead is offer the old rate of $40/hour, and the corresponding $240/month or $1200/year for all services paid for before March 1, 2004.

MORE ON RACING FRONT-DRIVE CARS, AND ON LOAD TRANSFER

Racecar Engineering based my column for the December 2003 issue on my March 2002 newsletter, which dealt with some aspects of racing front-wheel-drive cars. Some readers have written in response to this column, and one of the questions also relates to the basics of load transfer, the topic of the August 2003 newsletter and November 2003 column.

I am a new (1 year) racer who bought a used VW Sirocco mini stock. We race on a slightly banked ¼ mile paved oval. The suspension is VERY stiff, and the right side springs are a higher rate than the left. I do not know the rates, and I plan to install new springs this winter.

Two of the fastest fast guys at the track tell me they use soft suspension, and what perplexes me, their right side springs are softer (by 100 lb.) than the left side springs. Their logic is that the inside wheel will now carry more load as the car moves downward in the corner, therefore giving more equal tire loading and less push on our FWD cars.

Does this make sense to you?

The reasoning doesn't, but running softer springs on the right does, when the track is close to flat – or at least for many cars it does.

Having softer springs on the outside wheels does not reduce overall load transfer. If the turn is slightly banked, and we are comparing setups with the same average spring rate, there will be slightly more roll with a left-stiff setup than with a right-stiff one. Since roll slightly increases load transfer, the left-stiff setup will actually have slightly more overall load transfer – but not by enough to worry about.

On a perfectly flat turn, the car will roll about the same amount with a left-stiff setup of the same average stiffness, but the right wheels will be better able to follow bumps – at the expense of the left ones. Since the right tires provide more than half of the cornering force, it is worth more to have them follow the road surface better.

There will not necessarily be less understeer (push) with a left-stiff setup, in steady-state cornering.

Now, if we were to take your car, which you say is relatively stiffly sprung overall, and we soften just the right front spring, that will reduce understeer. If instead we soften just the right rear, that will increase understeer. If we soften both the right front and the right rear, the effect on understeer will depend on how much we soften each one.

If the track is banked modestly enough so that the left-side suspensions both extend in the turns, we will also reduce understeer if we soften the left front or stiffen the left rear, and we will increase understeer if we stiffen the left front or soften the left rear.

If the track is banked steeply enough so that the left-side suspensions compress in the turns instead (the car still rolls rightward; the right-side suspensions compress more than the lefts), then the effects of left-side spring changes reverse. More left front spring reduces understeer. More left rear increases understeer. Effects of right-side spring changes do not reverse on steep bankings. They just get bigger.

As we encounter steeper banking, it becomes easier to keep the inside rear wheel on the ground. Recall that stiffening the rear suspension relative to the front reduces understeer, at least up to the point where the inside rear wheel lifts. If the turns are banked, we can go much stiffer on rear roll resistance, either with springs or anti-roll bars, before we encounter the limitation of wheel lift. This is very important when racing front-drive cars on ovals.

One other effect of left-stiff springing is that it loosens the car when braking or decelerating, and tightens it under power. This is true for both front-drive and rear-drive cars. So if you enter the turns while slowing, as is usual in oval-track racing, the car will be freer (have less understeer) on entry with a left-stiff suspension. And since a push, once initiated, tends to persist, a freer car on entry may also be freer mid-turn.

One drawback to watch out for when pursuing this approach is that the car will be more prone to inside front wheelspin on exit. This may or may not be a problem on your oval, with a small engine. In road racing and autocross, or on a very small oval with tight turns, it is definitely a factor. I don’t know if your rules allow you to use a spool in that car, or if you are allowed to run a limited-slip. If you are, then you definitely need to pay attention to front tire stagger. Your tire rules may or may not allow you much choice of stagger, but the car will be sensitive to it. With a spool or limited-slip, a larger tire on the right will help the car turn. With an open diff, you want a bigger tire on the left instead.

Finally, be aware that the car is sensitive to static wheel loads, just like a rear-drive car. These work about the same as they do in a rear-drive car. Less diagonal (RF + LR) percentage frees up the car, just like stiffening the left rear spring or softening the right front.

SINGLE OR DUAL REAR BRAKES FOR FSAE?

I am a Formula SAE team member. As you know, some teams use one differential-mounted rear brake disc and some use two discs mounted on the driveshafts or outboard on the wheels. I am having a bit of a time rationalizing the idea of using one rear brake. I spoke to the Wollongong team president at the '03 Detroit event, and he explained that a single rear brake causes corner entry understeer (bad for small-radius SAE courses!) by effectively "locking" the diff. You said basically the same thing in the FSAE message boards. We are using the common Torsen 1 type diff. Does this diff lock in the same way when applying a braking torque as if you were applying an acceleration torque? Looks as if it would. I realize that a clutch diff could act differently under brakingdepending on ramp angles.

Situation:

- *Maximum torque bias ratio 80:20*
- *Diff mounted disc*
- *Car approaches corner. Brakes are applied. Enter left hand turn.*
- *Weight is transferred to front and to right side. Left rear wheel loses most of its normal load and traction.*

Will 80% of the braking force be sent to wheel with traction, while the wheel with low normal load receives 20%? Does the unbalanced braking or the locked diff create the understeer, or is it both?

Now with 2 outboard mounted discs it is obvious that I will always have 50% available rear braking force at each wheel (but different normal loads).

Is this correct? What equations can I use to calculate my braking force, or acceleration force distribution through a locking diff?

P.S. Are you a design judge, and where can I get info on FSAE-sized cam and pawl diffs?

Taking the last items first, a few people have talked with me about the possibility of my doing design judging, but this has all been purely tentative. My understanding is that SAE is no longer even paying expenses for judges, so the judges are actually taking a loss on the activity. I try to get paid if I'm going to have to work.

The only ready-made diffs for FSAE that I know of come from either Quaife or Gleason. These are both Torsen style. The UNC Charlotte team made their own in 2003, using Gleason gears. This was done purely to reduce weight.

You could probably make a ZF-style clutch diff, with the ramps on the pinion shafts, but I think you'd need to do the clutches and carrier yourself, and maybe use spider and side gears from a small car. You could probably also make a Detroit locker style diff yourself, but you'd probably have to make everything.

All things considered, I think it's easier to make or find two small brakes instead.

Somebody may have some equations that describe the behavior of limited-slip diffs, but I don't. Even modeling the friction in these units at known speeds and forces is a bit tricky, because the friction forces are a combination of Coulomb and viscous friction, in proportions that vary with load, speed, temperature, and lubricant properties.

Not only that, the forces are **history-sensitive**! That is, to model or understand the unit's behavior at a specific instant, we need to not only know the speeds and input forces at that particular instant, we need to know the events immediately preceding that instant. For example, suppose we have either a clutch or a Torsen diff, with no preload. If there is no force on the diff, and we jack one wheel off the ground, and apply rotation to the carrier, we just spin the airborne wheel, get no locking, and transmit no torque to either wheel. But if the car is in motion, and the diff is transmitting torque, and then one wheel gets airborne, there is torque on the diff, and therefore loading on the clutches or worm gears, and therefore locking or friction in the unit, at the time the wheel goes airborne. That means we will continue to transmit torque to the wheel that's on the ground, as long as there is no interruption of input torque.

Now, examining the situation you've posited, where we apply the brakes before a turn and continue braking while initiating a left turn, with a single brake acting through a Torsen: First of all, yes the diff does act the same when transmitting reverse (braking) torque. If the diff locking or transfer torque is less than half of the brake torque, both rear wheels are retarding the car (exerting a rearward force), but the torque on the outside wheel is half of the brake torque plus the transfer torque, and the torque on the inside wheel is half of the brake torque minus the transfer torque. If the transfer torque is large enough, it may exceed half of the braking torque. In that case we may have a forward force at the inside wheel. In either case, in a left turn we have a rightward yaw moment due to the diff locking effect, and this moment adds understeer.

It is more common for the inside front tire to reach the limit of grip during trailbraking than for the inside rear to do so. But if the inside rear locks, the outside rear is seeing whatever torque it took to lock the inside rear, plus the transfer torque. In an FSAE car with a single rear brake and a Torsen, there will be more transfer torque if there is continuous braking up to this point than if the brake is applied with the inside rear already unloaded due to cornering.

With two rear brakes, the rearward force is equal on both rear wheels up to the point where the inside one locks. If there is a limited-slip diff, there will be transfer torque. This may be very small, or if there is substantial preload and engine braking, it may be considerable. Even with an open diff, when one wheel locks, the braking forces are no longer necessarily equal on both sides of the car. We can say with certainty that the force at the unlocked wheel is at least as great as on the locked wheel, and possibly greater.

This also applies on the front wheels of a rear drive car, and the above remarks also apply to the front end of a front-wheel-drive or four-wheel-drive car, and to the front end of a rear-drive car with frictional device connecting the front wheels to prevent lockup and flat-spotting of the inside front tire. In all of these cases, the presence of limited-slip or anti-lock transfer torque creates a yaw moment that adds understeer.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January/February 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

VIDEOS STILL AVAILABLE

I still have available videos of my lecture, “Minding Your Anti”, presented March 2003 at UNC Charlotte. Price is $50.00, which includes shipping and handling worldwide. North Carolina residents please add 7½ % sales tax.

RAISING MY RATES

For the first time in three years, I have decided to raise my hourly rate for consulting. The new rate will be $50/hour, which is still reasonable compared to what other consultants have told me they ask. Retainer rates will likewise go up proportionately. A month will be $300; a year will be $1500.

I was considering having the new rate take effect at the turn of the year, but what I’m going to do instead is offer the old rate of $40/hour, and the corresponding $240/month or $1200/year for all services paid for before March 1, 2004.

REVERSE ACKERMANN OR TOE-IN ON OVALS

I race stock cars and am from the old school of using about 1/8” toe-out. Recently, I’ve heard of successful stock car racers using significant straight-ahead TOE-IN (e.g. ½” toe-in). And racers who used to believe in running shorter steering arms on their left front spindle to get more Ackermann are now doing the opposite to get reverse Ackermann. They are using tire steering plates [turn plates] to adjust Ackermann and have the toe they want when the tires are steered, but still use a very small amount of straight-ahead toe-out. It’s all related to tire slip angles, tire temps, optimized handling, etc. I would appreciate a newsletter addressing this topic.

The December 2002 newsletter did address Ackermann a bit, more in the context of road racing. To recap from that issue:

> There isn't a universally agreed way to express how much Ackermann (toe-out increase with steer) a car has. The closest thing we have is to take the plan-view (top-view) distance from from the front axle line to the convergence point of the steering arm lines, divide the wheelbase by that number, and express the quotient as a percentage. If the steering arms converge to a point on the rear axle line, that's said to be 100% Ackermann. If they converge to a point twice the wheelbase back, that's said to be 50%. If they converge to a point 2/3 of the wheelbase back, that's said to be 150%. If they are parallel, that's zero Ackermann. If they converge to a point twice the wheelbase ahead of the front axle, that's said to be –50%.
>
> Supposedly, with 100% Ackermann, the front wheels will track without scuffing in a low-speed turn, where the turn center (center of curvature of the car's motion path) lies on the rear axle line in plan view. This is actually not strictly true, even for the simplest steering linkage, which would be a beam axle system with a single, one-piece tie rod. With either a rack-and-pinion steering system or a pitman arm, idler arm, and relay rod or center link, we can't fully predict what the Ackermann properties will be at all, merely by looking at the plan view geometry of the steering arms. The whole mechanism affects toe change with steer.
>
> Even knowing what instantaneous toe we want in a specified dynamic situation is not simple. We don't necessarily want equal slip angles on both front tires. For any given steer angle, the turn center might be anywhere, depending on the situation. All the infinitely numerous possible situations will have different optimum toe conditions. Therefore, there is no relationship between steer and toe that is right for all situations.
>
> The toe we have at any particular instant results not only from Ackermann effect, but also from static toe setting and toe change with suspension movement (roll and ride Ackermann).
>
> Because of these complexities, there is no single obvious way to define what constitutes theoretically correct Ackermann. It is possible to come up with a rationally defensible definition for your own purposes, but there is no standard rule, and it is unlikely that there ever will be.

With oval track cars, we can have additional complexities. As the questioner notes, it is common to use unequal-length steering arms on oval track cars, usually shorter on the left. In fact, this is the only way to get positive Ackermann on a front-steer stock car, unless the rules allow fabricated spindles and we accept a much larger scrub radius than we'd like. A shorter left steering arm only gives positive Ackermann when the wheels steer left, at the expense of exaggerating negative Ackermann when steering right. This is not an option when the turns go both ways.

The questioner referred me to a website where an expert says that tire slip angle is a property of the tire, which should be available from the tire manufacturer. That is incorrect. Slip angle is an **operating condition** of a tire at a particular instant: the angle between the wheel's aim and its direction of actual travel. A tire does not have a single slip angle. It has a measurable slip angle, when tested on a machine, at a specific load (or normal force), a specific desired lateral force, a specific camber, a specific pressure, a specific temperature, a specific rim width, and a specific wear condition. Change any of these factors, and the slip angle changes. Change the properties of the road or simulated road surface, and the slip angle changes.

The body or sprung mass can also be said to have a slip angle, since it also has a definable direction of aim and direction of travel. An aircraft or watercraft also has this type of slip angle. Slip angle as measured by recently introduced GPS-based data acquisition systems is body slip angle. In a large-radius turn, body slip angle is very similar to rear wheel slip angle, provided there is no large amount of roll steer, ride steer, or static steer at the rear wheels. Front wheel slip angles can be very different from the body slip angle, since the front wheels can steer.

Defining the tire's direction of travel can be a bit enigmatic. At first blush, we might suppose it would simply be the same as the body's direction of travel. For situations where we don't need tremendous precision, and for large-radius turns where the vehicle's yaw velocity is small, this can be an adequate approximation. But when the body has a yaw velocity – and it must have some in any steady-state cornering condition – the tires will not be traveling in exactly the same direction as the body's origin point or CG.

As an example, suppose we are running on a quarter-mile oval where the turns are half of the lap distance. These would be fairly tight turns by oval track standards, about 105 foot radius. Let's suppose for simplicity that the car's CG, or our chosen origin point, is midway along the wheelbase and centered side to side. If the car has a 9' (108") wheelbase, the wheelbase subtends or spans an arc of about 5 degrees on a 105' radius. If the car is driving through this turn very slowly, the tires not sliding significantly, then the center of curvature of the car's path will lie on the rear axle line as seen from above. The body's origin or center of gravity will track outside the midpoint of the rear axle. The midpoint of the front axle will track still further out, and the front wheels will track outside the rears. The front wheels will be steered an average of about 5 degrees to the left. The left front will be leading the right front, and will need about ¼ degree more steer angle than the right front if we want least scuffing, tire wear, and rolling resistance. That's about 1/8" toe-out as measured with typical toe plates, or 1/16" total as measured at the wheel rims using a string or laser.

If we measure the car's body slip angle using a GPS-based data acquisition system, it will tell us the car has a negative slip angle: it is traveling about 2.5 degrees left of the direction it's pointing, while making a left turn!

Now suppose the car is going faster, and the rear wheels need to run at about 2.5 degrees of slip angle to keep the car on course at the speed it's running. The center of curvature will now lie on a line perpendicular to the car's centerline, intersecting the centerline at the CG or origin. The GPS will now tell us we have a slip angle of zero. That won't be what the tires are feeling.

The front tires will now need only half as much toe-out as in the previous case, or about 1/8 degree, for their slip angles to be equal. The rear tires will need about 1/8 degree toe-in for their slip angles to be equal. At the front, the left wheel leads the right slightly. At the rear, the left trails slightly.

Okay, now let's raise the speed again, to a point where the rear wheels have about a 5 degree slip angle. The GPS will now tell us we have a positive slip angle: the car is travelling about 2.5 degrees to the right of the direction it's pointing, while making a left turn. The center of curvature now lies on the front axle line. Neither front tire leads the other. For the front tires to have the same slip angle, they now need to have no toe-in or toe-out with respect to each other. At the rear, the left now trails enough so the rear tires would need ¼ degree toe-in to have equal slip angles.

If we go still faster, the center of curvature moves ahead of the front axle line. The front wheels now need toe-in for their slip angles to be equal.

Now suppose we try a similar comparison, except the track is 2 miles long, and the turns have a radius of 820 feet. The same car's wheelbase now subtends an arc of only about 0.6 degree. The center of curvature is now on the front axle line at only 0.6 degree rear wheel slip or 0.3 degree body slip. Any time the car is cornering hard on this turn, the center of curvature is well ahead of the front axle line, and the front wheels will require toe-in for equal slip angles.

Do we want equal slip angles? Not necessarily.

Although we cannot describe a tire's properties as they relate to slip angle behavior with a single number, we certainly can meaningfully discuss them, and also generalize about them to some extent. For a desired set of conditions, we can measure them on a tire testing machine. For a given set of conditions, a tire will have some slip angle at which it develops maximum lateral force. If we increase slip angle beyond this, lateral force drops off. This is what happens when a tire reaches its limit and breaks away. For street radials at typical loads and pressures, this occurs somewhere around 6 degrees.

In general, bias-ply tires develop peak cornering force at higher slip angles than radials. Narrow tires develop peak cornering force at higher slip angles than wide ones. Tires at low inflation pressures develop peak cornering force at higher slip angles than at high inflation pressures. Tires with thick or deep tread develop peak cornering force at higher slip angles than tires with shallower treads. This is particularly true with treaded tires, but the effect is measurable with slicks also.

As we increase normal force on a given tire, other conditions constant, the slip angle for peak lateral force increases. For most applications, this means that in a left turn, the right front tire develops peak lateral force at a higher slip angle than the left front. This is why we might want toe-in when cornering, or more toe-in than required for equal slip angles.

If both front tires are operating at peak cornering force together, that should be the greatest total cornering force available from the pair. However, we should also think about the longitudinal forces, as these also affect the car's balance. If both front tires are optimized for lateral force, the right one will not only be making more lateral force than it would with more toe-out or less toe-in, it will also be making more drag. This will tend to add understeer. So the toe for least understeer may be a little different than the toe for greatest front lateral force.

We also need to remember that if both tires peak together, they break away together too. So part of the price for greatest peak lateral force is more sudden breakaway. That means that exploiting the lateral force capability may be harder for the driver.

Looking at effects on tire temperatures, in oval track racing the right front tire can easily overheat. The left front seldom gets hot enough to be a problem. Running less toe-out or more toe-in through the turns doesn't help this, it makes it worse.

Whatever slip angles we decide we want on the front tires, we can definitely say that we are going to want less toe-out or more toe-in as the center of curvature moves forward relative to the car. In other words, the larger the radius of curvature or the more the rear tires are sliding, the more toe-in or the less toe-out we want.

What are the implications of this for the optimum relationship between the steering arm lengths? Do we want more positive Ackermann, or less negative, when steering left, or when steering right? I think the tradition of having the left steering arm shorter, if anything, is sound. When the wheels are turned left the most, the center of curvature is furthest aft. When the car is crossed up, steered right in a powerslide or catching a slide, the center of curvature is ahead of the car and we need toe-in. Some experts say having toe-in or less toe-out when countersteering "pins the front end" and spins the car. This is based on the misconception that the wheels always drag more when toed in. As we have seen, this is not the case when the center of curvature is ahead of the car, and the left front is trailing the right front.

So here are my recommendations:

1. Whatever your strategy, the combination of static toe and Ackermann has to give you a good toe value for your prevailing conditions. Wrong Ackermann with a toe setting that compensates is better than improved Ackermann with static toe that doesn't suit.
2. For road racing or street use, where a wide variety of conditions will be encountered, a combination of substantial positive Ackermann and moderate static toe-in is the way to go. This is the prevailing factory approach for road cars, and also, I'm told, the approach taken by Cornell on their winning FSAE cars.
3. For pavement oval-track applications, toe-in may make sense, if combined with positive Ackermann when steering left, especially on high-speed tracks, provided that right front tire temperature considerations don't overrule the decision. Making the right steering arm shorter than the left does not make sense, although it may work as a way to crutch static toe-out on a high-speed track.

It is worth noting that this whole question has been subject to fads throughout the history of racing. In road racing, the first successful rear-engined cars used negative Ackermann. This was in the late 1950's and very early 1960's. The designers claimed this was because the more heavily loaded outside tire required a larger slip angle. In the early '70's, a bit of static toe-out, with zero Ackermann, was the most popular choice. By the 1990's, it was becoming commonly recognized, especially in CART, that high-speed ovals demanded much less Ackermann than road courses, and street circuits demanded the most. Far as I've heard, nobody has yet tried static toe-in for the high-speed ovals, but it's possible my latest information isn't entirely current.

MANAGING THE BEHAVIOR OF THE UNBALANCED

I recently read your article, Load Transfer Basics, in the November issue of Racecar Engineering (Vol.13 No.11)[based on the August 2003 newsletter] and was very interested in the concepts you discussed. I am currently an engineering student in Australia and I am doing some work experience with a Porsche racing team out here. I would be very grateful if you would be able to help me with several questions I have concerning load transfer and suspension set up.

Firstly I was wondering if you could recommend any text books relating to load transfer and suspension set up so I can pursue some further research in this area.

Secondly during a race meeting held a month ago it was mentioned that the way to get faster lap times is to "balance" the car by getting the static load on the front left suspension divided by the load on the back right equal to the opposite diagonal. In other words getting the following ratios equal:

Front Left : Back Right = Front Right : Back Left

Our team thought that this sounded a reasonable proposition but we are unable to see exactly how this works. It seems to me that if this theory does in fact make the car more "balanced" it assumes that the left handed corners are similar to the right handed corners. Since this is often not the case surely the car must be preloaded to suit the important corners on the track, in which case the broad statement of balancing the car isnt helpful...?

I like the chapter in Milliken & Milliken's *Race Car Vehicle Dynamics* by Dave Segal on the subject (Chapter 18). The book is published by SAE, and available from their bookstore at www.sae.org. I would also of course not pass up the opportunity to plug my own video on the subject, mentioned at the beginning of this newsletter.

Another way of stating the relationship you were told would be to say that the rear percentage taken diagonally is the same for both diagonal wheel pairs.

For a car with 50% left, this works. You would have a car with 50% diagonal, the same rear percentage on both sides, and the same left percentage at both ends. Such a car should corner with similar balance in right and left turns.

However, as you note, sometimes even in road racing we want the car heavier on the side it turns toward the most, on a particular track. It may also happen that we get something other than 50% left unintentionally, due to packaging or rules constraints. We still will normally want the car to corner similarly in both right and left turns, even though it will be faster when turning toward its heavy side.

My recommendation for this is to start with what I call an unwedged car. My definition of this is a car with equal rear percentage on both sides, and equal left percentage at both ends. This is different than what your instructor taught you, and also different from 50% diagonal weight.

Let's consider an example. Suppose we have a 2000 lb. Porsche with 60% rear, and we are racing on a road course with predominantly right turns. Suppose the car has right hand drive, and we have enough ballast so we can get 55% right. An unwedged car, as I define it, would have 55% of 60%, or 33%, on the right rear (660 lb.); 55% of 40%, or 22%, on the right front (440 lb.); 45% of 60%, or 27%, on the left rear (540 lb.); and 45% of 40%, or 18%, on the left front (360 lb.). The car has the same 45% left at both ends, and therefore LF/RF = LR/RR. Also, it has the same 60% rear on both sides: LF/LR = RF/RR. This setup imposes no static torsional load on the frame.

Note that we can get either of these equations from the other by algebraic manipulation. (Starting with the former, we multiply both sides by RF, and divide both sides by LR.) They are different forms of the same expression. However, there is no way we can manipulate either equation to obtain the form LF/RR = RF/LR. That is a different equation.

If we look at the ratios in your instructor's equation, in the above example, we get: LF/RR = 18%/33% = 360/660 = .545; RF/LR = 22%/27% = 440/540 = .815. Not close at all.

Suppose we look at the above example with reference to the target more conventional wisdom would set, namely 50% diagonal. That would mean 2(RF+LR) = RF+RR+LF+LR. Our example doesn't fit that rule either. Its diagonal percentage is 27% + 22% = 49%. Close, but not the same.

Diagonal weight would be 980 lb. To increase that to 1000 lb. and make the diagonal 50%, we would have to adjust the suspension to transfer 10 lb. left-to-right at the front and transfer 10 lb. right-to-left at the rear (assuming equal track widths, for simplicity). That would mean that the static left percentage at the front would be 350/800 = 43.8%; left percentage at the rear would be 550/1200 = 45.8%. If the car has equal roll resistance in both directions front and rear, we can see that the change we've made to the static load distribution will help the front wheels in a right turn and the rear wheels in a left turn. We would therefore expect more understeer in left turns than in right turns, which is probably not what we want.

Now, suppose we wanted to set this car up to your instructor's rule. Could we even do it? Well, yes, we could, but it would take a pretty freakish setup to do it. We'd have to have 600 lb. on the LF, 900 lb. on the RR, 200 lb. on the RF, and 300 lb. on the LR. This is obviously not what we want. Our diagonal percentage would be 25%. Our left percentage would be 75% at the front, 33% at the rear. Yow! The car would push like a dump truck in right turns and spin if you sneeze in left turns.

Rather interestingly, if the car has 50% left, all three rules work. If the car has either 50% left or 50% rear, the 50% diagonal rule works. If the car has neither 50% left nor 50% rear, my rule works, or gets you closest.

Even with my rule, small adjustments may be needed. If there is more understeer turning right than turning left, add load to the RF and LR. If there is more understeer turning left than turning right, add load to the LF and RR.

INDEPENDENT OR BEAM AXLE FRONT SUSPENSION FOR PAVED OVAL?

We race primarily on quarter-mile asphalt ovals – no banking, no bumps. The car weighs 700 kilos (1500 lb.) and runs on 10-inch Hoosier slicks front and rear. Everyone uses A-arm front suspension, with a mandatory beam axle rear. Would there be any advantage running a beam type front end? If the steering rack is mounted on the axle, would the driver feel roll-induced steer?

The main advantage of a beam axle front end for an oval-track application is that it reduces overall weight, and also usually reduces the torsional stiffness needed from the frame. Even if you have a minimum weight requirement, reducing car weight is advantageous because you have more ballast to move as desired. This is especially beneficial if you have no limit on left percentage, or if an engine setback limit prevents you from getting as much rear percentage as you'd like.

A beam axle also eliminates camber change due to roll, which is good. However, on an oval, independent suspension can be set up with sufficient static camber so that you can also obtain desired camber when cornering with an independent system.

Back when I was doing consulting for fun and experience, I worked with a west coast super modified team. At that time, there were no rules on suspension design. Beam axles have now been made mandatory front and rear for these cars. (Interestingly, in the US, other classes prohibit beam axles and mandate independent front ends. There is no universal consensus as to which is better.) Back in the '80's, people ran both styles, and the usual practice was to mount the rack on the axle on the beam axle cars. This works fine. The driver does feel a little roll steer, but it's not bad, and it's in the right direction, i.e. roll understeer.

The main advantage of the independent setups in this class was that the car was more controllable at the point of inside front wheel lift. With a beam axle, especially with a high roll center and a soft wheel rate in roll, when the inside front wheel comes off the ground, the roll center rises as the axle rises, causing the wheel to rise further. The wheel tends to rise a lot then, causing unfavorable camber on the outside front, and therefore understeer. It is difficult for the driver to maintain an attitude where the inside front tire is carried just a little way off the ground. With independent, this is much easier. However, I believe that with a low roll center and a high wheel rate in roll, a beam axle could deliver similar controllability. And if your cars never reach the point of carrying a wheel, the whole issue is moot.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

VIDEOS STILL AVAILABLE

I still have available videos of my lecture, “Minding Your Anti”, presented March 2003 at UNC Charlotte. Price is $50.00, which includes shipping and handling worldwide. North Carolina residents please add 7½ % sales tax.

RAISING MY RATES

For the first time in three years, I have decided to raise my hourly rate for consulting. The new rate will be $50/hour, which is still reasonable compared to what other consultants have told me they ask. Retainer rates will likewise go up proportionately. A month will be $300; a year will be $1500. New rates take effect March 1.

DIFFERENTIATING DIFFERENTIALS

A 2200 lb., 300 bhp, rear-wheel-drive car is doing circuit racing on a short, twisty asphalt circuit with mostly right-hand turns. It runs the same rubber on all four wheels. Which of the following differential setups would be quickest around the track, giving good turn-in and good traction out of the corners?

1. *A spool type locked rear*
2. *A Positraction rear with 250 lb. preloading*
3. *A Detroit Locker rear*
4. *A Torsen rear*

All of these options have their adherents. One thing that complicates the picture is that the choice interrelates with setup and driving style.

Conventional wisdom is that spools are a bad idea for road courses. To get good steady-state cornering with a spool, the car needs to have tire stagger, or alternatively the power available and the nature of the track surface have to be conducive to powersliding. Tire stagger is generally an impossibility if the turns go both ways, although maybe if most of them go one way it is possible to accept poor behavior in the few that go the wrong way. If, as the questioner posits, the tires are truly identical all around, in circumference as well as all other properties, then we definitely don't have stagger.

Driven more or less normally, a car with a spool tends to understeer, or push. It does this much more in tight turns than in large-radius ones. We can free it up by putting lots of roll resistance in the rear suspension, but when it's right in the slow turns, it's too loose in the sweepers. One cure for this is to have a lot of aerodynamic downforce at the rear. Whether we can get that will depend on the bodywork rules.

There are situations where the locked-axle push is helpful. If we are trying to brake and turn at the same time, the car tends to oversteer. This limits how hard we can brake while turning. The car also will oversteer under power, if there is enough power applied, because the rear tires are using a lot of their available grip to make forward force and have less grip available for making lateral force. Again, if the car is tight to begin with, the driver can feed it more power before it goes power-loose.

This means that a spool can work on a road course if we have a driver who trail-brakes deep into the turns, and then gets on the power hard right away. In most of the slower turns, the car never sees steady-state cornering when driven this way. Mid-turn speed isn't necessarily best with this approach, but we get good entry and exit speed, and a late brake application point and an early power application point. Consequently, elapsed time, or average speed, on any straightaways before and after the turn improves.

Not surprisingly, the driver best known for making this work is the one who popularized trail-braking when most drivers were still completing their braking before turning: the late Mark Donohue. Donohue could win races on a road course with a spool, even in a car with huge tires and a massive rear wing, such as the Can-Am Porsche 917-30. Other drivers would get into cars set up to his liking, and not be able to do anything with them.

We might reasonably suppose that a driver could simply learn this, and adapt. Actually, drivers' abilities to do this vary considerably, and even those who can learn require practice to use a new technique well. Then again, we may have a driver who learned to drive this way, and has to learn a new style to do anything else. Not only do personal preferences differ on the question of trailbraking, so do driving schools.

To perform the best, a race driver needs to be able to drive the car without thinking, and focus his/her conscious mind on observation. This means that it's not easy to learn and unlearn new driving styles to adapt to different cars. For some talented individuals, it's merely difficult. For others, it's impossible. And for all drivers, just having to think about technique costs speed, all by itself. Therefore, when we set up a race car, we need to accommodate the driver, and not simply write off to stubbornness a driving style that doesn't suit our setup. If we have a driver who is in the habit of driving like Donohue, a spool may be worth considering. If not, that argues against the spool.

The spool is simpler and lighter than any alternative. Its simplicity is a plus for both cost and reliability, although a spool is generally harder on axles than any form of limited-slip.

For those unfamiliar with the Positraction, or Posi for short, it is a clutch pack style limited-slip, usually with a single clutch pack establishing friction between the right side gear and the carrier. The clutch pack is preloaded either by having a dish in one pair of discs so they act like Belleville washers, or by having coil springs bearing on the clutch pack. Added clutch loading is applied to the pack by the spreading force on the side gear when torque acts on the ring gear. At a given ring gear torque, this spreading force depends on the tooth profile and the diameter of the side gears. That makes the preload the only adjustment.

Preload is measured in lb.ft. of torque, rather than pounds – though it is common to say pounds for short in casual conversation. The procedure is to jack up one wheel, or jack up both rears and have a helper hold one, put the transmission in neutral, and measure the torque required to turn one wheel. This requires an adapter to allow a torque wrench to turn the wheel. For best accuracy, it's best to turn both wheels at once with the torque wrench first, to measure brake and bearing drag. Then you divide that value by two to find the drag for one wheel, and subtract that from your torque reading when turning just one wheel.

250 lb.ft. in this test is a lot. Typical values for stock road cars are more like 50. If the car is on racing slicks with a coefficient of friction of 1.3, and a rear wheel has a static load of 550 lb. and an effective radius of one foot, the tire's breakaway torque under static load is 715 lb.ft. For the inside tire in a corner, with a substantial portion of the traction circle being used for cornering, 250 lb.ft. could easily be enough to overpower the tire. And under power, we get more clutch pack loading. In such a situation, a heavily preloaded Posi can act a lot like a spool.

However, when the tires can overpower the clutch pack, the Posi acts tamer than a spool. Like a spool, but less severe. This will tend to save the axles to some extent, as well.

It might be worth mentioning the ZF style differential, which the questioner didn't list. The ZF and the Posi are both sometimes called Salisbury differentials, but they differ a bit. Like the Posi, the ZF has multi-disc clutches that can be preloaded and are loaded additionally by engine torque. The difference is that in a ZF, the clutch discs are outside the carrier, and the carrier is split into two halves at the pinion shafts. The pinion shafts have angled flats on them that bear against mating flats on the carrier halves. The angle of these flats can be varied, by installing different parts having different angles. This allows us to adjust the severity of the lockup under power. There are similar flats that spread the carrier halves under reverse torque (engine braking). The angles of the flats for power and decel can be varied independently, allowing the unit to lock more or less strongly on deceleration compared to acceleration.

The Torsen, or Gleason, is also two designs that have very similar properties. Both use worm gears in place of the spider and side gears of a conventional diff. The worm gears provide a very smooth, yet strong lockup under load, yet turn very freely with no load, provided they are not preloaded. Gleasons can be preloaded. One problem with preloaded Gleasons is that the preload is highly sensitive to gear wear.

If I were to make a general-purpose recommendation for road racing, it would be the Gleason. It has the ability to lock strongly, yet smoothly, with little or no preload. This makes it very driveable.

It does have some drawbacks. It is generally the costliest of all the types we're considering here, although not prohibitively so. The power and decel lockup are not separately tunable, as with the ZF. If the unit is not preloaded, it will not prevent one wheel from spinning if that wheel is very lightly loaded, or is airborne, or is on a very slippery surface.

The Detroit Locker is not really a differential in the sense that the Salisbury and Gleason are. That is, it has no gears at all, and there is no way it can be set up to split torque equally between the two output shafts, while letting their speeds vary, even at very low torques. The locker contains a center element consisting of a dog ring driven by the carrier, like the spider gears and pinions in a Salisbury or an open diff. This central dog ring has dogs on both sides. These mate with driven dog rings on either side, which drive the axle shafts.

The driving dog ring can float a bit side to side. It is held centered by two conical coil springs. When the driving dog ring is centered, it engages both driven dog rings, and we have a locked axle. If the driving dog ring moves to one side or the other, it moves more deeply into engagement with the driven dog ring it moves toward, and if it moves far enough it disengages from the other dog ring. We then have drive to only one wheel.

For a wheel to disengage, it has to overrun the carrier. For this reason, the locker is sometimes called a ratchet. This isn't really accurate, but the unit is somewhat similar to a pair of ratchets, each driving one wheel, in that it drives the slower wheel and lets the faster one overrun. It differs from a pair of ratchets in that only one wheel can overrun at a time. In decel, the slower wheel sees the engine braking. In a race car, this promotes very free turn-in. If the driver likes to finish braking and then turn, this can work well. If the driver likes to do heavy trail-braking, it may be a disadvantage.

When the driver gets on the power, the inner wheel drives, up to the point where it spins. As soon as the inside wheel reaches the speed of the outside wheel, the unit locks and drives both wheels. The lockup is not smooth at all. The dogs are either engaged, or they're not. The unit cannot slip. This requires the driver to develop a feel for when the unit is going to lock, and anticipate the change in car behavior at lockup. It also rewards decisive driving. That is, lockers are somewhat unpredictable if the driver is on and off the throttle trying to balance the car. It can be hard to predict whether the rear will be locked or unlocked when the power is reapplied after a brief lift. So the locker responds best to a driver who gets on the power and stays on it.

One peculiarity of the locker is that when backing up, it drives the faster wheel rather than the slower one. This means that we don't have drive to the wheel with traction if one wheel is off the ground or on very slick ice. Normally, this is of no concern in road racing, but it is something to consider for street or off-road use.

The locker develops less heat than a Salisbury or Torsen, since it has no slipping friction elements. The rear end can still get hot, due to the friction in the ring and pinion gears. And if it gets hot, the springs in the locker can lose their temper. This will cause erratic locker behavior. For this reason, it is customary in racing to replace the springs frequently.

Lockers are universal in NASCAR, because the rules require them. Salisburys and Torsens are prohibited. This rule originated with a prohibition against all limited-slips and spools. The inspection procedure was to test the rear end by turning one wheel with the transmission in neutral, as described above for testing a Posi. The wheel had to turn freely. Since a locker would pass this test, people started using them, and they were such an advantage over an open diff in a stock car that they became universal. NASCAR saw that the racing was better when the cars could put power down with both rear wheels, so they never prohibited lockers. Now, they are specifically written into the rules by name.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2004

Please note: This document has been revised to correct an error. Corrected text appears in red.

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

EFFECT OF CALIPER MOUNTING POSITION

What effect on wheel loading does the positioning of the calipers in a leading or trailing location have – i.e. mounted at 3 and 9 o'clock positions? Does a trailing caliper add or subtract load on the front tires? In a rear independent suspension, does a leading caliper add or subtract wheel loading, and is it the same in a live axle situation?

The short answer is no. Caliper location has no effect whatsoever on wheel loading. Having the caliper's mass lower or higher does have a very minute effect, because it affects the CG location a tiny bit, but there is no difference between a 3 o'clock mounting position and a 9 o'clock position.

However, there is an effect on bearing loads. It might seem counterintuitive that we can change the bearing loads and not change the tire loads, but that is in fact the case. As the questioner appears to have considered, the disc tries to carry the caliper upward if the caliper is trailing, and downward if the caliper is leading. That reduces bearing loads if the caliper is trailing, and increases bearing loads if the caliper is leading. However, these forces are reacted entirely within the hub/bearing/spindle/upright/caliper/disc/hat assembly, and do not change the loads on other parts of the car.

We can think of it like this: Gravity acts downward on the car, with additions and subtractions due to inertia effects and aerodynamic effects. The road surface holds the car up. Or, we may say the road holds the tire up; the tire holds the wheel up; the wheel holds the hub up; the hub holds the bearings up; the bearings hold the spindle up; the spindle holds the upright up; the upright holds the suspension up; the suspension holds the sprung mass up. If the caliper exerts an upward force on the upright and a downward force on the disc, that just means the brake is helping the bearings and spindle hold the upright up. It doesn't change the total support force, only the load path within some of the unsprung components.

It is worth noting that in braking there are also horizontal forces acting through the wheel bearings. The car is trying to keep going forward at a constant speed. The road surface is exerting a rearward force on the car, through the tires, wheels, hubs, bearings, spindles, uprights, and suspension. We can reduce the bearing loads due to this component if we mount the caliper below center, or increase the bearing loads if we mount the caliper above center. In fact, the horizontal force may be greater than the vertical force on the tire. With racing slicks on dry pavement, the horizontal force may be 1.3 or more times as great as the vertical load on the tire. So for least bearing loads during braking, the caliper should be somewhere in the upper rear quadrant – around 5 o'clock or 7 o'clock, depending on which wheel we're looking at, and from what direction.

Now, do we actually want maximum cancellation of the bearing loads by the brakes? We might suppose so, but actually there is an argument for not having maximum cancellation. The effective radius of the brake (roughly the radius to the middle of the pad) is often less than half of the tire effective radius. This means that the force at the caliper is more than twice the rearward force at the tire contact patch, and it may also exceed the vector sum of the vertical and horizontal forces at the contact patch. Consequently, the caliper force may not only reduce the bearing loads, but reverse them. If there is any free play in the bearings, or deflection in the components, this load reversal may result in a vibration or a small variation in the steer angle of the wheel. So there is a case for building the components nice and strong, and positioning the calipers so the bearing loads will not reverse.

Of course, as a practical matter, if we are using purchased calipers we need to mount them with the bleed screws at the top, or very nearly so, just to facilitate good brake bleeding without requiring the calipers to be dismounted. This may well outweigh any theoretical considerations. If we are designing from a blank sheet of paper, we don't face this constraint, but most of us, most of the time, are designing around purchased calipers.

Another practical constraint is packaging, particularly of the steering arms and cooling ducts.

There are some ways in which we can affect wheel loads by the design of the brake system and the suspension. I am referring here to the longitudinal "anti" or "pro" effects: anti-dive or pro-dive in the front suspension, anti-lift or pro-lift at the rear. With independent suspension, it makes a difference to these effects whether the brakes are inboard or outboard. With a beam axle, it makes a difference if the calipers are mounted directly to the axle, or on birdcages or floaters that rotate on the axle and have their own linkages.

However, with all of these, we cannot significantly alter the loading on the front or rear wheel pair, nor on all four wheels. We can change the way the sprung mass moves in response to braking, and this may have small effects on CG height, with corresponding small effects on overall load transfer. But the big effects come from having geometry differences on the right and left sides of the car. These may be present even in supposedly symmetrical road racing cars, because no car stays symmetrical when it rolls. In oval track cars, we often design in, or adjust in, asymmetry even in the static condition. Such asymmetry can produce significant changes in diagonal percentage when braking, and we can use these to tune corner entry behavior.

All such effects are independent of the "clock" position of the caliper mount.

MORE ON WEIGHT DISTRIBUTION

In the April 2004 issue of Racecar Engineering [see November 2003 newsletter], *you mentioned that a 52/48 front to rear percentage works well on medium to high speed tracks. How does it change when you go to a short track? Is there an optimum left/right percentage and cross weight for medium speed track with moderate banking (18 degree corners and 12 degree straights)?*

Regarding rear percentage, it depends to some extent on the design of the track, but assuming that we are required to use equal size tires front and rear, and assuming the bodywork rules permit only very limited downforce, something close to 50/50 works well for mid-turn speed. More rear helps braking. More rear also helps forward acceleration, provided that the car is traction-limited.

So if that short track is bowl-shaped – short straightaways and long turns, no really straight running, small speed variations between mid-turn and end of straight – we want close to 50/50. If the track is paperclip-shaped – tight turns, long straights, two drag races and two hairpins per lap – we want more than 50% rear, especially if the tires are narrow and hard, and the car has lots of power.

If the track has more banking, that tends to make the car less prone to wheelspin, at least during corner exit. That reduces the need for rear percentage greater than 50%.

Regarding left percentage, there's no such thing as too much, at least within most rules. Ordinarily, you are limited either by an outright limit on left percentage, or by a minimum right side weight, or by rules on engine and frame offset. In west coast supermodifieds, these days you have a choice of two left percentage limits. If you use the lower left percentage, they let you run a bigger wing.

Having more left percentage does change the car's behavior, so it only makes the car faster if the setup is suitable. It can happen that a team will reduce left percentage and find speed, but this is because they haven't figured out how to make the left-heavy setup work.

When the car has lots of left percentage, it tends to turn right when braking and turn left under power. That makes it tighter on entry and looser on exit. So the key to making the car work is to compensate for this, but not overcompensate.

Regarding diagonal percentage, there is no ideal amount just based on the track and type of car. It interrelates with the rest of the setup. If you have more roll resistance at the rear compared to the front, you need more diagonal, to keep the same amount of understeer. If the car goes loose on slick, add diagonal percentage and add rear roll resistance or reduce front roll resistance. If the car goes tight on slick, do the opposite.

SPRING SPLITS WITH BIG BARS

I just read your article in the Racecar Engineering October issue [see July 2003 newsletter], *about stock car setups where the right rear spring is the stiffest, due to the front end having soft springs* *and a stiff anti-roll bar. You say that you have some clients using this type setup without the right-stiff spring split. How are you doing that?*

What effect does the left rear spring have on these setups, and does a stiffer LR tighten or loosen the car on entry?

We are running on a D-shaped 5/8-mile track, with a maximum banking of 6 degrees, minimum of 2.8 degrees. The surface is old, worn out, and bumpy. The car is a Howe XL Late Model with coilovers all around. It weighs 2900 pounds, 56% left side, 58% diagonal.
Spring rates: LF 200 RF 200 LR 150 RR 200
Front bar: 1.250 dia., 9.5inch arm Track bar height: 11.5 inches

To answer the first question, of course if the RR spring is the stiffest on the car you do have a right-stiff rear spring split, but you can have the car rear-stiff in terms of roll resistance with geometry, or a rear anti-roll bar – or with a stiff LR spring if the turns are flat enough so the LR extends in the turns.

Those are definitely soft springs for the weight of the car, even allowing for using coilovers, which improve the motion ratios compared to a big spring car. I can't actually calculate the wheel rate contribution from the front bar without knowing its length and the motion ratio from the arm end to the wheel, but it appears pretty substantial compared to the springs. The Cup cars are using much stiffer ones, though. Their diameters can get as large as 2 inches in some cases, and the wheel rates in the rear are also higher than yours. They sometimes use rear anti-roll bars too.

As to what happens in your case when you stiffen the left rear spring, all I can say for sure is that any time the LR spring is compressed relative to static position, a stiffer spring will add diagonal and usually tighten the car. Any time the LR spring is extended relative to static, a stiffer spring will reduce diagonal and loosen the car. These effects can sometimes reverse during early entry, if the car is being slowed mainly by the rear wheels.

To really know what that spring is doing, you need to have data acquisition. The best way is to have motion sensors and electronic data logging like the big boys use, but you can also improvise with video cameras. Either mount a camera under the car aimed at the spring, or clamp a piece of welding rod to the axle near the spring, with the rod poking up into the interior through a hole. Then mount the camera in the interior and aim the camera at the rod.

My guess would be that the spring is extended most of the way through the turn, with the greatest extension occurring during late entry, and the least extension during exit. Based on that assumption, I would predict that a stiffer LR would loosen the car through most of the turn, with the greatest effect about ¼ of the way through, and the least effect on exit. If you add more diagonal or more front roll resistance to tighten the car back up, you may have a condition where entry is looser, mid-turn is similar to before, and exit is tighter. Remember, this prediction is based on a number of assumptions, so if the driver reports different results, don't automatically assume he's wrong.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE ON LAST MONTH'S TOPIC

Last month's newsletter contained a minor error. To reduce the longitudinal or X-axis forces on the wheel bearings during braking, the caliper should be below center rather than above. I am sending out revised April issues with this issue, and the archived April issue incorporates the correction. Thanks to Eric Zapletal in Australia for catching me on this.

Also, Ramon Mendoza at Bridgestone/Firestone pointed out that caliper location has some effect on steering feel. Centrifugal force acting on the caliper during cornering creates a torque about the steering axis. When the caliper is ahead of the steering axis, this adds steering effort. When the caliper is behind, it reduces steering effort. In general, the increase is better than the decrease, because the driver is somewhat better able to feel the lightening and heavying of the steering as the front tires dance back and forth across the limit of adhesion. We might also conclude that having the caliper closer to 6 o'clock or 12 o'clock might be desirable in this regard, as it would reduce the magnitude of the effect.

THE NEW NASCAR TIRE

A number of people have written me with questions about the new NASCAR Nextel Cup tire.

NASCAR has a new tire. The common thread is a new soft sidewall. What happens to spring and shock rates to adjust for sidewall flex?

On the NASCAR commentaries the announcers say the teams are adjusting their tire pressures by ½ to 1 pound increments. I am assuming the increase in pressure makes the tire stiffer, hence making that corner stiffer – lowering the pressure vice versa. I would assume that these small adjustments are from the optimum pressure required to give good tire footprints with even temperatures. Would the same scenario apply to other radial tires, including street tires, or are the NASCAR tires unique in their sensitivity to changes?

Darrell Waltrip described the old tire (hard sidewall and compound) versus new tire (soft sidewall and compound) as hard Jello versus soft Jello. Larry McReynolds said that teams were looking at control arm angles. It would seem to me that the tire needs to be tilted out at the top (positive camber) to compensate for the sidewall flex.

I don't have any special information on these tires. I'm relying here on what's publicly available.

I understand that the teams are using significantly higher pressures with the new tires. I've heard figures as high as 7 psi more. I am therefore unsure whether the sidewalls really do flex that much more, as the tires are actually run.

All tires are quite sensitive to pressure. Part of the reason small differences matter so much in racing is that the competition and on-the-limit operating conditions make small changes in car balance more noticeable to the driver.

A tire has an optimum pressure for greatest lateral grip and another, lower, optimum pressure for longitudinal grip. Above or below optimum, grip diminishes. This effect is more important than the effect that pressure has on tire loading by softening or stiffening a corner of the car.

The key to understanding tire pressure settings in stock car racing is to recognize that the tires are all overinflated when hot, and this is unavoidable because if they are run any softer they are unmanageable right after leaving the pits. If tire warmers were allowed, this might not be so. But under existing rules, softer inflation improves grip once the tire is hot – not so much because it makes the carcass a softer spring, but more because it results in hot pressure closer to optimum.

Anyway, to adjust spring rates for a more compliant tire, yes you would go stiffer. Or maybe you wouldn't if the rules include a minimum spring rate and a minimum ride height, and you're trying to get the car to run lower. But barring special considerations, you'd go stiffer.

As for camber and control arm geometry, in general more compliant tires want more aggressive camber when cornering, and are also more tolerant of camber when running straight. On an oval track car, that means more positive camber on the left front and more negative on the right front.

Pavement stock cars were already surprisingly aggressive in terms of "camber gain". Front view swing arm lengths substantially less than the track width have been the norm for years. This is much shorter than used in any other kind of racing. I don't know how much shorter they can go.

One point I've thought the NASCAR folks are missing has to do with left front suspension geometry. On banked tracks, and even on relatively flat ones when running big bars and soft springs, the left front suspension compresses in the turns, rather than extending as it would in pure roll. In such a condition, a geometry that places the front view instant center near the centerline of the car actually hurts cornering camber instead of helping it. To get the car behavior and tire temperatures we want, we then have to run extremely aggressive static positive camber. We could actually get more favorable camber change in the turns with an instant center to the left of the car. That would involve having both arms slope up toward the frame, the upper one more steeply. If that isn't possible with legal spindles, a long front view swing arm with an instant center far from the car (implying control arms close to level) would give a smaller unfavorable camber change than the short front view swing arms usually seen.

Note that this does not apply for the right front, nor does it work for road courses – which brings us to a related question from another reader:

CONTROL ARM ANGLES IN FORMULA 1

I've just been watching F1 qualifying [this was written in September 2003]*, and I noticed that the front upper wishbones on the Jaguar slope down from the chassis to the uprights. I had seen some*

cars with them roughly level before, but the ones on the Jag had a very pronounced drop to them. This seems to fly in the face of conventional wisdom of camber gain in bump. Is it because the tire construction is different, or are the suspensions getting to the point where there just isn't enough travel to worry about camber gain? After seeing this, I looked at the other cars again. Most wishbones were angled in the "normal" direction, with the Ferraris' being the most pronounced. Is Jag onto something (points don't seem to reflect this)?

I must confess that I share the questioner's mystification regarding the front upper control arm angles we have been seeing on F1 cars over the last ten years or so. This season I'm noticing a similar slope on the Williams cars. More commonly, the arms are very close to level, much more so than they were up to around 1994.

It is true that geometry makes less difference when the suspension moves very little. I am fond of quoting Colin Chapman's famous aphorism: "Any suspension will work if you don't let it."

In spite of that, I find it hard to imagine why a designer would want camber to go toward positive as the suspension compresses and toward negative as the suspension extends – except, as noted above, on the left wheel in an oval track suspension, when the left suspension compresses in the turns.

Tires do have different preferences regarding camber, but none of them like to be tilted out of the turn. They all like to be tilted in the direction of the turn at least a little, and they all like to be close to upright when running straight. Front view swing arm length is therefore a compromise between creating a camber change that compensates for adverse camber change due to roll in cornering (I prefer to call it camber recovery rather than camber gain), and minimizing camber change when roll is absent, or over bumps when cornering.

If the front view swing arm length is "negative" (instant center outside the track width, on the same side of the car as the wheel), we get camber change greater than the roll angle in cornering, in the wrong direction, and also get camber change in ride. The camber change in ride can be beneficial in certain circumstances -- specifically in braking on the front wheels, when negative static camber is used. Of course, this approach also requires more static negative camber to get acceptable outside wheel camber in cornering, and this produces more unfavorable camber on the inside wheel than we'd have with more conventional geometry. Overall, I don't see a gain.

If the control arms are close to horizontal, that makes the system insensitive to ride motion, which is good, especially in high-downforce cars. I am still inclined to allow a bit of camber change in ride, to reduce camber change in roll – just not a lot – and limit roll with high wheel rates in that mode.

It should be mentioned that camber characteristics also depend on the angle of the lower arms, which are hard to see from trackside on an F1 car due to the wings. It would be possible in theory to have the upper arms slope down toward the spindles, and still have the camber go toward negative in compression, if the lower arms sloped the same direction but more steeply. However, this would produce a very high roll center, and I very much doubt that the F1 teams are doing that.

ARM ANGLES IN TRIANGULATED FOUR-BAR BEAM AXLE REAR

I run a figure-8 car – we have to turn left and right. We have to run a stock four-link rear end. Ours is out of a '78 Impala. We have moved some of the suspension points, but my question was what angle I should run the lower trailing arms. I'm trying to get the tops around 10 degrees.

For readers unfamiliar with this setup, this is a live axle locating linkage with four semi-trailing links, and no purely lateral links. The upper links converge to a point above and slightly behind the axle. The lower links converge to a point well forward of the axle, and lower. These two convergence points define an axis about which the axle moves in roll. The roll center is considered to be the point where this axis of rotation intersects the axle plane (vertical plane containing the axle centerline).

There is no theoretically ideal angle for any of the links. Ordinarily the lower links slope up a bit toward the front. In a three-link rear, this would produce roll oversteer (rear wheels point toward the outside of the turn when the car rolls), but in this type of linkage we just get reduced roll understeer. If we level the lower links out in side view, we actually get more roll steer.

The side-view angles of the links also affect anti-squat under power and anti-lift under braking. Leveling out the lower links reduces both of these effects. That can help the car if wheel hop under power or in braking is a problem.

My general-purpose recommendation is to run all the links close to factory angles, or level them out a little. I wouldn't make any of them steeper. I wouldn't aim the lowers down at the front. There are no miracles available here, but as long as you don't make the geometry extreme in any regard, you won't have a disaster.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE ON F1 CONTROL ARM ANGLES

Last month we discussed F1 control arm angles. This question from a reader relates to that. Additionally, after more examination of photographs of the 2004 Williams, I need to correct a misconception held by this reader, and, until recently, by me.

With regards to the recent discussion about the F1 control arm angles, could they be using the "negative" swing arm to achieve a below-ground roll center? I suppose there are arguments for both below-ground and above-ground roll centers at the front – seems to be a preferential design variable. I notice that the Williams does not droop limit (observation during pit stops) [The questioner means that the car does have some droop travel from static ride position, although not necessarily to the point where the springs completely unload.], *so the geometry would be susceptible to lateral migration of the front roll center in roll, though with modern design tools this can probably be optimized.*

In the attached picture, the front legs of the wishbones are nearly parallel, but it is hard to discern the front view swing arm geometry without knowing the attachment of the wishbones to the tub.

I am inserting the questioner's photo on the next page. It shows the Williams front end , from the front, with the car on stands, wheels off, and suspension at full droop. The appearance is that the upper and lower control arms are approximately parallel in this condition, and both are lower at the outboard end than they are where they attach to the tub.

After studying other pictures of the car, taken from other angles, I have concluded that the angle of the upper control arms is not what you'd think from a frontal shot. Those airfoil-shaped links with the Michelin decals on them are not the front legs of the top wishbones; they are the tie rods, or steering links. And they are not parallel to the front legs of the wishbones, which lie immediately behind them. At static condition, the upper and lower control arms are very close to horizontal. The front legs of both the upper and lower wishbones are very nearly transverse, and there appears to be little anti-dive, so the geometry of the front legs of the wishbones very closely approximates the front-view projected geometry of the suspension.

So, unless I'm still missing something, for small displacements from static we have geometry that produces little camber change in ride, camber change in roll approximately equal to angular roll displacement, and apparently a substantial amount of bump steer (!).

I hestitate to attribute a car's racing record to any single observed design feature, because I know that results come from the whole package. However, even with the modest amounts of suspension movement seen in F1, significant bump steer can't be good. If the bump steer characteristics are what they appear, that could easily explain the team's difficulties making the car work this year.

Looking at other aspects of the geometry, it is evident that the length of the front view projected control arm is shorter for the lower arm than for the upper, which is unusual. I am certain that the reason for this is aerodynamic, not mechanical. The designer has chosen to compromise suspension geometry in pursuit of aerodynamic efficiency, particularly that of the center portion of the front wing. The idea is to get all aerodynamic obstructions in line with the "tusks" that carry the front wing, and have as large a clear region as possible aft of the wing's mid-section. For an F1 car, there is a rational case for ordering the priorities this way.

That said, what are the mechanical effects of having the lower arms shorter than the uppers?

First off, it's not good but it's not a disaster. Practically every strut suspension in the world approximates a short-and-long-arm system with a long upper arm and short lower.

We do get large changes of force line slope with suspension movement. (Force line is my term for the line from contact patch to instant center.) The roll center, when defined in the customary way as the front view force line intersection, can move around quite a bit.

When the force line intersection is near the ground, or when the force lines are close to horizontal, small changes in force line slope can produce dramatic changes in the lateral whereabouts, and even the height, of the force line intersection. This is true even when force line slope change rates are moderate. Some people suppose that this dramatic force line intersection movement translates to similarly dramatic changes in car behavior. However, this is not necessarily so. In real life, the location of the force line intersection can migrate wildly, with no adverse effects, provided that no abrupt changes in force line **slope** occur. This demonstrates that the practice of treating the force line intersection as a roll center or as a "moment center", about which we can take moments to predict the car's tendency to roll, is not scientifically sound.

Indeed, there are cases where there is no force line intersection (when the force lines are parallel to each other but not horizontal), and one case where there are an infinite number of force line intersections (force lines both horizontal).

Some may say this is heresy, but actually there is widespread recognition in the automotive engineering community that the whole notion of roll centers needs to be re-thought.

Full discussion of the issue is impossible in a document the size of this newsletter. My video "Minding Your Anti" ($50.00 from me, shipping included) explores it in depth. We can cover some basics here, however.

The two front or rear wheels in an independent suspension system each generate an individual upward or downward jacking force when lateral force acts at the contact patch. This force may act to oppose roll, in which case the individual suspension is said to have anti-roll. Or the force may act to exaggerate roll, in which case the individual suspension is said to have pro-roll, or negative anti-roll.

The magnitude of these jacking forces depends on two things: the slope of the individual wheel's force line, and the magnitude of the lateral force acting on the contact patch. This means that the amount of overall anti-roll or pro-roll moment from a wheel pair depends not only on the control arm geometry but also on the tire loadings, cambers, and slip angles. Change the anti-roll bar stiffness or the springs, and you not only change the elastic, or spring-derived, component of the roll resistance, you change the geometric component as well because you change the wheel loadings.

The effects we are discussing here are sometimes called "lateral anti" effects, because they are directly analogous to anti-dive, anti-squat, and anti-lift effects in side-view suspension geometry.

The location of the force line intersection doesn't matter at all! We can define some relationships between the force line slopes and their intersection, of course, but we cannot say, for example, that the car necessarily has a lot of geometric anti-roll if the force line intersection is high, or a lot of geometric pro-roll if the intersection is below ground. We also cannot necessarily say that the suspension has little anti-roll in a banked turn if the force line intersection is toward the inside of the turn, or a lot of anti-roll if the intersection is toward the outside of the turn.

Even the force line slopes don't matter as regards any tendency toward roll resulting from car-vertical forces.

There is a way to assign a roll center height for modeling and discussion purposes that takes account of these realities. The video explains how to do that.

Anyway, returning to the question of what happens when we have the lower arms shorter than the uppers, if the arms are all horizontal when roll is absent, as approximated by the Williams, then in the unrolled condition, neither the inside nor the outside wheel has any anti-roll or pro-roll. As the car rolls, the outside wheel acquires rapidly-increasing pro-roll (downward jacking), and the inside wheel acquires rapidly increasing anti-roll (also downward jacking). Since the outer wheel generates a progressively larger share of the lateral force as we corner harder, we can confidently say that the net result is pro-roll for the wheel pair, increasing relatively rapidly with increasing roll. We may model this, or think of it, as a roll center that is at ground level to begin with, and drops relatively rapidly as the car rolls.

As for camber properties, in roll the outside wheel starts out with a camber recovery rate of zero (wheel tilts out of the turn an amount equal to roll angle) and gets worse from there – in other words, the camber recovery goes negative (wheel tilts out of the turn an amount greater than roll). On the inside wheel, camber recovery starts at zero and improves from there. In ride, camber goes toward positive in both bump and droop. Of course, any static or initial camber value will be superimposed on these effects. And again, the foregoing discussion assumes an initial condition where all arms are horizontal.

CLUTCH USE WHEN DOWNSHIFTING

In the July 2002 newsletter, we explored the question of whether to downshift while braking for a turn. I pointed out that we really can't do it any other time, since we need the car in the proper gear for the turn as soon as braking ends, and we generally can't downshift it before we get it slowed down.

A letter in the February 2004 *Racecar Engineering* notes that it is not common anymore to double-clutch, or double-declutch as they say in England, while heel-and-toeing to match revs while downshifting. The writer asks whether there is actually any penalty or benefit to letting the clutch pedal up while matching revs.

Normally, the reason for doing this is to match the speed of the meshing parts inside the gearbox, and thereby reduce wear in the box, especially to the dogs or synchros, as opposed to merely getting a smooth-feeling shift that won't upset the chassis or break the tires loose.

To some extent, it depends on the type of clutch and gearbox. Many multi-disc racing clutches drag a bit even when released, and if that's the case and the rotating parts are light, it may not matter much if the clutch pedal is up or down while the revs are being matched.

If the gear ratios are close, the mismatch in revs without double-clutching is less than if the ratios are widely spaced.

The situation where double-clutching is perhaps the most beneficial is with a wide-ratio transmission, or a skip-shift (two gears at once) with a close-ratio box, and a single-plate clutch that does not drag at all when released.

I think rev-matching is definitely beneficial with both dog and synchromesh boxes. I can't think of any reason it wouldn't be desirable in terms of gearbox life.

Another option popular with some drivers, particularly with dog boxes, is to not use the clutch at all. If done right, this can work quite well, but the consequences can be relatively severe if the driver misses the rev match. Clutchless shifting does have the advantage of allowing left foot braking. This permits a quicker transition from braking to power, at least in theory. It also offers an advantage in speed of response with a laggy turbo motor. The driver can feed the engine some throttle before getting completely off the brakes, and thereby help the turbo get spooled up.

It is also fairly common to upshift clutchless, and use the clutch, either once or double-clutching, for downshifts only.

All these techniques have their adherents. And any one of them, done right, beats any one of them, done wrong.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
July 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SETUP FOR LOWER GRIP

All the books talk about setting the car up in a general manner but do not mention anything about what to do with regard to the grip level at different circuits. I run in a classic cars series where the aerodynamic downforce is very small and we have to use the same type of tires for all of the season. I've set the car up at a track where the grip level is considered to be high. The car goes well there and the times are good, but when I've been to slicker tracks I don't know what to do. Should I make the car softer?

In general, a lower-grip surface does call for softer settings.

However, there can be exceptions. If you are restricted to a relatively hard tire, and if the tire needs to heat up to work properly, a stiffer setup may help get the tire up to temperature. Given a free hand, you'd go softer on both the compound and the suspension.

Many street tires, and some tires used in vintage racing, do not get stickier as they get hotter, in which case you would generally go softer with the suspension.

Another factor that enters into this is how bumpy the surfaces are. Bumpier surfaces require softer suspension. So if the slicker tracks are also bumpier, the decision is pretty easy. We then have two factors that would tend to demand softer suspension settings. On the other hand, if you have a situation where the grippy tracks have bumps and the slick ones are smooth, things are not so clear-cut.

You don't mention what kind of car you have, or what class you're running. It would be possible, particularly in vintage racing, to have a car that is favored by a grippy track merely because it has a power advantage against the cars it's classed with, and also has a handling disadvantage. In such a case, you may just have to live with the situation.

If you have not tried softer suspension at the slicker tracks, and you are less competitive there, I would say trying softer settings would make sense. Keep good records, and you can always go back to your old settings if the softer ones don't help.

LOWERING BLOCKS ON TRUCK ARMS

What is the effect of using lowering blocks on a truck arm suspension?

First of all, there is no effect to the anti-squat or anti-lift properties, assuming that the pivots at the front of the arms are not moved, and assuming the ride height is adjusted back to its previous setting.

In most big-spring cars, the springs act on the truck arms, near the axle. In such cases, the car will be lowered by the lowering blocks. If there are jacking bolts at the springs, the ride height can be reset to the previous position, assuming sufficient available adjustment.

On most stock car truck arm suspensions, the shocks and the Panhard bar attach to the truck arms. It would be mechanically possible to build a truck arm suspension without this characteristic, however.

In some classes, coilovers are allowed on truck arms. Usually, the coilovers attach where the shocks would attach on a big-spring car. But again, they wouldn't have to.

If you have coilovers attached to the truck arms, the same thing happens to the ride height as if you had big springs on the truck arms. If you have coilovers that attach to brackets on the axle tubes, ride height adjustment won't be affected.

With coilovers or big springs, if the shocks attach to the truck arms, lowering blocks will cause the shocks to be extended more at a given ride height. This means the shocks have more compression travel and less extension travel available. This may hurt, or help, or make no difference. If you use bump rubbers on the shocks to keep other components from hitting, you may need bigger bump rubbers.

If the Panhard bar attaches to the truck arms, the end of the bar that attaches there will be lowered. On racing truck arm suspensions, you usually have at least a couple of inches of adjustment on this attachment. As with ride height, if sufficient adjustment is present, the Panhard bar can be returned to its original setting, but the overall range of possible settings moves lower.

Assuming the lowering blocks have no taper, they will change pinion angle slightly, upward. This effect will be small, and should not affect the car's behavior.

So is there any advantage to using lowering blocks? If you want a lower rear roll center or more bump travel from the shocks, maybe. Otherwise, the blocks are added weight, with no benefit.

OHLINS HIGH-FREQUENCY PISTON

Are you familiar with the Ohlins high-frequency shock piston? The claim made by Ohlins, and by those who use them, is that they offer more grip.

Great description of the pistons' benefits, huh?

Ohlins is not the kind of company to build such an odd piston without a reason, but I have yet to hear a scientifically sound explanation. Any input?

Last year I attended a seminar presented by Ohlins, where they showed us these pistons. (If you get a chance to attend an Ohlins seminar, do it – they're good.)

This piston provides what might be called relative position (or displacement) sensitivity. This would contrast with absolute position sensitivity, meaning sensitivity to position of the piston in the shock body. This relative position sensitivity is sensitivity to the piston's position **relative to last velocity reversal**. The shock is softer for a few millimeters of motion after any reversal of piston motion.

This effect is more or less the opposite of "stiction", the phenomenon where the friction is higher in the first little bit of motion.

The shock therefore has a softer action on low-amplitude disturbances. It might be a bit more descriptive to call it a low-amplitude piston, but high-frequency disturbances tend to also be low-amplitude. Or maybe it should be called a negative-stiction piston. That would cause some head-scratching.

Here's how it works: under, or inside the teflon piston sealing ring is an o-ring. Unlike many designs that use the o-ring to load the sealing ring against the wall, in this design the o-ring is a loose fit in its gland. The o-ring can float up and down, and will also let fluid pass around it.

The lands that form the top and bottom surfaces of the o-ring gland have holes drilled through them, rather like vertical gas ports in a drag racing engine piston. When the o-ring seats against the top or bottom land, it closes off these ports. When the ports are open, they add bleed to the shock, and soften its action.

The o-ring is moved partly by its own inertia, but mainly by oil flow. When the piston comes to rest after moving in one direction for a while, the o-ring will be seated and the ports will be closed. As soon as the piston starts moving the other way, oil flow unseats the o-ring and starts driving it toward the other land. When it gets to the other land, it seats against the holes there. When the o-ring is floating from one land to the other, the ports are open and the shock has increased bleed.

On a series of small-amplitude disturbances, the o-ring might be floating between lands most of the time. The piston's unique properties would have the greatest possible effect under these conditions.

Most race cars have considerable stiction in the suspension due to the large number of sliding-contact pivots. One reason passenger cars use rubber bushings at the pivots is that they don't have stiction. So this piston to some extent compensates for the stiction not only in the shock itself but in the rest of the suspension.

WHICH SHOCK TO SOFTEN?

I have a quick question about my DIRT modified. To get off the corner harder, would you use an easy-up (soft rebound) shock on the right front, the left front, or both?

Quick answer would be right front.

The basic rule here is that anything that adds an extension force or reduces a compression force (extension damping force acts in the compression direction) adds load to the near tire and the diagonally opposite one, and unloads the other two. If that adds diagonal percentage, we tighten the car. If it reduces diagonal percentage, we loosen the car.

We are making a few assumptions here, though. We are assuming that you're trying to tighten the car. If you've already got a power push, making the car tighter won't help you come off harder. Instead it will make you be gentler with the throttle to keep the front end stuck.

We are also assuming that the surface is smooth enough so that sprung mass motion is the main cause of suspension motion. This may not be the case at all, especially on dirt.

The effectiveness of using shock valving this way will also depend how much extension velocity the right front has.

My own preference is to try to work with the springs instead, and valve the shocks for optimum roadholding rather than trying to trick the shocks into controlling wheel load distribution.

WHY THE CHAINS, SIR?

When NASCAR teams use a chain for one of the sway bar links, are they using it as a lost motion device, allowing wheel travel before the bar rate becomes active?

More common than a chain nowadays is an adjustable pad on the end of the sway bar, bearing on a pad on the lower control arm. Chains are still seen sometimes in the lower divisions, where original-equipment-style bars are required. But the basic idea is the same either way: have a connection that transmits force in only one direction. The bar only resists rightward roll, unless it's preloaded, in which case it does resist leftward roll up to the point where it unloads.

The intent here is to help keep the car from going quite so loose when the driver gets the left front wheel on the apron of the track, which is sometimes abruptly flatter than the banked turn.

Usually, the bar is run snug or slightly preloaded at static condition. That means that the bar acts just like it normally would in a left turn. When the car is cornering, the bar has substantial load on it. The one-way connection, be it a pad or a chain, will only go slack if the left front wheel hits the apron hard enough to put the front suspension into a left roll condition – left front deflection greater than right front. This leads me to question the use of these devices, especially since they make the car loose when turning or spinning to the right, which can happen during a crash or when avoiding one. Nevertheless, they are very popular.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
August 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL CENTER MIGRATION, SOME MORE

The discussion of roll center definition in the June newsletter prompts this question from a reader:

Would you be able to discuss the effects of front lateral roll center migration for an oval track car with a solid axle rear end (NASCAR style) – perhaps an example on a short track where there are low speeds and aerodynamic loads, and higher amounts of vehicle roll where the left side suspension could be travelling into rebound?

The questioner mentions that he is an engineer for a major car manufacturer, and expresses a desire to remain anonymous, which is my usual practice in any case. Knowing that the questioner is an engineer, I am going to assume here that the reader is conversant with the basics of roll center theory as usually understood, and will not start at square one.

As those who read the June newsletter will know, I do not believe that the intersection of the front view projected force lines can properly be considered the roll center, or moment center, or anything of the kind. There are situations where you don't get big modeling errors if you use the force line intersection as a roll or moment center, and other cases where you get huge errors. This merely illustrates that an incorrect analysis method can coincidentally produce correct or nearly correct answers in certain cases, despite the incorrectness of the method itself.

I have also said that the roll center, properly assigned, should be considered a point in side view (of the car), and its lateral position should be considered undefined. It lies in the transverse plane containing the wheel center in all cases, or, in side view, it lies straight down or straight up from the wheel center. So we really need only one number to define its position, namely its height. This height is not the same as the height of the force line intersection. Rather, it is the mean height of the two force line intercepts on a line I call the **resolution line**.

The resolution line is a vertical line in the front view, positioned according to distribution of lateral force generated by the two tires. For example, if the right front tire is generating 75% of the front lateral force, the front suspension resolution line is 75% of the track width away from that tire.

Unfortunately, we do not know this distribution of lateral force exactly, in most cases. We have to estimate it. That means our modeling of the suspension's behavior is only as good as this estimate. This is unfortunate, but ignoring the fact doesn't make it go away. The behavior of the suspension really does depend on the distribution of lateral force. To predict the jacking forces each of the individual wheels generates, and thereby calculate an anti-roll or pro-roll moment, we must not only know the suspensions' geometry, but also the forces at the contact patches. Any analysis method that takes this into account, even using an estimate for the lateral force distribution, is better than a method that ignores this factor altogether.

What we're doing here is directly analogous to modeling longitudinal anti effects – anti-dive, anti-squat, anti-lift – in side view. It is widely recognized that when modeling longitudinal anti effects, we have to know the front/rear distribution of longitudinal force, or try to estimate it with reasonable accuracy. For example, for braking short of lockup, we use the calculated brake bias. If the front brakes make 70% of the rearward force, we construct our resolution line 70% of the wheelbase back from the front wheel center. We then look at where the front wheel side view force line intercepts this resolution line. We take the height of this intercept as a percentage of sprung mass center of gravity height, and that is our percent anti-dive. We can do the same for the rear wheel, and that's our percent anti-lift. When these are both 100%, the car will not pitch at all in braking, regardless of wheel rates, nor will the whole car jack up or down.

We can likewise define a percent anti-roll for the right and left wheels in an independently suspended front or rear pair, and we may also average these to define a percent anti-roll for the wheel pair. The average height of the intercepts makes a good value to use for roll center height – much better than using the force line intersection height, though in some cases the two values may be similar. The average height of the intercepts, or roll center height, may also be described as the sprung mass c.g. height times the percent anti-roll for the wheel pair. Also, the height of each of the intercepts, as a percentage of sprung mass c.g. height, is that wheel's percent anti-roll.

This definition of the roll center provides a number that can be accurately used for load transfer and roll angle calculations, for it is a valid measure of the suspension's geometric anti-roll properties.

Assigning a roll center location is useful not only for modeling or analysis, but also for discussion. To use the method I'm advocating for discussion, it is useful to have default assumption for lateral force distribution. I think assuming that the outside wheel generates 75% of the force is appropriate, absent better information.

Note that suspensions only generate geometric anti-roll or pro-roll moments in response to car-horizontal (lateral or longitudinal) forces. Force line slopes, force line intersections, and force line intercepts of the resolution line do not affect any tendency to roll, or resist roll, in response to car-vertical forces. (Spring splits, or wheel rate splits, do affect this. So does an offset c.g., or static left percentage other than 50%.)

Of course, the questioner here is not asking about effects of changes to the location of the roll center as I prefer to define it. He is asking about effects of lateral migration of the roll center as most people conceive it, namely the intersection of the front view force lines.

And in fact we can say some things about the position of the force line intersection, and the conclusions we can draw from it, for particular classes of situations. We can't necessarily say the car has more tendency to roll, other things held constant, if the intersection moves to the inside of the turn, nor that the car has less tendency to roll if the intersection moves in the direction of roll, even in a banked turn. However, we can make some more complex and qualified statements, for particular sets of conditions or assumptions.

To begin with, there are certain situations where we don't know much at all from the force line intersection. If the force line intersection is at the contact patch center for either of the wheels, we know that the opposite wheel's force line is horizontal, and therefore the opposite wheel has no anti-roll or pro-roll. However, the force line for the wheel on top of the intersection could be at any angle, and therefore this suspension could have any amount of anti-roll or pro-roll. In this situation we can't say anything about the overall amount of anti-roll or pro-roll in the geometry from the location of the force line intersection, nor can we infer the location of the roll center as I define it, without additional information.

Parallel lines do not intersect. When the force lines are parallel, there is no force line intersection. In this situation, users of force line intersection as the roll center will either say the roll center is undefined, or that it has disappeared, or – arbitrarily – that it is on the vehicle centerline, at the average height of the two force line intercepts of the centerline, which will be at ground level. However, the parallel force lines could be at any angle, relative to car-horizontal. We don't know what that angle is; all we know is that it's the same for both of them. We know that one wheel has anti-roll and the other has pro-roll, but we don't know how much. We know that the anti-roll and pro-roll forces are equal if the tires are making equal car-lateral force, which would equate to a roll center at ground level. But if the tire forces are unequal – and they usually will be – we cannot say how much overall anti-roll or pro-roll the system has, and we cannot define the roll center my way, without additional information.

There is a third unique class of situation – or, if you like, a special case of the parallel force line situation – the one where both force lines are horizontal. In other words, the force lines are not only parallel, but coincidental. In this case, we cannot say what the lateral location of the force line intersection is. We may say there is an infinitely large number of them. We do know, however, that all these points are at ground level, and we can say with certainty that the suspension has no anti-roll or pro-roll, regardless of the magnitude of car-lateral forces at the contact patches. We can also say that the resolution line intercepts of both force lines are at ground level, no matter where the resolution line lies. Therefore we can define a roll center height my way, at ground level, despite the fact that we cannot define a single force line intersection. In this case only, we can do this without knowing, estimating, or assuming lateral force distribution.

For all cases except the above three classes, we can calculate the slopes of the force lines from their intersection. Knowing this, and a known, estimated, or assumed lateral force distribution, we do know enough to assign a roll center my way, and to say something about the overall anti-roll or pro-roll characteristics.

We can also say which direction the force line intersection will move, for a known roll displacement, if we know one more characteristic of the suspension: how the force line slope changes with suspension movement.

William C. Mitchell, in his SAE paper no. 983085 entitled *Asymmetric Roll Centers*, introduces a definable parameter that is useful in discussing this. He calls it the **incline ratio**. I find this nomenclature to be suggestive of a different meaning, and I think Bill deserves recognition for coming up with the idea, so I call it the **Mitchell index**.

By either name, we calculate this number as follows: We look at the centerline intercept of the force line, and we note its rate of height change as the suspension moves in ride. We express the rise and fall of the intercept as a proportion of the ride motion, and that's the incline ratio, or Mitchell index. If the intercept moves up and down at the same rate as the sprung mass, we have a Mitchell index of 1. If it doesn't move at all, we have a Mitchell index of zero. If it moves up when the sprung mass moves down, we have a negative Mitchell index. If it moves down when the sprung mass moves down, by a lesser amount, we have a Mitchell index between 1 and 0. If it moves down when the sprung mass moves down, by a greater amount, we have a Mitchell index greater than 1.

The case the questioner raised in the June newsletter, where a short-and-long-arm suspension has the lower arms shorter than the uppers, illustrates a Mitchell index substantially greater than 1. Likewise, a strut suspension has a Mitchell index greater than 1. A pure trailing arm suspension has a Mitchell index of 0. Most short-and-long-arm layouts have Mitchell indices fairly close to 1 or a bit greater. With unusually short upper arms, stock car front ends can have a Mitchell index a bit less than 1. To get a Mitchell index of zero with a short-and-long-arm suspension, we need either very long lower arms, or very short uppers. The lengths have to be in accordance with Olley's Rule: the lengths of the control arms have to be inversely proportional to their height above ground level, usually as measured at the ball joints. For typical stock car lower arm and spindle (upright) dimensions, that means upper arms somewhere around six to seven inches long, rather than the lengths of 9 inches or more commonly seen. Not surprisingly, Mitchell indices less than zero are uncommon.

The Mitchell index can be different for the right and left wheels, and in oval-track stock cars it usually is, though not by much. It also varies some as the suspension moves, but it does not undergo large, sudden changes. We also end up defining it differently if we take the centerline as being where the frame builder marked it, or as being at the midpoint of the front track, or as being the edge view of the longitudinal c.g. plane. These nuances aside, if we consider roll to be angular motion about the ground intercept of whatever centerline we've defined, then we can say certain things about how the force lines and their intersection will move in particular combinations of ride and roll, based on the Mitchell indices of the two individual wheel suspensions.

If the Mitchell index is 1, the force line slope doesn't change in roll. If the Mitchell index is 1 for both right and left wheels, the force line intersection doesn't move in roll.

To take the most common category of cases first, suppose that, at static condition, the force line intersection is above ground level and between the wheels. In this condition, if the Mitchell index is greater than 1, the force line intersection always moves laterally opposite to the direction of roll. The force line for the outside wheel (right wheel in a left turn) loses inclination, while the force line for the inside wheel gains inclination. In the case the questioner cites, the intersection would move to the right. If the Mitchell index is less than 1, the force line intersection moves toward the outside wheel instead. The outside wheel force line gains inclination, while the inside wheel force line loses inclination.

In general, the former case implies a decrease in overall geometric anti-roll, and the latter implies an increase in overall geometric anti-roll, even with no change in force line intersection height, because the outside wheel generates more lateral force. Correspondingly, the roll center, defined my way, drops in the former case and rises in the latter case, even with no change in force line intersection height.

Now let's change things a little. Let's suppose the force line intersection is between the wheels but below ground level. This is actually not an uncommon condition in stock cars, especially drop-snout cars on banked turns.

Now, if the Mitchell index is greater than 1, the force line intersection moves toward the **outside** wheel in roll! The outside wheel is still losing anti-roll, or should we say gaining pro-roll. The inside wheel is still gaining anti-roll, or losing pro-roll. So the change in roll resistance is still the same as when the force line intersection was above ground, but the lateral migration of the force line intersection is in the opposite direction – toward the outside wheel.

If the Mitchell index is less than 1, again the change in roll resistance is the same as with an above-ground intersection – it increases. And again, the lateral migration of the intersection is in the opposite direction – toward the inside wheel.

This illustrates that we cannot infer the change in roll resistance knowing only the direction of lateral migration of the force line intersection, even supposing that the intersection height isn't changing.

The wildest migrations of the force line intersection occur when the force lines are close to horizontal, and close to parallel. Small changes in force line angle will make the intersection move all over the place. Small ride motions can make it move from above ground and way out to the right to below ground way off to the left. Does this mean the geometric anti-roll or pro-roll moment is varying all over the place, or that the car's properties in a banked turn are varying all over the place? Not at all, because the force line slopes and individual wheel anti-roll and pro-roll are not changing much. And, correspondingly, the height of the roll center as I define it doesn't change much.

All of this holds true regardless of whether the turn is banked, and regardless of what kind of suspension is at the other end of the car.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
September – November 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ERRATA AND ADDENDA

Thank you to readers who have pointed out some detail discrepancies in my recent columns in *Racecar Engineering*, which are drawn from this newsletter.

In the April newsletter I originally stated that for least bearing loads with an outboard brake, the caliper should be somewhere in the upper rear quadrant of the disc. Shortly after sending out that issue, I sent out a correction saying that should be the lower rear quadrant. This correction was supposed to be incorporated when the material was published in the magazine, but unfortunately the original version was what ran (July 2004 issue). So a correction is in order on this point, for those who read me in the magazine.

In my October 2004 column, drawn from the June 2004 newsletter, there was some disagreement between what I said about the Williams Formula 1 suspension and the second picture that ran in the magazine, showing the suspension from above and behind. I had said the forward portion of the upper wishbone was not aligned with the tie rod, whereas in the picture it appears to be aligned. Some readers have understandably called this to my attention.

The picture in question was chosen by the magazine, not supplied by me. I based my comments on other pictures, which I did not have in electronic format. If anybody is to be faulted here, I am. Anyway, it appears there have been two versions of the suspension. In the version shown in the magazine, the bump steer and aerodynamics appear good, but the camber control properties appear poor. In the version I was looking at when I wrote the text, the control arm appears to have been leveled out by moving the forward pickup point down a bit, without the steering rack being lowered to match.

Without inside knowledge of the team's internal affairs, I am of course speculating as best I can from partial evidence. The best explanatory theory I can devise is this: the original version had the wishbone and the tie rod as shown in the October column's illustration. Perhaps the camber control was consciously compromised to get the nose higher and aid airflow underneath it, which was also partially the object of the tusk nose design. To get the floor of the nose up, the driver's feet had to go up. That forced the steering rack, the pedals, and everything else to go up. The wishbone then had to agree with the steering, to prevent bump steer and get good airflow over the tie rod and wishbone.

When it was found that the car lacked front grip, and other fixes didn't cure the problem, the team tried moving the wishbone pickups down. Moving them down a little on the existing tub was feasible, but moving the steering rack was tougher. So as a temporary experiment and hopefully a temporary solution, the team decided to accept some bump steer and rely on driver skill to deal with that, and see if improved camber control would help the grip.

Raising the upper ball joints and outer tie rod ends was not an option, because the wheel rim was in the way. Lowering the rack was not an option, because the driver's feet (or maybe other elements of the car) were in the way. So the team made the only modification they could, under the circumstances.

I also noticed elsewhere in the October issue that as of the Hungarian Grand Prix the team had abandoned the tusk nose design, although the tub has not been re-done.

I understand they are going to try an entirely new approach for 2005.

TIRE LOAD SENSITIVITY AND WEIGHT TRANSFER IN TRAILBRAKING

In the Forum (letters) column of that same October issue was a letter from a physicist, responding to an earlier article on the theory of cornering line and trailbraking by Erik Zapletal. Erik had correctly noted that forward load transfer (weight transfer, in customary vernacular) tends to improve the lateral acceleration capability of the front wheels, at the expense of the rear ones. The physicist took issue with this, and pointed out that adding weight to a wheel pair reduces lateral acceleration capability, because due to the phenomenon we call tire load sensitivity, the coefficient of friction diminishes as we add load. Mr. Zapletal replied that this is true, but load sensitivity is a minor effect. (He also noted that other factors enter into this, including aerodynamics, brake bias, and camber changes.)

Who is correct? Both are, partly. But I think I can explain the matter a little better.

By the way, Erik is a sharp guy and it was he who first pointed out my aforementioned error on brake caliper location. Hopefully, I am returning a favor and shedding some light here, not being a pain.

This question illustrates perfectly why I prefer to speak of load transfer rather than weight transfer. The effects of forward load transfer under braking are quite distinct from the effect we get if we move mass forward in the car. Moving mass forward in the car adds understeer. Forward load transfer in braking adds oversteer. Both effects can be said to relate to tire load sensitivity.

When we move mass forward in the car, we increase both the normal (vertical or perpendicular to the road) force on the front tires and the centrifugal (inertial) force the tires must overcome to produce any given lateral acceleration. Consequently, the front end's lateral acceleration capability depends not on whether the force capability of the tires increases, but on whether it increases at the same rate as the normal (and centrifugal) force. The ratio of the tire's force capability to the normal force is the coefficient of friction. This diminishes with increasing normal force, so in any situation where we add weight to the front end, the lateral acceleration capability diminishes.

When we brake, however, most of the load increase on the front wheels does not come from mass moving forward on the wheelbase, although a small amount of such motion does usually occur. The increase in front wheel loading comes primarily from the forward pitch couple which inevitably results from the tires exerting a rearward force at ground level and the car's inertia exerting an equal and opposite forward force above ground level. To prevent the car from somersaulting, the front tires exert an increased support force against the ground, and there is a corresponding decrease in support force at the rear, creating an equal and opposite anti-pitch couple. Because the center of mass has not moved appreciably on the wheelbase, the front wheels are not required to overcome an increased centrifugal force per unit of centripetal acceleration in proportion to their increased normal force. The normal force increases, while centrifugal inertia force for a given car-lateral (centripetal) acceleration remains largely unchanged. Consequently, lateral acceleration capability for the front wheel pair increases.

In both cases, the normal force increases and the coefficient of friction decreases. But in the former case, the centrifugal force per unit of acceleration increases with the normal force, whereas in the latter case it does not.

The latter case may be said to be similar to what happens when we add aerodynamic downforce. We add significant normal force, or load, without adding significant mass.

WHY ARE WIDE TIRES BETTER?

It has been recognized for about 40 years now that wide tires provide more grip, at least when we are not limited by aquaplaning. One might suppose that this effect would be be well understood by now, on a theoretical level as well as a practical one. Yet the matter seems to be receiving a lot of attention from various authors lately. This seems to be due in part to the need for mathematical tire models to be used in computer simulation. I have encountered the question at least twice in the past month, once in a seminar presented by Paul Haney, based on his recent book about tires, and once in Paul Van Valkenburgh's November *Racecar Engineering* column. The issue has also come up in my work as an advisor to the UNC Charlotte Formula SAE team.

On the face of it, one might wonder why there is any controversy about this, and also why it took people until the 1960's to try wide tires. More tire, more rubber on the road. More rubber on the road, more traction – right? Why wouldn't this be obvious?

Essentially, there are two reasons it wasn't obvious. First, according to Coulomb's law for dry sliding friction, friction is independent of apparent contact area. It depends instead on the nature of the substances in contact, the normal (perpendicular) force, and nothing else. Second, a tire's contact patch area theoretically doesn't vary with its width anyway. If we widen the tread, the contact patch just gets shorter, and the area theoretically stays the same.

Let's consider each of these notions. Coulomb's law applies quite accurately to hard, dry, clean, smooth surfaces. However, a tire tread is a soft, tough, sometimes tacky substance in contact with a hard, rough surface. When two hard, smooth surfaces are in contact, they actually touch only at a small percentage of their apparent or macroscopic contact area. Friction depends on molecular bonding in the small microscopic contact zones. As normal force increases, the microscopic contact area increases approximately proportionally, and consequently friction is directly proportional to normal force.

With rubber on pavement, however, there is not only the usual molecular bonding but also mechanical interlock between the asperities (high points) of the pavement and the compliant rubber. Sliding then involves a combination of shearing the rubber apart and dragging the asperities through it as the rubber reluctantly oozes around the asperities. The interface somewhat resembles a pair of meshing gears. With gears, when we increase the size and number of teeth in mesh, we increase the force required to shear off the teeth. It would be reasonable to expect a similar effect with the interlock between the tread and the pavement.

With increasing normal force, this interlock gets deeper, as the asperities are pushed further into the rubber. However, we might reasonably expect that at least beyond a certain point, the asperities are pushed into the rubber to pretty nearly their full depth, and further increase in normal force does not proportionately increase the mechanical interlock. With greater macroscopic contact area, it should take a greater normal force to reach this region of diminishing return.

A tire typically does show characteristics that would match this hypothesis. It will often have a range of loadings where its coefficient of friction is almost constant; where friction force is almost directly proportional to normal force. Above this range, the tire exhibits much greater load sensitivity of the coefficient of friction. The curve of friction force as a function of normal force goes up almost as a straight line for a ways, then begins to droop at an increasing rate.

Of course, the contact patch does not remain the same macroscopic size as load increases. It grows as we add load. Nevertheless, this contact patch growth is evidently not enough to keep the coefficient of friction constant.

The contact patch growth is interesting in itself, and a bit counter-intuitive. A tire can be considered a flexible bladder, inflated to some known pressure, and supporting a load. If such a bladder is extremely limp when uninflated, like a toy balloon, and we inflate it, place it on a smooth, flat surface, and press down on it with a known force, the area of contact with the surface is equal to the normal force divided by the pressure: $A = F_n/P$.

If a tire approximates this behavior, then it follows that the contact patch area depends only on the load or normal force and the inflation pressure. If we make the tire wider, then at any given load and pressure the contact patch doesn’t get bigger, it just gets wider and shorter.

Accordingly, much discussion of the reasons a wide tire gives an advantage focuses on reasons we might expect a wider tire to yield greater lateral force than a narrower one, assuming similar construction and identical pressure, tread compound, and load.

One theory, advanced by the late Chuck Hallum and evidently picked up by Paul Van Valkenburgh in his recent column, is that a tire is primarily limited by thermodynamics. It generates drag when running at a slip angle. The drag times the speed equals a power consumption, or rate of energy flow. This energy is converted into heat. For the system to be in equilibrium, the heat must be dissipated as fast as it is generated. Even short of the point of true equilibrium, the tread compound needs to be kept below a temperature where it softens to the point of being greasy rather than tacky. If the contact patch is shorter, that means that each square inch of tread surface spends less time getting heated and more time getting cooled.

Also, when a tire is operating near its lateral force limit, the front portion of the contact patch is “stuck” to the road and the rear portion is a “slip zone” in which the tread moves across the pavement in a series of slip-and-grip cycles. The slip zone grows as we approach the point of breakaway. Beyond the point of breakaway, the entire contact patch is slip zone. The slip zone generates less force and more heat than the adhering zone. A shorter, wider contact patch is thought to have a larger adhering zone and a smaller slip zone at a given slip angle, and wider tires are also known to reach peak force at smaller slip angles. Therefore, a wider tire is not only better able to manage heat, but also generates less heat at a given lateral force.

This all makes sense, but it fails to explain why wide tires give more grip even when stone cold.

There is little doubt that they do. If you have a street car with four identical tires, and you replace the rear tires and wheels with ones an inch wider, using the same make and model of tire, with no other changes, the handling balance will shift markedly toward understeer. You will see this effect at all times, from the first turn in a journey to the last. Surely this effect is not coming from heat management.

Paul Haney explains this by the larger-adhering-zone theory described above. The tire makes more efficient use of its contact patch, even if the contact patch isn’t larger.

As much sense as the above theories make, they ignore some real-world effects that have a bearing on the situation.

First of all, the degree to which tires follow the $A = F_n/P$ rule varies considerably. A very flexible tire, at moderate load, may have a contact patch as large as 97% of theoretical. A fairly stiff tire may be well below 80%. We are all aware of run-flat tires currently being sold, which will hold up a Corvette with no inflation pressure at all. As P approaches zero, F_n/P approaches infinity. If A does not approach infinity, and the tire does not go flat, the contact patch area as a percentage of theoretically predicted area approaches zero.

One might suppose that the effect of carcass stiffness would be significant mainly in street tires, with run-flats being an unrepresentative extreme. Yet I have seen dramatic differences in carcass rigidity in different makes of racing tires intended for the same application. The Formula SAE car run by the University of North Carolina Charlotte uses 10" wheels. Hoosier and Goodyear both make 6" nominal-width tires for the application. The stiffnesses of these tires differ dramatically. The Hoosiers are much more flexible than the Goodyears. The Goodyears are so stiff that they will support the front of the car (without driver), with little visible deflection, when completely deflated – run-flat racing tires! How closely do these tires approximate $A = F_n/P$ in this load range? Not very closely at all.

My point here is that tire stiffness, vertically, laterally, and otherwise, is not purely a function of inflation pressure, so it is a bit risky to try to infer contact patch size from pressure and load. Therefore, we don't necessarily know that two tires differing only in width do have the same contact patch area at the same inflation pressure and load, or even that tires of the same size do.

Anyway, if it is approximately true that $A = F_n/P$, it follows that a wide tire will have greater vertical stiffness, or tire spring rate, than a narrow one, at any given inflation pressure. It will also have a smaller static deflection at a given load, which is why the contact patch is shorter. The flip side of this is that for a given static deflection or tire spring rate, a wide tire needs a lower inflation pressure. Consequently, if we compare wide and narrow tires at similar static deflection or tire spring rate, rather than similar pressure, they will have similar-length contact patches and the wider one really will have more rubber on the road, just as we would intuitively suppose from looking at them.

As we make a tire wider, not only does vertical stiffness increase for a given inflation pressure, so does the tension in the carcass due to inflation pressure. A tire is a form of pressure vessel. We may think of it as a roughly cylindrical tank, bent into a circle to form a donut or torus. Borrowing from the terminology of pressure vessel design, we may speak of the "hoop stress" in the walls: the tensile stress analogous to the load on a barrel hoop. For a given inflation pressure, the hoop stress is directly proportional to the cross-sectional circumference, or mean cross-sectional diameter. When the carcass is under a higher preload, the tire acts stiffer laterally. This effect can easily be seen in bicycle tires. A fat bicycle tire will feel harder to the thumb than a skinny one, at any given pressure. If we try to inflate a mountain bike tire to the pressure we'd use in a narrow road racing tire, the tire will expand its bead off the rim and blow out. So when we compare narrow and wide tires at equal inflation pressures, the wider one will be stiffer laterally as well as vertically, and it will achieve this at no penalty in contact patch size.

Finally, there is the question of tread wear. As we have noted, if the contact patch is longer, it has a larger slipping zone near the limit of adhesion, and it also spends a greater portion of each revolution in contact with the road. Not only do these factors influence how hot the tire runs, they also influence how fast it wears. Therefore, assuming good camber control, a wide tire should last longer than a narrow one, with similar tread compound. The astute reader will see where I'm headed with this. If we need to run a given number of laps or miles on a set of tires, then with wider tires we can trade away some of the inherent longevity advantage, and run a softer compound.

Okay, summing up, what does a wider tire get us?

1. It runs cooler, and/or
2. it makes more efficient use of its contact patch by having a greater percentage adhering, and/or
3. it can run at lower inflation pressure and therefore actually have a larger contact patch, and/or
4. it can have greater lateral stiffness at a given pressure and therefore keep its tread planted better, and/or
5. it can use a softer, stickier, faster-wearing compound without penalty in longevity.

Note that most of these effects in turn play off against each other. We can blend and balance them, and get a tire that is somewhat cooler-running, has a somewhat lower operating pressure and somewhat larger contact patch, has somewhat greater lateral stiffness, and survives long enough with a somewhat stickier compound, all at the same time. That would explain an improvement in grip, wouldn't it?

REAR PERCENTAGE VERSUS YAW INERTIA

Setting up race cars is invariably a compromise. Most of these are well documented, but I have found little on the compromise between longitudinal weight distribution and polar moment of inertia.

I race a 250bhp, 900K V8 MGB which, unfortunately for a rear-wheel-drive car, has a frontal weight bias of around 60%. The fuel tank of around 40L overhangs the rear axle, which creates a moment around the rear axle and thus helps to remove weight from the front.

There is space in front of the rear axle to place two tanks where the batteries would have gone on the road car. This would have the effect of reducing the polar moment of inertia while slightly lowering the center of gravity and possibly allowing softer rear springing. It would also improve safety (as long as it is protected from the prop shaft running between the two tanks!). However, it would increase the front bias.

This would be quite a lot of work and expense, so I would be grateful if you could comment on the benefits or otherwise of this approach.

From a vehicle dynamics standpoint, I would opt for more rear percentage rather than less yaw inertia, especially in a car that is so nose-heavy now.

Within limits, yaw inertia can be coped with by driving technique. In some situations, it can even make the car faster. Overall, though, less yaw inertia is better, particularly when the course demands high yaw accelerations, as when negotiating chicanes or street-circuit turns that come in quick succession. But steady-state handling and ability to put power down are more important. A rear-drive car with only 40% rear weight and a powerful engine is seriously traction-challenged, especially when exiting turns. Any further reduction would not be good.

With more front percentage, you will actually have to stiffen the rear suspension, at least in roll, with respect to the front. Otherwise, you will be adding understeer. The inside rear will then be extremely light when cornering. It probably is now.

One situation where yaw inertia can make a car faster is where the car may unexpectedly encounter a slippery spot in the middle of a turn. If the slippery patch is short enough so the front and rear of the car hit it separately, the car will experience understeer and then oversteer in very quick succession: it will do a wiggle. If it has little yaw inertia, it will do a big wiggle. If it is close to the limit, it may spin. If the driver wants to allow a margin of safety to increase the chances of catching the wiggle before it becomes a spin, for a given level of risk the driver must stay further from the limit in a car with little yaw inertia.

For this reason, in the days of high-speed open-road racing, many engineers regarded yaw inertia as desirable, and this was thought to be one of the advantages of a front-engined car with a transaxle, and a problem for the rear-mid-engine layout.

Even back then, everybody recognized that weight distribution change with fuel burn-off was not good. The Lancia-Ferrari of the mid-1950's, with its pontoon fuel tanks between the front and rear wheel on each side, appears designed to get much of the fuel amidships longitudinally, while preserving high yaw inertia. As Ferrari developed the car after taking it over from Lancia, they moved the side tanks inside the body, reducing yaw inertia, and this is generally thought to have improved the handling.

It is worth noting that the Lancia-Ferrari predated foam-baffled fuel containers. Even with some sheet metal baffles, there must have been considerable fuel slosh in those tanks, and that can't have been good for controllability.

As regards crash safety and fire risk with the fuel inside the wheelbase versus outside, there are pros and cons both ways. If the fuel is within the wheelbase, it is less likely to spill when the rear takes a hit. On the other hand, if it does spill, it is more likely to spill into the driver's compartment. And it can still spill, as recently demonstrated in Dale Earnhardt Jr.'s crash in the Corvette at Sebring.

At the recent SAE Motorsports Conference in Dearborn, Michigan, I had the opportunity to ask a very distinguished panel of safety experts about the safety aspects of fuel location. Gary Nelson of NASCAR said that they strongly considered having the fuel ahead of the axle in the new NASCAR chassis they are developing, but decided against it. The reason, he said, was that the greatest risk of fuel fires occurs in refueling during pit stops. The present rear location was considered preferable from that standpoint. This may be less of a factor where the races are short and there are no fuel stops.

ROLL SPRINGING PRELOAD IN MONOSHOCK SETUPS

I race a 1998 Dallara F398 F3 car in an amateur series in the UK. The car is fitted with a monoshock front suspension with Belleville washer stacks to control lateral movement of the rocker. I have attached a couple of photos and a copy of the setup page from the Dallara build manual. [To keep file size manageable for e-mailing, I am including only one photo in the newsletter. As it appears here, it's on its side. Rear of the car is to the left, left side of the car is up. The assembly shown is located on top of the footwell area of the monocoque tub.]

My questions concern spring preload. I have read many books and can find little information on the subject.

Question 1: The front spring is preloaded so that in a static condition there is sufficient load to support the weight of the car, i.e. there is no front droop travel. Can you explain to me the effect of increasing the amount of preload compared to, say, fitting a stiffer spring with less preload? Does this cause the car to behave differently in low-speed corners when there is little downforce, compared to high-speed corners?

Since the coilover only acts in ride, its action only affects the wheel rate in ride, and has no effect on the wheel rate in roll. It does affect ride height, and that has effects on both downforce (higher ride height makes less downforce, as a rule) and geometric roll resistance (roll center height – higher front ride height means more front roll resistance, as a rule, because the roll center usually rises and falls with the sprung mass).

With a setup as the questioner describes, with preload just equal to static load, if there is no downforce, the wheel pair cannot extend synchronously (in ride) past static position, but one wheel of the pair can if the other moves into compression a similar amount. If there is slight downforce, as at low speed, there can be a little ride movement in the extension direction. When there is a lot of downforce, the suspension can move quite a bit in ride extension before it reaches static position and tops out.

What happens if we preload the suspension more in ride? Assuming non-spring components to be rigid, we then get no ride motion at all in either direction, until force in the compression direction (from aero, road irregularities, or rearward acceleration, or any combination of these) overcomes the preload.

If, on the other hand, we add spring rate and not preload, we get less ride motion in compression for a given load increase, but we always get some.

So a stiffer spring begins to move with any load increase. A softer spring, with more preload, doesn't yield at all until the preload is overcome, but then moves more per unit of load increase.

Well, do we want the suspension to move more or not? It's a tradeoff. If the suspension moves more, wheel loads change less, but camber, suspension geometry, and aerodynamic properties change more. If we cinch down the suspension so it can't move, then we keep camber and geometry from changing, but wheel loads will vary greatly over irregularities, and the wheels may even become airborne. We get good aerodynamic consistency, up to the point where the wheels come off the ground. Beyond that point, ride height and pitch angle can change very rapidly.

Suppose we have to keep the car at a certain height at a certain peak speed, with a certain aero package. Suppose we have a choice of a soft ride spring, heavily preloaded, or a stiff ride spring, set to zero droop but with no other preload, or the stiffer spring allowed to move in droop.

What would give better behavior in low-speed corners? If there are bumps to contend with, the stiff spring not preloaded will be best, because it will alow the wheels to follow the bumps best. The stiff spring set for zero droop will be next best. The soft spring heavily preloaded will be worst, because it will essentially be locked solid in both directions.

If, on the other hand, the track is very smooth, the heavily preloaded ride spring may actually confer a slight advantage, because it will control camber a little better due to the lack of ride motion.

These are the same basic advantages and disadvantages we see with soft and stiff setups generally.

Question 2: Roll control of the front suspension is achieved by using Belleville washer stacks, which can be configured to various spring rates. The build manual recommends minimum and maximum amounts of preload that should be applied to each stack configuration. I understand that if the preload is lost on one of the stacks during roll, then the total stack stiffness suddenly halves, which is why there is a minimum preload. I guess that the maximum value is ther to ensure that stacks do not go solid. Apart from the above, am I right in thinking that changing the Belleville preload will have no effect on the car?

Assuming the washers themselves are fairly linear, as suggested by the listing of a single rate for each stack configuration, there should be no change in handling properties or rates with variations in preload, provided nothing unloads or goes solid.

You might be able to test the washers to see how linear they really are.

It will be apparent that you could create stacks that are non-linear if you wanted to. To do that, you would have to have some of the washers doubled, tripled, or quadrupled, and others not. None of the factory stacks include anything like that.

I wouldn't really suggest trying deliberately non-linear stacks, but if you did, that would create a situation where the system would become more sensitive to preload. If you had enough preload to compress the softer portions of the stacks to solid when the rocker is centered, you could create a stepped falling rate. If you had the softer portions not compressed to solid with the rocker centered, you could get a stepped rising rate instead.

I imagine you don't race on ovals, but this could get interesting when using asymmetical setups with diagonal percentages other than 50%.

WHY SO FEW REAR MONOSHOCK SETUPS?

With monoshocks appearing on more designs at the front, is there any reason, other than packaging, that they are never used on the rear of a car?

The Dallara system above is typical in that the roll mode is undamped. I do think it is possible to incorporate a damper in such a mechanism, but it would involve making your own, and the relatively small travel might be a disadvantage.

The Dallara information sheet has a partial sectional view of the rocker mechanism, showing something that looks like a needle valve, which conceivably might be part of a damping mechanism, but the questioner informs me that it's part of a position sensor, and there is no roll damping.

With a conventional two-coilover suspension, we get damping in both ride and roll.

If a car has no roll damping at either end, it will oscillate in roll and possibly in warp (opposite roll at the front and rear). Really, we'd like to have roll damped at both ends, and have at least the low-speed roll damping adjustable.

But if we are going to damp only one end of the car in roll, it should be the rear. Here's why: rear roll damping creates an anti-roll moment at the rear when roll velocity is outward, early in the turn. It creates an anti-de-roll (in other words, pro-roll) moment late in the turn when the car's roll velocity is inward. That loosens the car (adds oversteer) during entry, and tightens it (adds understeer) during exit. This tends to compensate for the tendency of yaw inertia to create understeer during entry and oversteer during exit.

Actually, we don't always want this, or if we do, we can have too much of a good thing. Driving style, brake bias, driver preference, and other factors enter into this, so really we'd like adjustable roll damping at both ends of the car. But it's better to have the roll damping loosening entry and tightening exit than vice versa, for the majority of setups, drivers, and situations.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2004

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL CENTER BELOW GROUND

Your October column [based on the June 2004 newsletter] *addressed a reader's question about the possibility of a roll center being deliberately placed below ground level. I've thought quite a bit about what I call "anti-jacking", and I've seen no real references to below-ground roll centers. If we put the roll center below ground, wouldn't this tend to put more force on the inside wheels than the outside wheels? Because of the non-linearity of the coefficient of friction, wouldn't this give, effectively, more total friction, and therefore more centripetal acceleration?*

Taking the last question first, if we assume that the inside and outside tires are symmetrical and identical, and are running at identical absolute camber, with the same camber direction relative to the turn (e.g. x degrees positive on the inside wheel and x degrees negative on the outside wheel), and assuming the car has 50% left weight at static, then we'd ideally like no load transfer inward or outward. If the outside tire has more favorable camber than the inside one (common in road racing), or the outside tire is larger (common in oval track racing, rules permitting), then we'd ideally like a bit more than 50% of the load on the outside tire.

As a practical matter, however, we always get more load transfer than we'd like, and we're always trying to reduce it. We don't really want more than half of the load on the inside wheels, but we're always trying to move the situation in the direction of more load to the inside.

The total load transfer for both wheel pairs, at a given lateral acceleration, depends only on the center of gravity height and the track width. The only way to get zero load transfer at both ends of the car would be to have the c.g. at ground level. The only way to get load transfer inward at both ends of the car would be to have the c.g. below ground level. This is of course impossible for a car as we would normally conceive it, running on a flat road. It would only be possible if the car could hang in a trench between two tracks, or if the wheels ran on elevated rails.

About all we can do with roll center height (or, more properly understood, with geometric anti-roll/pro-roll), springs, and anti-roll bars is to control what share of the inevitable load transfer is borne by the front wheel pair compared to the rear. We also can control the amount of roll, which has a small effect on total load transfer because the c.g. moves outward slightly as the car rolls.

We normally think of the sprung mass as a rigid body, supported on two roll-compliant support systems (the front and rear suspensions and wheel pairs). The two roll-compliant systems resist the rigid body's roll moment in parallel with each other. Each suspension system absorbs a portion of the roll moment proportional to its overall roll resistance.

That overall roll resistance has three components: geometric (from the suspension linkages), elastic (from the springs and anti-roll bars), and frictional (from the dampers and the mechanical or Coulomb friction in the system).

There is another component of load transfer as well: the unsprung mass component. This is commonly thought of as acting only through the tires, and not through the suspension. It would therefore be unaffected by suspension geometry. Actually, this is not strictly true with independent suspension, although it is correct for beam axles. A reader recently sent me some very insightful information about this, which we will take up in a future issue. The unsprung masses in an independent system do usually create some moments on the sprung mass, through the suspension linkage, which are resisted by the springs and anti-roll bars and are affected by both elastic coefficients and linkage geometry. We will ignore these effects for now, in the interest of simplicity.

If the roll center for the front or rear wheel pair – understood in the usual way, as the intersection of the front view force lines – is below ground level and between the wheels, that implies geometric pro-roll on both wheels. The outside wheel's linkage generates a downward jacking force in that wheel's suspension, and the inside wheel's linkage generates an upward jacking force. The resulting couple acts to roll the sprung mass outward, exaggerating roll. Considered in isolation, this does add load to the inside wheels, and remove load from the outside wheels.

However, if this moment were not resisted somehow, the sprung mass would accelerate outward in roll and wouldn't stop: the car would turn over. So it falls to the springs, anti-roll bars, and any frictional forces to resist the overturning moment, unassisted by any anti-roll forces from the linkage geometry. Additionally, the springs, anti-roll bars, and frictional effects must resist the pro-roll moment created by the linkages.

Therefore, if we have pro-roll geometry at both ends of the car (roll axis below ground), the elastic component of the roll resistance just gets very large, and there is still a net anti-roll moment from the suspension as a whole. Since the suspension is supported only by the tires, any moment generated in the suspension reacts through the contact patches and creates a load change there. That means there is net outward load transfer, even if the geometric moments are the "wrong" way.

Actually, it is theoretically possible to get inward load transfer, with a c.g. above ground level, at one end of the car only. I doubt that there is any real-world situation where we'd want such a setup, but it is an interesting hypothetical curiosity. If we used geometric pro-roll, combined with a wheel rate in roll of zero or nearly zero, at one end of the car only, we could achieve inward load transfer at that “soft” end. To get sufficient ride stiffness, we’d have to use a springing system that acted only in ride, such as the Z-bar at the rear of a Formula Vee or the third spring on a modern high-downforce car.

The other end of the car would then have to resist not only the overturning moment of the entire car, but also the pro-roll moment from the suspension at the “soft” end of the car. We would then get outward load transfer at the “stiff” end greater than the total outward load transfer for the car, and a small inward load transfer at the “soft” end, equal to the difference.

This would work up to the point where the “stiff” end lifted a wheel. The car would then immediately flop over onto the bump stops at the “soft” end. The “soft” end would then no longer be soft, and we would start getting outward load transfer at both ends of the car.

As for net downward jacking, in most cases if the front view force line intersection is below ground and between the wheels, we will get some net downward jacking. However, it is possible we could get net upward jacking if the inside wheel has a substantially steeper force line slope than the outside wheel. That might create an upward force from the inside wheel’s linkage greater than the downward force from the outside wheel’s linkage, despite the smaller contact patch force at the inside wheel.

This would imply a force line intersection off-center toward the inside wheel, though still within the track width. That is not necessarily an unrealistic case. We could easily encounter it in a lowered production car strut suspension, in a rolled condition. This is an interesting case, because it illustrates that there are situations where the car does not roll about the force line intersection, or even do anything close to that. If it did, it would have to move downward rather than upward with roll – and if the upward jacking on the inside wheel exceeds the downward jacking on the outside wheel, the car clearly doesn’t do that at all.

ROLL MOMENTS FROM LONGITUDINAL ANTI

Some people tell me that anti-dive and anti-squat act to stiffen the suspension when forward or rearward forces are present at the wheels. Does that mean these effects add roll resistance? How does this really work?

Anti-dive at the front wheels does impose a bit of a roadholding penalty, because it requires the contact patch to move forward as the suspension compresses, at least if we imagine the imagine the wheel as locked. Or, we might view this effect as requiring an increase in wheel rotational speed with respect to the caliper, as the suspension yields to a bump. The effect varies somewhat with the abruptness and height of the bump, the outside diameter of the tire, whether the hub moves forward in compression or not, and how hard we’re braking. However, anti-dive, even in an amount that completely eliminates dive (100% anti-dive), does not completely lock up the suspension as some authors have suggested. It merely acts counter to our desire to have the wheel move backward relative to the car, as well as upward, when the wheel hits a bump.

At the rear of the car, things are a bit different. Anti-lift in braking and anti-squat in forward acceleration cause the contact patch to move rearward in compression rather than forward, while of course the bumps still come at the wheel from the front. So rear longitudinal anti actually improves the system's ability to yield to a bump.

Jacking forces, whether lateral or longitudinal, do not in themselves add to wheel rate or subtract from it, provided that the jacking forces do not change with wheel movement. The jacking force simply acts in parallel with the wheel rate or elastic forces, which are displacement-dependent. That doesn't mean the jacking forces can't create roll moments or affect wheel loads. They definitely can.

While anti effects do not necessarily vary as the suspension moves, it is very common for both longitudinal and lateral anti effects to vary with suspension displacement. Most often, both lateral and longitudinal anti diminish as the suspension compresses, and increase as the suspension extends. This is not so in all cases, however. A counter-example would be the trailing arm front suspension on a VW beetle. The arms are equal-length and parallel to each other, and at static condition they slope down a bit toward the rear. As the suspension compresses, the arms quickly reach horizontal, then begin to slope upward to the rear. The suspension goes from decreasing pro-dive to increasing anti-dive. The direction of change is consistently toward anti-dive as compression increases.

A NASCAR front end is an extreme case of the opposite, and more common, tendency. It changes rapidly toward pro-dive with compression, because the lower control arm is a semi-leading arm (pivot axis angled dramatically outward at the rear in plan view) while the upper control arm is almost a purely transverse arm (pivot axis close to longitudinal in plan view).

If the slope of the suspension's longitudinal force line varies with suspension displacement, then assuming a constant longitudinal force at the contact patch, the jacking effect can act in a manner analogous to a spring force: it may increase or decrease according to displacement. It won't necessarily increase with compression, however. If it does increase with compression, as in the case of the VW, it can loosely be thought of as adding wheel rate. If it decreases with compression, as with the NASCAR suspension, it can be similarly thought of as subtracting wheel rate.

On the face of it, we might suppose that if the front wheels have the same amount of anti-dive – that is, the same longitudinal force line slope – then their longitudinal-force-induced jacking forces will lift both the right front and the left front corners of the car with the same force, and this will not create any roll moment, although it will create a pitch moment. Therefore the anti will neither wedge nor de-wedge the car.

This is true, but we must remember that the longitudinal forces at the two contact patches may not be equal. In fact, if we are cornering, they are unlikely to be equal.

The longitudinal forces at the front contact patches actually come from two sources. One is braking, which may or may not be present. The other is the induced drag that a tire produces when running at a slip angle, which is present any time we're cornering. That induced drag varies with the amount of load on the tire, and it is also affected to some degree by camber and toe. Generally, though, it is safe to say that the induced tire drag is greater on the outside wheel. Therefore, the jacking force on the outside wheel will be greater for a given force line slope than on the inside wheel. That will tend to wedge the car (add diagonal percentage, i.e. outside front tire load plus inside rear, as a percentage of the whole), and tighten it (add understeer).

Braking forces, on the other hand, tend to be more nearly equal. Theoretically, if we are short of the point of lockup, with no front tire stagger, and the brakes are as identical as we can make them, the braking forces will be identical at the two front wheels. Actually, with no tire stagger, the outside tire will act slightly smaller if we are braking while cornering, because it will deflect more vertically and therefore have a reduced loaded radius. Reducing the loaded radius reduces the effective radius (increasing the revs/mile), though not by the full amount of the deflection change. This effect will make the braking force slightly larger on the outside wheel.

If we are braking and cornering at the same time, we will have both a drag component and a braking component. If we are braking hard and cornering gently, the rearward forces at the front contact patches may be fairly equal, especially if the car has some toe-out. If we are braking gently and cornering hard, the rearward force may be substantially greater on the outside wheel, especially if there is some toe-in.

When we are off the brakes entirely, and cornering hard, we can say fairly confidently that the rearward force will be greater on the outside front.

Can we say, then, that adding anti-dive makes the car tighter? Well, we almost can. If we add anti-dive only on the outside wheel (right front, for oval track), that will tighten the car. It will do this even when we're not braking. If we add anti-dive evenly on both front wheels, that may also tighten the car, due to the greater rearward force on the more heavily loaded tire. Any such effect will tend to be more pronounced in hard cornering than in hard braking.

However, if we increase anti-dive only on the inside wheel (left front, for oval track), that will loosen the car (add oversteer) instead. This effect will be present whether we are braking or just cornering. So this would be a situation where we'd be increasing overall anti-dive yet adding oversteer.

One might suppose that adding anti-dive on just the inside or outside wheel is impossible for road racing, but as we have noted, suspension layouts vary as regards how anti-dive changes with suspension movement, and such effects can be used to control the left/right balance of anti-dive when the car is in a rolled condition, even when the car has to turn both ways. Such effects are often hard to manipulate on an existing car, but they deserve consideration in the design phase.

All of the above is based on the principle that adding diagonal percentage tightens the car and reducing diagonal percentage loosens it. When applying these principles, it is also important to bear aerodynamics in mind. More anti-dive will cause the front of the car to ride slightly higher through the turns, particularly with soft front springs. If static ride height or valance height are not adjusted for this, the greater ride height when cornering may add understeer purely by reducing front downforce.

Now, what about anti-lift and anti-squat at the rear? As at the front, the jacking forces will depend on both the force line slopes and the magnitude of the forces at the contact patches. And, as at the front, any roll moments created will depend on the difference in jacking forces at the right and left sides. Two things are different at the rear: we can have forces forward or rearward (this whole discussion assumes rear wheel drive), and we have various kinds of differentials (or lack of) that can influence the relative magnitude of the longitudinal forces, and in some cases even their relative direction.

Like the front tires, the rears generate drag when running at a slip angle. However, it is unusual for that to be the only longitudinal force. The rear tires are almost always either propelling the car or retarding it. Even in roughly constant-speed cornering, the rear tires are making enough forward force to overcome the front tire drag and the aerodynamic drag.

The rearward forces at the rear contact patches when braking or trailing the throttle will tend to be fairly equal if we have an open differential. If we have a spool or a limited-slip, however, any rearward force will be greater on the faster (usually the outside) wheel. When under power, again the forces will be fairly equal with an open diff, but any locking effect will result in more force at the slower (usually the inside) wheel. At least, that holds true up to the point of inside wheelspin. Then the outside wheel may make more forward force than the inside.

All of this makes it fairly complex to predict the distribution of longitudinal force at the rear. However, we can say this much: in braking, more anti-lift or less pro-lift on the inside rear loosens the car (adds oversteer); more anti-lift or less pro-lift on the outside rear tightens the car (adds understeer). Under power, more anti-squat or less pro-squat on the inside rear tightens the car (adds understeer); more anti-squat or less pro-squat on the outside rear loosens the car (adds oversteer). Effect of more anti-lift or anti-squat geometry added evenly on both sides depends on the distribution of longitudinal force between the two rear contact patches.

Distribution of longitudinal force also affects handling balance because it creates yaw moments. In general, we can state that more longitudinal anti of any type intensifies these effects. For example, more induced drag at the outside front creates a yaw moment promoting understeer. If there is more anti-dive, there is also an increase in diagonal percentage, which intensifies the tightening. If there is more forward force at the inside rear, that creates an understeer-adding yaw moment. If there is ample anti-squat, again we get an increase in diagonal percentage, intensifying the effect. More rearward force at the inside rear creates a yaw moment that adds oversteer. More anti-lift there reduces diagonal percentage, again intensifying the effect. So we may say that, in general, increased longitudinal anti geometry makes a car more sensitive to its tires' load and force distribution.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January/February 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

IDEAL SPINDLE OR STEERING GEOMETRY

I have a fairly simple question for you, but I don't know how simple the answer will be. In any literature I read, we are always dealing with things like kingpin inclination (KPI) and scrub radius. Usually these cannot be eliminated in a sedan with big tires, and we are told in the books to "not have too much". But if I have more design freedom, and it is possible to get zero scrub, and zero or very close on kingpin inclination, would this be a favorable setup? I know that scrub + caster can give weight jacking, and that KPI gives a self-centering effect, but otherwise both of them seem like something I want to get rid of. So to sum up, is it a good idea to completely eliminate scrub, or KPI (or any other variables) entirely if the design allows?

To help readers who are less conversant with steering geometry, I am inserting some comments of mine from the August 2002 newsletter, which explain various steering geometry parameters and their effects:

The *steering axis* is a line about which the wheel steers, usually through the two ball joint centers of rotation in an independent suspension, or the kingpin axis in a beam axle. This line can be defined by the point where it intersects the ground and by its angular orientation. These are commonly described in terms of the X and Y coordinates of the ground intercept, with respect to a local origin at the contact patch center, and the transverse and longitudinal angles relative to ground plane horizontal.

The front view distance from ground intercept to contact patch center, or local Y, is called *scrub radius*, or *steering offset*. It would make more sense to call the top view distance from ground intercept to contact patch center the scrub radius, but most people use the term to mean the Y or transverse component of this. This quantity is generally considered positive when the contact patch center is outboard of the ground intercept.

The side view distance from ground intercept to contact patch center, or local X, is called *trail*, or sometimes *caster trail* or *mechanical trail*. It is positive when the ground intercept is forward of the contact patch center.

The front view angle of the steering axis from ground vertical is called *steering axis inclination* (SAI), or sometimes *kingpin inclination* (KPI). It is positive when the steering axis tilts inboard at the top, which is almost always the case.

The side view angle of the steering axis from ground vertical is called *caster*. It is positive when the steering axis slopes rearward at the top.

These parameters are controlled partly by the design and adjustment of the control arms, and partly by the design of the spindle, or spindle/upright assembly, together with the hub and wheel.

The term *spindle* can mean either the stub axle, or pin, that carries the bearings, or the assembly including this pin and the upright, especially when these are one piece.

The spindle or spindle/upright determines two important parameters: *spindle inclination* and *pin lead* or *pin trail*.

Spindle inclination is the front-view inclination of the steering axis, relative to pin or wheel vertical, as opposed to ground vertical. Spindle inclination approximately equals SAI minus camber. Spindle inclination is almost exactly identical to SAI when camber is zero. It is exactly identical when both camber and caster are zero.

The steering axis and the wheel axis do not have to intersect, unless we want the right and left uprights to be identical parts, with bolt-on steering arms and caliper brackets. The steering axis can pass behind the wheel axis, as it does on a bicycle. The perpendicular distance between the two axes is called *pin lead*. This is equivalent to the dimension we call fork rake on a bicycle. If the steering axis passes in front of the wheel axis, that's *pin trail*. So pin trail is negative pin lead, and vice versa.

Effective pin length is the distance, along the wheel axis, in front view, from the steering axis to the wheel centerplane. This distance depends on the wheel and hub as well as the spindle/upright.

We now have sufficient vocabulary to describe and discuss basic steering and spindle geometry. If we can specify all the quantities above, we have enough data to construct a stick model of the basic steering geometry.

We may want to add steering arms. For purposes of spindle/upright design, we can define the position of the outer tie rod end with respect to the pin and the steering axis. We may define a height from pin axis to tie rod end center of rotation. To do this in a manner appropriate for drawing the upright, or inspecting it when removed from the car, this should be the vertical dimension in side view, assuming zero caster and camber -- in other words, we are projecting to the wheel plane, and taking the steering axis in side view as our local vertical.

In such a side view, we may construct a horizontal line from tie rod end to steering axis. This is our *side view steering arm length*.

We may project a top view from the side view, and locate the lateral position of the tie rod end. If we have a longitudinal line corresponding to the side view steering arm described above, we may construct a transverse line from it to the tie rod end, and measure that distance. This we may call *steering arm offset.* It will usually be outboard for a front-steer layout, and inboard for a rear-steer layout. I don't know what sign conventions other people use, but I generally call outboard positive for front steer and inboard positive for rear steer. Thus positive offset is the direction that gives us positive Ackermann.

In terms of coordinates, we are establishing a local origin where the side view steering arm meets the steering axis. The side view steering arm length is our local X, and the steering arm offset is our local Y.

This doesn't mean there's anything wrong with assigning global or front-suspension coordinates to the tie rod ends when doing an overall front end layout. I'm just pointing out that at some point you will have to deal with the spindle/upright/steering arm unit as a sub-assembly, off the car, and it helps to be able to measure and discuss it that way too.

Now we have a fairly complete vocabulary to describe steering geometry, so we can discuss what effects these parameters have.

Trail causes *lateral* forces at the contact patch to produce a torque about the steering axis. This causes the steering to seek a gravitational/inertial center. The driver feels lateral cornering force through the steering. He also feels the lateral force that the tires must generate to make the car run straight on a laterally sloping, or cambered, road surface. It is worth noting that this is only one component of the self-centering forces the driver feels. Another is the tire's own *self-aligning torque*, which is present whenever the tire runs at any slip angle. This will provide some feedback of cornering force even in the total absence of trail. This effect is sometimes described as mimicking trail. The amount of tire self-aligning torque, divided by lateral force, is sometimes called *pneumatic trail*. Note that this is a calculated value which depends on tire properties, and not an actual steering geometry parameter.

One important distinction between the forces from trail and tire self-aligning torque is that tire self-aligning torque is not a linear function of lateral force. It builds at a decreasing rate as lateral force increases, and at a point a bit short of maximum lateral force it actually begins to decrease. This means that if our car has little or no trail, the steering will start to go light a bit *before* the point of tire breakaway. Some argue that this is a good thing, especially for a passenger car, because it gives the driver a signal to ease up short of the point of actual loss of control. In a race car, this type of steering feel requires that the driver be accustomed to driving just a controlled increment beyond the point where the steering wheel tells him/her that the limit of adhesion has been reached. If the driver is used to having more trail, he/she will often find this very difficult.

Trail also causes a small lateral movement of the front of the car with steer, in the direction of steer. We might call this *steer yaw*. It can rationally be argued that this improves turn-in, both by yawing the car promptly and by causing the rear wheels to develop a slip angle promptly.

Scrub radius or steering offset causes *longitudinal* forces at the contact patch to generate a torque about the steering axis. If right and left scrub radii are equal and longitudinal forces at the right and left wheels are equal, no net torque at the steering wheel results. The driver feels the *difference* between the longitudinal forces at the front wheels. The driver feels one-wheel bumps, brake pulsations, and crash impacts where one wheel hits something, in direct proportion to scrub radius.

A car with a lot of scrub radius is sensitive to wheel imbalance and tire and brake imperfections, has a lot of "wheel fight", and has greater tendency to injure the driver's hands in one-wheel crash impacts or curb or pothole impacts. A car with very little scrub radius is less subject to these problems, but the steering will tend to be numb and uncommunicative.

A car with large scrub radius *may* steer more easily at parking speeds, depending on other parameters, provided the brakes are not applied. This is because the wheels can roll as they steer rather than purely scuffing. With the brakes applied and the car stationary, a car with a small scrub radius steers more easily.

Caster causes the front wheels to lean in the direction of steer. With a given spindle/upright geometry, more caster implies more trail.

Caster combined with trail causes s*teer drop* or *steer dive*. The front of the vehicle drops as the wheels steer away from center, if caster is equal on right and left. This tends to cause an anti-centering force at the steering wheel. It is the reason why the front wheels of a dragster at rest tend to flop to one side or the other.

Caster combined with scrub radius causes the car to drop as the wheel steers forward (toes in), and lift as the wheel steers rearward (toes out). When this occurs on the right and left wheels as one steers forward and the other steers rearward, the result is *steer roll.* The car leans away from the direction of steer. The wheel loads also change. The car de-wedges: the inside front and outside rear gain load; the outside front and inside rear lose load. This effect can help the car turn in slow corners, especially with a spool or limited-slip differential. In excess, it can create low-speed oversteer and over-sensitivity to steering angle. In general, cars running on lower-speed tracks need more steer roll, and cars on fast ovals should have very little.

The camber change associated with caster is favorable, particularly for road racing cars, which usually cannot get favorable camber on both front wheels any other way. We can have too much of this good thing, but that's extremely uncommon.

Steering axis inclination (SAI) causes both front wheels to gain positive camber as they steer away from center.

SAI combined with scrub radius causes *steer lift*. The front of the vehicle rises as the wheels steer away from center. This induces a self-centering force in the steering which seeks vehicle center rather than inertial/gravitational center. This is particularly useful in passenger cars because it reduces the car's tendency to follow road camber, and therefore reduces the need for the driver to pay close attention in casual driving on roads with varying slope. The centering force also tends to suppress steering shimmy.

In race cars, the camber change associated with SAI is unfavorable on the outside wheel. The self-centering force increases steering effort, which is a factor for any vehicle without power steering. It also creates what could be considered a false message to the driver about the lateral forces present at the contact patches. There is therefore a rational case for using more caster and less SAI in a race car.

With the packaging constraints we usually face, more SAI generally implies less scrub radius. The main limitation will often be how far outboard we can place the lower ball joint without having it too close to the brake disc. If the wheel has generous negative offset, we may instead be limited by the wheel rim hitting the control arms in some combinations of suspension motion and steer. Either way, we often cannot place the entire steering axis as far outboard as we would theoretically like to. Using SAI allows us to at least get the ground intercept further outboard in such cases. With MacPherson strut front ends, large amounts of SAI are necessary if we are to obtain any camber recovery in roll.

Consequently, in many cars we see SAI used for reasons not directly related to SAI's own dynamic effects.

A full discussion of **Ackermann effect** (increase of toe-out with steer) is beyond our scope here, but we can at least say that in low speed turns with the wheels steered into the turn, the car generally needs toe-out on the front wheels. For high-speed sweepers or ovals, the front wheels generally need toe-in instead. The key determining factor is whether the *turn center* -- the instantaneous center of curvature of the car's path -- is ahead of or behind the front axle line. Other determining factors include the tendency of the loaded wheel to want a larger slip angle than the unloaded one, and what yaw moments we wish to create with the tire drag forces.

The attitude of the front wheels at any given instant depends on both the static toe setting and the change in wheel-to-wheel toe with steer. This means that optimum Ackermann depends on static toe setting.

It should be clear, then, that there is no such thing as perfect Ackermann properties. But we can at least say some definite things about what geometric parameters will do to Ackermann. In particular, increasing steering arm offset increases Ackermann effect.

Ackermann for oval track cars is often asymmetrical. The side view steering arm length is less on the left wheel than on the right. This produces more Ackermann when steering left than when steering right.

We should mention that if we are willing to tolerate a bit of additional complexity, there are ways around some of the tradeoffs in steering geometry. For example, it is possible to create a self-centering force by springing the steering system. This can mimick the self-centering that we get from SAI, without the adverse effects on camber. We can also damp the steering to reduce kick and shimmy.

We can get small SAI and small scrub radius at the same time by using compound control arms (two single links replacing the usual wishbone or A-frame) and dual ball joints. This gives us an instantaneous virtual ball joint outboard of the linkage itself. We can adopt this arrangement at the upper end of the upright, or the lower end, or both.

When using dual ball joints, it is important not to splay the two links too widely. Otherwise it may be possible for the linkage to snap over center if it has to take an impact near full lock. As a rule of thumb, it is probably advisable to have the front link around 25 degrees forward from transverse and the rear link around 25 degrees back.

Ackermann also gets interesting when using dual ball joints. As we steer, the steering axis moves rearward with respect to the ball joints and tie rods on the outside wheel and forward with respect to the ball joints and tie rod on the inside wheel. This means we lose Ackermann as we steer if the steering linkage is behind the wheel axis (rear steer layout), or we gain Ackermann as we steer if we have a front steer layout. Conversely, it is difficult to get the outer tie rod ends far enough outboard to have positive Ackermann at small steer angles with front steer. With rear steer, it is easy to get the tie rod ends inboard far enough to have initial positive Ackermann, but it is hard to avoid having too much initial positive Ackermann. It becomes very important to evaluate the steering geometry through the full range of steering motion. In many cases, we find that outer tie rod end packaging limits our steering axis location, rather than ball joint and control arm packaging.

I might also mention that we can get small SAI and small scrub radius together by using an actual kingpin, as on a beam axle, and bushings rather than ball joints at the top and bottom of the upright. The kingpin can then be placed much further into the wheel and brake than would otherwise be possible. This was actually the most popular way to build independent front ends up until the mid to late 1950's.

Okay, returning to the current question, is it desirable to completely eliminate SAI/KPI, and/or scrub radius/steering offset? As the comments above indicate, these questions relate heavily to steering feel and the answers are therefore to some extent a matter of personal preference.

The nature of the track influences the decisions too. There is a stronger case for minimal SAI and scrub radius on a high-speed oval than on a street circuit. There may also be a stronger case for minimal SAI and scrub radius for off-road racing than for pavement competition.

The high-speed oval requires steadiness and freedom from vibration at high speeds. We don't want the car's properties or wheel loads changing much with steering wheel movement. We want stability and predictability more than responsiveness or communicativeness. The driver will tend to drive more in "open-loop" mode, steadily holding a smooth line, by eye, through the wide turns rather than responding reflexively to constantly changing information received through the steering wheel. He/she will still need to feel whether the front wheels have grip, and may still want the steering to seek center, but the need to sense the road wheels' exact position is much diminished.

On a street circuit or tight road course, the driver will tend to operate more "closed-loop", relying on input from the steering wheel to tell him/her where the wheels are on the road. The driver senses this by seeing the bumps in the road at a distance as they approach, and then feeling through the steering what the wheels actually run over. This is true even in a car where the driver can see every bit of the front wheels, because at speed the driver has to be looking ahead at where he/she wants to go, and not down at the wheels. Speed on a tight course depends heavily on using all the road, and that depends on being able to position the car precisely.

Road courses also tend to be bumpier and contain more varied surfaces than high-speed ovals. Consequently, the driver relies on the steering for information about the ever-changing grip level. This really is more a matter of trail than scrub radius, but in braking in conditions where there is more grip on one side of the car than the other, scrub radius may make it easier to detect lockup on the slicker side. On the downside, when the steering pulls toward the side with more traction, the driver needs to resist the pull, and maybe even steer against it a little, to keep the car pointed straight. For this reason, Volkswagen used to advertise the negative scrub radius on their cars as an aid to stability in braking.

If the car has really stiff suspension, as with current high-downforce cars, the driver will have little problem sensing which wheel hits a bump, even if he/she can't feel it through the steering wheel. On the other hand, a car with soft suspension has greater need for communicative steering, because it doesn't transmit so much information by other paths.

If we are running off-road, precise vehicle placement may be somewhat less critical, although there are situations where it is important. The wheels will be hitting big bumps all the time, and it becomes important to avoid beating up the driver. If the vehicle has a beam axle in front, steering oscillations and kick from lateral movement at the contact patches on one-wheel bumps are particularly troublesome. Consequently, there is a case for a small scrub radius for off-road use.

There are a lot of vehicles out there that successfully use zero scrub radius with zero SAI. Most of them have only one front wheel, but they work fine. The faster ones do often use steering dampers. I am referring of course to motorcycles. It should be equally possible to have similar geometry on two front wheels, all packaging considerations permitting.

Is that the best solution? As previously noted, it depends on what you want, or what the driver wants, in terms of steering feel. If you want light steering and minimal feedback of bumps, wheel vibration, and brake pulsation, try zero scrub radius and zero SAI. If you want a "natural" feel of which wheel is hitting a bump, you probably want some positive scrub radius. If you are trying to unload the inside rear wheel in tight turns, you want a lot of scrub radius.

It is possible to have a large scrub radius with little or no SAI. This combination is prone to oscillation, particularly at low speeds, and particularly with large caster settings. However, in a race car, it may make sense to either live with that, or damp and/or spring the steering to control it, rules permitting.

It is also possible to have large SAI with small scrub radius, or even negative scrub radius. Indeed, ample SAI is generally the easiest way to get a small scrub radius. There is a case for this sort of geometry in street cars, especially with front wheel drive and strut suspension, or with a beam axle front end. It is not a good idea for a situation where we are designing for racing, with rear wheel drive and generous design freedom.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March/April 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

FINER POINTS OF ANTI-ROLL BARS ON STOCK CARS

I race primarily on two ovals, one a 1/3 mile with 7 degree banking and short straights, and one a 5/8 mile with 15 degree banking and long straights. I am not clear on some of the finer points of setup and adjustment of front anti-roll bars for a specific track, especially for a car running at two tracks such as these.

Most cars in my area run a solid link to the anti-roll bar on the right, and a slider link on the left, where the left front A frame in droop would cause tension on the link, but the left front in bump would remove any tension. I have run many cars with this setup, and a few with solid links on both sides. I don't really understand why there would be a preference for one over the other, or in what situation the solid links on both sides would be preferred over the slider setup.

I have found from home-made data acquisition that once a car is up to speed on the 1/3 mile oval, the left side never returns to ride height, as the car always has considerable g's throughout the lap. Many competitors run play in their anti-roll bars, as much as ¼ to 5/8". While it is easier to duplicate a setting if there is a specific gap to measure rather than a preload when making last-minute changes in tire size or wedge adjustments, it seems to me it might be better if there is less spring in the front of the car and less or zero gap in the bar.

For most bar play settings, the gap is gone as soon as the car comes up to speed, and the gap doesn't return until the cool-down lap at the 1/3 mile oval. Running play in the bar at a track with long straights, where the chassis returns to zero roll angle on the straights, seems like it would create a variable roll rate at the front, as the initial roll resistance would be only on the RF spring. Then the anti-roll bar and the resistance from the extension of the LF spring would kick in, greatly increasing the front roll resistance.

The car running the 7 degree banked turn would be in true roll, with the right frame rail down and the left rail up, while at the 15 degree track, the right rail would be down, but the left rail would be at approximately static ride height in the turn. It would seem that the bar would be half as stiff dynamically for the same right side travel at this track.

Another concern I have is that I'm not always sure in which situation it would be best to change the bar rate as opposed to the front spring rate. I realize the bar controls only roll, but in many situations where a rate change is desired, I have a hard time deciding which would be best, bar or spring.

Concerning small adjustments in bar play made at the track, consider a situation where a car running at the 1/3 mile, 7 degree track has a bar rated at 800 lbs. per inch of movement at the right front, at the location where the bar link is mounted, and ½" of static play. Suppose an adjustment is made to the bar reducing play to 3/8", or 1/8" less than before. This would presumably increase roll resistance by 100 lbs. For a particular right front travel, but would not increase the roll ***rate****. How would this compare to changing to a bar that rates at 900 lb/in, with the original ½" play retained?*

It has long been common on oval track stock cars to use some type of "soft link" on the front anti-roll bar. These may take the form of a slider, a chain, or a pad on the anti-roll bar arm that bears against a pad on the lower control arm. The last of these is the form commonly seen in the upper divisions of stock car racing. Chains and sliders are common in "street stock" type cars, where the suspension components actually come from production passenger cars.

All these variations do the same thing: they create a connection that transmits force in only one direction. The anti-roll bar resists rightward roll only. The bar may be preloaded, in which case it will resist leftward roll up to the point where the preload is relieved. (Some sanctioning bodies prohibit preloading the bar. The reason for this is a mystery to me.)

The reason the soft link is there at all is to make the front suspension more compliant when the driver gets the left front wheel on the flat apron of a banked track. In this situation, the car tends to go loose (develop oversteer) and spin. This is partly because of the leftward yaw moment created when the left front encounters increased resistance, and partly because the car de-wedges (load increases on the left front and right rear, and decreases on the right front and left rear). If the suspension is more compliant, both effects are reduced, especially the de-wedging.

However, even in this situation the front suspension only acts softer beyond the point where the left front suspension compresses enough to put slack in the bar. If the bar has no play and no preload, the front suspension must go into leftward roll (left front compressed more than right front) before the soft link has any effect. And the soft link will make the car looser if the driver needs to turn right to avoid a wreck. So the soft link is a mixed blessing.

Wheel rates and arm end rates of anti-roll bars are a confusing subject for many. The main cause of this confusion is that there is no agreed convention as to what an inch of arm end or wheel motion means. Does it mean an inch of motion at just one wheel? An inch of motion in opposite directions at both wheels? An inch of difference between the two wheels?

I find it convenient to express the rates of all springing devices, including interconnective ones such as anti-roll bars, in terms of pounds per inch per wheel. This agrees with the way we always express the rate of non-interconnective springing devices, i.e. the main springs or ride springs. We then have numbers that easily tell us what change of anti-roll bar rate equates to what change of spring rate.

However, it has become customary to rate anti-roll bars by testing them in devices originally intended for testing non-interconnective torsion bars. The bar is placed in a fixture that holds one end of the bar stationary, and the arm on the other end of the bar is moved one inch. The force required to do this is measured, and that is taken as the bar's rate in pounds per inch.

This isn't wrong. But it can create confusion when we try to translate it into pounds per inch per wheel, because what we have from the test is pounds per inch per wheel *pair*. An inch of roll per wheel pair is only half an inch of oppositional motion per wheel. Therefore the number from the testing machine is only half the bar's equivalency to wheel springs, if the bar and the springs acted at equal motion ratios.

On the other hand, it is common for anti-roll bar manufacturers to express the rate of their bars the way I like to see it, as pounds per inch per arm end, not pounds per inch per arm end pair. This sometimes leads to acrimonious exchanges between bar manufacturers and customers who test a bar they just bought and find that it "rates" at only half the advertised value.

Often, we encounter situations where we do not have "pure" roll motion. Pure roll, for a front or rear wheel pair, would mean equal amounts of suspension motion at each wheel, in opposite directions. Pure ride would mean equal amounts of motion at each wheel, in the same direction. Ride and roll are the two modes of motion for a front or rear wheel pair. Any possible motion of that wheel pair can be resolved into some amount of ride and some amount of roll.

Applying this to the questioner's example of a situation where a car in a banked left turn rolls purely by compressing the right front suspension, and the left front neither compresses nor extends, that is a condition of equal ride and roll. If the right front compresses an inch and the left front doesn't move, we have half an inch per wheel of rightward roll, plus half an inch per wheel of compressive ride. On the left wheel, the effects are subtractive and exactly cancel. On the right wheel, the effects are additive, and we have twice the half-inch per mode, or one inch of compression.

If we compare this to a pure roll situation (zero ride motion) which creates an inch of compression at the right front, yes, the bar creates twice as much force. But the bar's rate is the same. The difference is that there is twice as much roll. If the roll were equal to the previous example, the right front compression would be only half an inch, the difference between right front and left front would still be an inch, and the force generated by the bar would be the same as in the first example.

Once we learn to separate the ride and roll components of the motion, and remember that the bar acts only in response to the roll component, things get much simpler.

Regarding how much of the front wheel rate in roll should come from the springs and how much from the bar, it mainly comes down to how much we want the front of the car to drop in the turns due to the banking. This in turn relates to both suspension geometry and aerodynamics.

It is normal in all stock car racing classes to have a minimum static ground clearance requirement. The car has to pass over a barrier of a certain height to get through tech inspection. Yet we would like the valance to run just off the track, and have as much forward rake in the car as possible, through the turns. This gets us the greatest aerodynamic downforce available, within the bodywork and ground clearance rules.

Additionally, with passenger car front suspension, we usually have insufficient camber recovery in roll for racing. The control arms are close to parallel at static, causing the wheel to change camber very little in ride, but a lot in roll. This is good for tire wear in gentle driving, but not good for hard cornering. In some cases, the rules will allow us to use extended upper ball joints to improve the geometry for racing. Sometimes this may be enough, but in most cases the geometry will benefit if we run the car lower. To do this, within the ride height rule, we need to soften the springs and stiffen the bar.

In the upper divisions of stock car racing, we can get as much camber recovery as we want, so only the aerodynamic factor argues for soft front springs. In general, with soft front springs and a big bar, we'll want the front view instant centers a bit higher and further from the wheels at static than we'd want with stiffer springs and a softer bar.

In either case, a flatter track calls for more of a soft-spring/big-bar approach than a steeply-banked track.

Regarding rate of a bar with slack in it: the bar has a rate of zero until the slack takes up. Then it has the same rate it would have without slack.

We can speak of the bar's rate as an instantaneous value, at any given point in roll travel, or we may speak of its average rate, over a specified interval of travel. When there is slack in the system for a portion of the interval, but not at the end of the interval, the bar has the same instantaneous rate at the end of the interval whether there is slack or not, but it has less average rate over the interval if there is slack.

Relating this to the questioner's example comparing an 800 lb/in rate with 3/8" of slack (let's call this Case A) to a 900 lb/in rate with ½" of slack (Case B), we may say that these two systems have the same instantaneous rate (namely zero) at any displacement less than 3/8"; the same average rate (450 lb/in) over the displacement interval 0" to 1"; the same force (zero) over the displacement interval 0" to 3/8"; and the same force (450 lb.) at 1" displacement.

Case A has less rate than Case B when the slack is absent; more force at displacements between 3/8" and 1"; and less force at displacements greater than 1". At ½" displacement, Case A generates 100 lb. more force than Case B. From there, the forces converge to a point of equality at 1", then diverge again, with Case B making more force, beyond 1" displacement. At 1½" displacement, Case B generates 100 lb. more than Case A.

In terms of car behavior, if all other setup parameters are equal we would expect Case A to have more understeer (or less oversteer) from 3/8" displacement to slightly less than 1" displacement; both setups to have similar understeer around 1" displacement; and Case B to have more understeer at greater displacements. We would expect Case A to have a lesser tendency to get looser (or a greater tendency to get tighter) as grip diminishes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May/June 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TOO MUCH LEFT PERCENTAGE?

My question is regarding left side weight percentage on oval track cars, specifically dirt Late Models. I have heard it generally stated that more left side is better in all situations, and I see a lot of paved track classes have limits on left side percentage. I understand the concept of load transfer and equal tire loading in steady state cornering.

My question is about the point of diminishing returns. My understanding would be that as grip decreases or banking increases, left side weight should be reduced to keep the left side tires from being more heavily loaded than the right sides. Is this a correct assumption, or have I missed something?

In theory, yes it is possible to have too much left percentage, and to have the left tires more heavily loaded than the right tires even at the limit of adhesion in steady-state cornering. In almost all cases, however, practical constraints or rules will stop us short of that point.

We can also have too much left percentage for the tire package short of that point, if the left side tires are smaller than the rights, or if the lefts are inflated to a much lower pressure than the rights.

Or, we might conceivably want more than 50% left dynamically, if the left tires are about as big as the rights, and we have a rule requiring a hard tread compound on one or both of the rights but not on the lefts.

Let's consider a simple, if not very typical, case. Suppose we have a car with a one foot c.g. height, a six foot track width, and identical right and left tires. Suppose that the overall coefficient of friction is 1.00. That would be about what we'd get from sticky street-legal radials. For this car to have 50% left dynamically at the 1.00g lateral acceleration that those tires will theoretically sustain, it would need 66.7% left statically. That's a wider, lower car than most, on tires with less grip than racing slicks.

If the same car is fitted with racing slicks that have a coefficient of friction of 1.30, the static left percentage needed to have 50% left dynamically increases to 71.7%.

If the car has a wing that acts equally on the right and left tires, lateral acceleration increases and the desired static left percentage goes up still more.

What happens if we put the car on a banking? It's a bit surprising. If the coefficient of friction stayed the same, the ratio of car-horizontal (y-axis, per SAE conventions) force to car-vertical (z-axis) force would be unchanged, although all forces would increase. This assumes the car is at the limit of adhesion both with and without the banking, not at an identical y-axis acceleration or an identical earth-horizontal acceleration.

However, due to the same tire load sensitivity that makes us want equal loading, on the banking the coefficient of friction will diminish, so the questioner's intuition is correct after all, and the optimum static left percentage will decrease.

In an earlier newsletter dealing with this question, I noted that if we do get to the point where left percentage is excessive for conditions, wedge or diagonal percentage adjustments will work backwards, and so will roll resistance adjustments. After that, a reader wrote in and said he had encountered this, with a go-kart on a very steeply banked dirt track.

Upon further discussion, it came to light that the kart had a much smaller tire on the left rear than on the right rear. This not only affected the optimum load distribution for the rear wheel pair, it also meant the kart had a lot of tire stagger. More load on the left rear increased the stagger-induced yaw moment on the kart, also causing more diagonal percentage to loosen the vehicle (add oversteer), contrary to what one might expect. This effect can easily occur in any car with a locked or partially locking rear end. This in turn affects our ability to infer whether left percentage is excessive, purely by noting how the car responds to adjustments.

I have also noted in earlier discussions that large left percentage makes a car tend to turn right under braking and turn left under power. This tightens the car (adds understeer) during entry and loosens it (adds oversteer) during exit. There are ways to counter this tendency with suspension design and tuning, but sometimes these are not legal, or the team doesn't understand them. In such cases, the car may well turn faster laps with less than optimal left percentage, even though it is slower in steady-state cornering.

These complexities can muddy the waters when tuning an actual car, but it is still fundamentally true that more left percentage is almost always better, provided we are able to work with the full package of consequences.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
July/August 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DYNAMICS OF THREE-WHEELERS

If you were building a three-wheeled vehicle, would you put some roll in the two-wheeled end, or not?

I am assuming that the questioner is asking whether there should be some roll compliance at the wide end, or whether the suspension should be essentially rigid in roll.

This is a reasonable question because, at least at low speed, a three-wheeler doesn't need suspension to keep all its wheels on the ground. Because any three points invariably lie in a common plane, a tricycle can trundle over very uneven ground at low speed with very little load change at the wheels, and without picking up a wheel, even if the entire chassis is rigid.

This fact led many of the earliest designers of motor vehicles to adopt tricycle layouts. Among these vehicles was the very first self-propelled off-rail vehicle, the Cugnot steam tractor of the 1760's (or 1771, if we go by the still-existing second model). Carl Benz's first gasoline buggy, in 1885, was also three-wheeled.

The trike layout was not universal, however, even in the early days. Benz's gasoline buggy was preceded by two models of internal-combustion, gasoline-fueled (or more accurately, benzene-fueled) cars built by Siegfried Marcus in 1865 and 1874. I don't know what the first of these looked like, but the second model, of which three were built, had four wheels. In 1879, George Selden applied for a US patent on the automobile. The model he submitted to the patent office had four wheels. It also had unitized body/frame construction and front wheel drive!

By the time automobiles became common, a wagon-style four-wheel layout had become the norm. This is not surprising, because this layout provides the best resistance to roll and pitch available within an envelope defined by a maximum length and width. Despite this, the tricycle layout has refused to die out completely.

The primary reason people have kept building trikes is economic. In most of the world, trikes are licensed and taxed as motorcycles rather than cars. Where cars are heavily taxed, this gives the trike a substantial price advantage. A secondary reason is that a trike can be made very light. Due to the aforementioned fact that a trike's tire contact patches always lie in a common plane, the vehicle's frame does not see the torsional loadings that a four-wheeler's does. Consequently, it can be built considerably lighter.

The tricycle layout brings problems, however. The main problem is poorer resistance to rollover. A trike can tip over by rolling about a line connecting the contact patches of the outer tire at the wide end and the single tire at the narrow end. For a given wheelbase and track, the vehicle's center of mass will unavoidably be closer to this line in plan view than it would be to a line connecting the front and rear outer-tire contact patches on a four-wheeler. Strictly speaking, the tipping motion we refer to here is not pure roll, but a combination of roll and pitch. Still, regardless of what we call the motion, the vehicle is limited by the easiest way it can tip.

The key to minimizing this problem is to put the c.g. toward the wide end as far as we can. If the single wheel is at the front, we need a rear-engine layout, similar to the VW-engined tricycles that are still fairly common in the US. If the single wheel is in the rear, we need a front engine, as in a Morgan trike. It is important to assure that the operator does not place any heavy cargo toward the narrow end.

It is best to drive the two wheels at the wide end, rather than the single wheel at the narrow end. Not only does this provide much better traction, but it further concentrates the masses at the wide end.

One problem we encounter when the c.g. is toward one end of the vehicle is that in hard longitudinal acceleration, the single wheel may lift, or become so lightly loaded as to impair directional stability. In a front-engine trike, the rear wheel will tend to lift in braking. In a rear-engine trike, the vehicle will tend to wheelstand under power. We can minimize this problem, and improve rollover stability, by making the wheelbase long, and by getting the c.g. as low as we can.

When choosing between the rear-engined and front-engined approaches, there is a safety advantage to the front-engine, front-drive layout. It has its best rollover resistance when decelerating, whereas the rear-engine, rear-drive layout is most likely to flip when the driver tries to lose speed upon entering a turn too fast. The front-engine, front-drive layout also provides much better crosswind stability.

Returning to the original question, what sort of characteristics should the suspension have at the wide end? First of all, it should not have large jacking forces. Either it should be an independent layout with a low roll center, or it should be a beam axle layout.

Particularly with an independent suspension, the wheel rate in roll needs to be substantial, but it should not approach infinity. If there is too much wheel rate in roll, the vehicle will se large roll accelerations, and large wheel load changes, when traversing one-wheel bumps at speed.

Barring a great increase in the popularity of trikes, we are unlikely to see a class for them in racing, except vintage racing where one does see Morgans and their contemporaries. We do, however, have racing for sidecar rigs. These are normally constrained by the rules to have two wheels in line, with the rear one driven, and a third wheel to one side. This is not the way to design a three-wheeler if we have a free hand, but it retains the connection to a roadgoing motorcycle with sidecar, and it provides thrilling, if dangerous, racing.

To optimize the sidecar layout, the wheels should again be spread as far in all directions as the rules will allow. The heavy side should, if possible, be toward the predominant turn direction. The c.g. should be away from the two-wheeled side, and fairly close to the single wheel in the fore-and-aft direction.

Even when all of this is carefully attended to, there will be no substitute for a good “monkey” or passenger, and the best possible helmet and leathers.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September/October 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

LOAD TRANSFER QUESTION

For argument's sake, let's assume the following:

1) *Total rear roll resistance is 600 lb/inch of rear suspension travel.*
2) *The static weight on each rear tire is 400 lb.*
3) *The car's rear suspension compresses 1 inch when cornering at the limit. Thus there is 600 lb of load transfer onto the outside rear tire.*

This car starts to corner and 400 lb of load is transferred from the inside rear tire to the outside rear tire. This leaves zero load on the inside rear tire. The car continues to speed up in the turn, reaching the limit of adhesion, and now there is 600 lb of load on the outside rear tire. Where does this additional 200 lb of load come from? Does it all come off the inside front tire? If you have rear anti-roll bar, it can actually push the inside tire up into the tire well. I guess this would be a negative load on the inside rear tire?

If I understand the question correctly, you are supposing that the outside rear suspension compresses one inch from static, implying that the tire gains 600 lb of transferred load, which would make its load 1000 lb.

If we are assuming that the car is in steady-state cornering, on a flat, smooth, unbanked turn, with no geometric anti-roll or pro-roll, and no aerodynamic downforce, you are describing an impossible case. Unless something adds load to the rear wheel pair beyond the static value, the outside wheel cannot have more load than the total for the wheel pair.

There can be no such thing as a negative tire load, unless the tire can somehow pull upward on the road surface. Short of creating a tread compound that is sticky beyond our usual conception, or nailing the tire to the road (either of which would make it very difficult for the car to attain enough speed to corner hard), that just can't happen.

It is also impossible for load to transfer from the front wheels to the rear wheels when the car is only accelerating laterally.

What will happen if the rear suspension reaches 100% load transfer, and then further lateral acceleration is applied to the car, is that the inside rear wheel will lift off the ground. The car will continue to roll, but without any further motion of the rear suspension. That implies that the rear ride height, measured from middle of the frame to ground, will increase as the wheel lifts.

The anti-roll bar will push the inside tire up into the wheel well only in the sense that it may prevent the inside suspension from reaching full droop – not in the sense of compressing the inside suspension beyond static position. The suspension's ride displacement from static will be zero. Its roll displacement will be two-thirds of an inch per wheel. The inside wheel will be off the ground, yet the suspension will be extended only 2/3" beyond static. The outside wheel will be compressed 2/3" from static. The average displacement of the two wheels, from static, will be zero.

We may say that in this situation, the rear suspension is <u>saturated</u> in terms of load transfer: it has absorbed all the load transfer that it can. Any further load transfer must be absorbed by the front suspension alone. This implies that the inside front wheel will lose load, but that load will not go to the outside rear; it will go to the outside front. The total load on the front wheels, and the total load on the rear wheels, cannot change.

Remember, though, that we made a number of simplifying assumptions here: purely lateral acceleration; no bumps; no banking; no geometric anti-roll or pro-roll; no aerodynamic downforce. In the real world, any combination of these might be present, meaning that we could very well have data acquisition traces showing an inch of compression from static on the outside rear.

To know how much added load we would need to get that added 1/3" of ride compression, we would need to know the rear suspension's wheel rate in ride as well as in roll. The required extra load wouldn't necessarily be 200 pounds. If the wheel rate in ride were 300 lb/in, we'd have that condition. (100 lb/wheel divided by 300 lb/in = 1/3 in/wheel)

If the only factor compressing the rear suspension is banking of the turn, and if the tires are racing slicks with a coefficient of friction around 1.30, we'd need about a 25 degree banking to generate 200 lb of extra load. A banking around 35 degrees would do this without the tires generating any cornering force.

If the turn is flat, and the only factor compressing the rear suspension is aerodynamic downforce, we'd need 200 lb of that at the rear axle if the wheel rate in ride is 300 lb/in. If the wheel rate in ride is less, these values decrease. If the wheel rate in ride is greater, the values increase.

The suspension geometry can generate a downward jacking force. This would be most likely in a lowered strut-style suspension, when most or all of the load is on the outside tire. In most cases, this will not be enough to compress the suspension a third of an inch unless the ride rate is very soft, but the effect could add to other effects to produce that much compression.

Forward acceleration will usually compress the rear suspension. In a front wheel drive car, it always will. We think of steady-state cornering as purely lateral acceleration, but actually there will be a car-longitudinal (x-axis) component, even at constant speed, because of the car's attitude angle or drift angle.

Since any or all of these effects can be present, it is entirely possible for the rear suspension to be compressed more than we would calculate for pure cornering on a flat surface. But something has to add ride compression for the condition described here to occur.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November/December 2005

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

OPTIMIZING ENGINE-OVER-DRIVE-WHEELS LAYOUTS

My question concerns taking a front-wheel-drive car and giving it the response of a rear-wheel-drive car.

It seems that for large sedans the trend in the market today is to build rear-wheel-drive platforms. The ideal seems to be to emulate the handling characteristics of the BMW and Mercedes RWD sedans. This is becoming a big deal as first Chrysler, and now Hyundai and the Chinese (new Cherry V8 of 4.8 liters), are heading down the RWD path.

Recently GM put the Pontiac GXP into production with a 327 cid V8 driving the front wheels. It can cut a 13.8-second ¼ mile and it's an automatic. It's faster than a standard Mustang! I discovered that the GXP was ready to produce ten years ago but GM couldn't afford to tool special parts for it then! Nor could they justify retooling the platform of the time in order to move the suspension hard points to more favorable locations. The recent new model introduction allowed the V8 option as they completely revised the platform for all model variants (Chev, Pontiac, Buick) and all engine options. It's an excellent car by all accounts (available only as a wrong-side-of-the-road model, so they can't export it – tsk tsk). [Questioner is in Australia.]

GM makes the point that front wheel drive is better in winter when traction is poor, for example in snow. Large parts of America (and many other countries) have difficult winters with snow and ice, so that's a valid argument.

This got me thinking about chassis response. I have a question about handling. Would it be possible to make a front-wheel-drive car behave like a RWD car (or feel like it had similar response) by applying some active rear wheel steer? I don't mean like Honda's mechanical system (where rear steer was a fixed function of front wheel steer regardless of speed, weight transfer, yaw rate, available traction, or chassis balance), but active. It wouldn't have to be much, perhaps just a degree or maybe even two at appropriate moments. What do you think?

This question implicitly raises a larger question: how can we get the most from an engine-over-drive-wheels two-wheel-drive car? To answer this, we need to look at the advantages and disadvantages of the whole engine-over-drive-wheels idea.

But first, to answer the question, yes it certainly is possible to use active rear wheel steering to point the rear wheels out of the turn in response to inputs from a steering position sensor, a throttle position sensor, a yaw acceleration sensor, wheel speed sensors, and any other inputs we can get a control computer to read and process. This would provide a yaw increase, or car rotation, in response to throttle application, as when throttle steering a rear-drive car. Getting exact reproduction of rear-drive behavior would probably be impossible, but getting the car to point the tail out when the driver applies power would be possible.

One reason that it would not be possible to duplicate rear-drive behavior is that we would have different behavior at the front wheels. At best, we'd be simulating the behavior of a throttle-steerable all-wheel-drive car, rather than a rear-wheel-drive car.

Another problem is that steering the rear wheels is not the same as powering them. Powering the rear wheels does not simply produce a slip angle increase; the effects are somewhat more complex. Light application of power actually plants the rear end, with front drive or rear drive. With rear drive, as throttle application increases, sooner or later we reach a crossover point, at which the use of the rear tires' friction envelope for propulsion starts to erode their lateral force capability faster than the normal force increase increases the lateral force capability. Then the tail starts to point out. With modest static rear percentage, this crossover point occurs relatively early. In tail-heavy rear-drive cars, there is a substantial range of throttle application in which power actually makes the tail stick rather than slide.

So the question arises: what sort of rear-drive response would we be trying to mimick? Rear-drive cars aren't all alike in their response to power application.

Regarding the Pontiac referred to: actually, the current Grand Prix GXP is the second V8, FWD car to bear the GXP designation. The 2004-2005 Pontiac Bonneville GXP has an engine closely related to the Cadillac Northstar V8, which has been used in FWD layouts branded as Cadillacs for many years now. This is a 4,565cc (279 cubic inch) engine with 4 cams and 4 valves per cylinder.

The current platform, using the 327 cid LS4 engine, is shared with the Chevrolet Impala and Monte Carlo. It is the first use of a pushrod "small block" family engine in a transverse front-drive layout.

This is little remembered now, but GM's first FWD cars were big V8-engined designs. The first was the Olds Toronado, in 1966. That was a really big, heavy car. As I recall, it weighed close to 5,000 pounds. A Cadillac variant followed, which had engines as large as 500 cubic inches. These cars had longitudinal or "north-south" engine mounting.

So putting large amounts of power through the front wheels can definitely be done. Power steering becomes a necessity, but power steering has become so commonplace, even on light cars with rear wheel drive, that we are generally not accepting an increase in cost or complexity if we include it.

With a transverse V8, the engine compartment packaging really gets tight. This makes it hard to find room for controls on either side of the car, to provide both left and right-hand-drive versions. It also becomes hard to find room for large-section structural rails beside the engine. The rails have to be narrowed to make room for the wheels to steer, and then the only way to make them adequately stiff is to make the walls thick, which adds weight. GM is using strut suspension, which helps some as the loads from the top of the struts can be fed into the cowl area. However, with struts and wide tires, the steering geometry doesn't work out very favorably. We have to accept either a large steering axis inclination, or a large scrub radius, or both.

Looking at the mechanics of putting ample power to the pavement, a FWD car doesn't look that much worse on dry pavement and street tires than a RWD, front-engined car. In fact, if both cars are relatively nose-heavy, the front-drive car may be better. For example, suppose both cars have a center of gravity 18 inches above the ground, and a wheelbase of 108 inches. At .50g forward acceleration, 1/12 of the car's weight, or about 8%, will transfer from the front to the rear.

If the front-drive car has 60% static front weight, it ends up with about 52% on the drive wheels. If the rear-drive car has 56% static front weight, it also ends up with about 52% on the drive wheels.

If the front-drive car has 62% static front weight and the rear-drive car has 58%, the front-drive car has 54% on the drive wheels at .50g, and the rear-drive one has 50%. That means the front-drive car will actually put power down better, assuming the coefficient of friction is such that the .50g figure is realistic. With 50% of the weight on the drive wheels dynamically, that would be a coefficient of friction around 1.00.

If we put slicks on the cars, things change. Assuming wheelspin still sets the limit, rear drive starts to look better, because more weight transfers rearward as the forward acceleration increases.

Conversely, if the surface is so slippery that it's a challenge to get the car moving at all, the front-drive car is clearly better. If the road slopes uphill, the front-drive car loses some of its advantage, but will generally still have an edge.

Note that we are comparing a fairly typical front-drive car to a decidedly nose-heavy rear-drive car. The more static rear percentage we have for the rear-drive car, the better rear drive looks.

Readers who like Porsches and Corvairs will be quick to point out that there is no law of nature that says an engine-over-drive-wheels car has to be front-engined. Assuming we are ruling out all-wheel-drive for reasons of cost, complexity, and weight, we really should be considering three layouts: front engine/front drive; front engine/rear drive; and rear engine/rear drive.

From the standpoint of propulsion, the rear-engine configuration has clear advantages over the other two layouts. The drive wheels can have 60% or more of the static weight, and this percentage increases when accelerating forward or climbing hills. Under hard braking, the front and rear brakes share the workload. In straight-line limit braking, the front wheels will do around 60% of the work, as opposed to 70% for a front-engined car with 50% or a bit less static rear weight, or 75-80% for a typical front-drive car.

The rear-engine layout can also be throttle-steered without any added contrivances. It can often do without power steering if desired. There is lots of room for the controls, making left and right-hand-drive versions relatively simple to accommodate.

The biggest problem with the two engine-over-drive-wheels layouts is the difficulty of achieving balanced cornering with four equal sized tires. This was more of a problem when tires were less reliable than they are now. Until recently, it was fairly important to have a single spare tire that would fit any corner of the car. Now, most cars have compact spares that are only suitable for limp-home use and don't match any of the regular tires anyway.

With four equal-size tires, we can build a FWD car that understeers, or a front-engine RWD car with balanced handling, or a rear-engined car that oversteers. The first two options are clearly preferable from a safety standpoint.

However, if we are willing to entertain the use of bigger tires at the heavy end of the car, some interesting possibilities open up. If we are careful with the body shape so that we maximize aerodynamic stability, and if we use a longish wheelbase, we can have a rear-engined sedan that will out-handle any front-engined design, and outperform it in the snow. This approach would be well suited to a transverse V8 powertrain. The problem of finding room for structural rails would be greatly eased since the wheels wouldn't have to steer. The car would be, just barely, mid-engined.

There would be some drawbacks. There could be two trunks, as in some mid-engined sports cars, but fold-down rear seats with a pass-through from the rear trunk would not be an option. The rear seat passengers would be subjected to more engine noise. Building a station wagon version would be problematic. Still, such a car could probably find an enthusiastic following among buyers who need to carry multiple passengers, yet give priority to performance.

Alternatively, we can also have a FWD car with little understeer. To make the most of this approach, we would want to have the engine well forward for at least 65% static front weight, bigger tires in front than in back, and the rear wheels way at the back of the car, as in a Citroen

DS21. This requires the designer and the buyer to cast aside accustomed notions of how a car should look, but it offers the promise of better handling and traction than existing two-wheel-drive cars, with a large, quiet, unobstructed passenger and luggage space in the rear.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

BIG TIRES ON THE FRONT

Last issue, I mentioned that there is a performance gain to be had in a front-wheel-drive car by making the car markedly nose-heavy and using larger tires in front than in back. This was in a rather lengthy response to a reader's question partly relating to the Pontiac Grand Prix GXP. I have recently noticed in magazine reports about this car that GM is in fact using larger tires on the front of the V8 versions of this car: 255 section on the front, 225 on the rear.

TIRES IN THE SNOW

I have heard two schools of thought on tire pressure for winter driving. This applies to driving on snow-covered roads. The first is that tires should be kept at the upper end of the manufacturer's specifications to help in cutting through the snow. The thinking is that the contact patch is smaller, hence more weight per square inch, as well as less sidewall deflection – which may decrease the potential for hydroplaning on the snow. The other is that running the pressure at the lower end allows for better bite for the tread in the snow, and more stability. What do you think?

I'm with the low-pressure school.

It's said that one measurement is worth a thousand expert opinions. Really, you'd think this might be measurable. Surely somebody has tried measuring, say, how steep a hill a vehicle can climb, at various tire pressures. I would be willing to defer to any actual measurement that contradicts my expert opinion.

That said, I offer my expert opinion, and that expert opinion is based on a lot of time spent in Wisconsin, where there are long, snowy winters.

First, tire pressure effects in snow are surprisingly subtle compared to other variables, and effects due to other variables are surprisingly large. That explains why there is controversy about inflation pressures, even though people have been driving cars through more than a hundred winters now.

One reason tire testing was moved indoors, with rollers or belts substituting for pavement, is that even on hard, dry pavement, weather, surface contamination, pavement temperature, pavement age, and other factors make enough difference that small variations in tire performance are hard to measure repeatably.

When we're dealing with snow, we have similar variability, exaggerated at least tenfold. Snow and ice come in dozens of different varieties and depths, and all of these have properties that are highly temperature-sensitive. Compared to snow, pavement is simple and consistent.

The explanation I have heard from sources that advocate high tire pressures is that the tire needs to penetrate the snow and get to the pavement, where it can find traction. But clearly, in most snow conditions that never happens, at any inflation pressure. If it did, we'd see bare pavement where the tire passed.

Instead, we see compacted snow, with an imprint of the tire tread. Or at least we see that if the tire is rolling, and not spinning or sliding – and if we are dealing with snow that has not already been compacted. We also usually see a small area alongside the track where the snow appears slightly raised, apparently having been pushed out of the way, perhaps just by the sidewall. This is not a lot of snow, however. Most of the snow stays put horizontally, and gets compacted vertically.

The tire evidently gets traction by packing the snow into a relatively solid form, and simultaneously interlocking with it. To break traction, the compacted snow projections residing in the tread grooves must be sheared off, and the layer of snow lying under the tread blocks must also fail in some manner.

The failure of the snow in the grooves is easily visualized as simple breakage. The failure of the snow under the tread blocks is a bit harder to visualize. It appears that the snow under the tread blocks contributes more to traction than one would imagine, because the tire's grip is greatly improved by siping the tread blocks. It also helps to roughen the surface of the tread blocks.

I do not claim to perfectly understand the mechanics of structural failure of snow in a tire contact patch, but I do know that it is normally a combination of breakage and melting. Ice (and snow is ice crystals) can be melted by mechanical pressure – or, stating it a bit differently, the melting point of ice is lowered by mechanical stress, either compressive stress or shear stress. Anyplace that the snow or ice liquefies, its mechanical strength disappears, and it turns into a lubricant. The closer the ice or snow is to its melting point, the less mechanical stress is required to turn it to liquid.

So, when we compact snow, we make it stronger, but only up to the point where we start to get localized melting. The unit loading required to reach this point depends on how cold the snow or ice is. Moreover, short of the point where we encounter melting by compression alone, we see an increased likelihood of melting by the combination of compression and shear. In other words, as unit loading increases, we gain hardness but lose melt resistance. The hardness gain is fairly independent of temperature. The melt resistance loss is heavily affected by temperature, or at least its importance is.

From this, we might logically expect that the ideal contact patch size would be smaller in really cold weather than when we're near thaw temperature.

I suspect that this is academic, however. I think the optimum contact patch size is far bigger than we can ever get with a tire. Consider the transportation devices that people have devised specifically for snow: snow cats, snowmobiles, snowshoes, cross-country skis. All of these operate by compacting the snow minimally, over a large area, and then trying to get maximum purchase on that large interface. For best performance on snow, or any other soft surface, we really want a belt or track, not a tire.

It would seem to follow that the more we can get a tire to act like a track, the better it should work. That would suggest a radial tire, at low pressure.

Note that it does not necessarily follow that we want a wide tire. It is generally agreed that for most winter conditions, a tire should be narrow. I think, based on the reasoning above, there will be winter conditions where a wide tire may be preferable. These may include bare ice and hard-packed snow, probably even shallow soft snow. But in snow of significant depth, narrow tires are better.

The reason for this doesn't have to do with an increase in traction when the tire is narrow, as such. Rather, it has to do with the force required to move the tires, which is less when the tires are narrow.

As the tire rolls forward, it is resisted by the snow in front of it. To advance, the tire must, in effect, climb a ramp of snow. The ramp of snow is not strong enough to support the tire, and it is continually collapsing under the weight of the car. The amount of collapse is fairly similar regardless of the width of the tire; for any practical tire size, we will compact the snow to a pretty solid state, no matter what. Yet the snow has substantial resistance to this compaction, and this translates to a resistance to the wheel's forward motion. The taller and wider the mass of snow we must compact, the greater the resistance to motion. The height of the snow we must compact depends on the snow's depth. The width we must compact depends on the width of the tire.

It would also seem that a narrow tire should provide more directional control, since it is better shaped to act like a blade or rudder.

From this reasoning, we might expect that the ideal tire for deep snow would resemble a bicycle tire. Such a tire would be easy to push along, and should have good directional stability.

However, it doesn't quite work that way with really narrow tires, as anybody who has tried riding a bicycle in deep snow will attest. The problem is that the ramp of compacted snow that the tire rides on is so narrow that the tire is forever sliding off the side of it into the soft snow alongside. As soon as the tire moves forward again, another narrow compacted ramp is formed beneath it, and again it slides off one side or the other – no predicting which side. The result is that the tire absolutely will not run straight.

So there is such a thing as too narrow. The tire needs to be wide enough to sit on top of the compacted ramp it is making for itself. A square-shouldered profile, or one with concave shoulders that compact a sort of retaining berm along the side of the main compacted ramp, also can be expected to help.

Returning to the question of inflation pressure, this also affects resistance to forward motion. And even this relationship is not as straightforward as one might think. Based on our experience with tires on pavement, on a smooth, hard surface, the higher the inflation pressure, the easier the tire rolls, at least within practical limits.

But on a rough surface, a softer pressure can actually roll more easily. For this to be so, the surface must have roughness as opposed to waviness: the ups and downs must come fairly close together. The tire rolls easier because it can yield to the bumps rather than having to climb over them. This was realized very early in the history of the pneumatic tire. John Dunlop immediately noticed that his new pneumatic tire would roll further across his bumpy back yard than a solid tire.

This has relevance to driving in snow because often the situation that gets us stuck is one where one or more wheels are in a fairly modest-sized depression, and we have to move the tire over the lip of the depression with the meager traction available. In at least some such situations, soft inflation will make getting over that lip easier.

It will be apparent that I am writing here from a mixture of practical experience and inference. I invite readers with further experience, or contradictory experience, to comment.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL AXIS INCLINATION

What is the influence of a roll axis inclination biased to the front suspension – meaning a front roll center always closer to the ground than the rear? At least in passenger cars, the roll axis is always inclined to the front except in some special cases, for example the BMW Series 1 which is reported by BMW to have the roll axis parallel to the ground.

I supposed I had an explanation, but after reading Race Car Vehicle Dynamics by Milliken my potential explanation has flown away. My explanation was based on the idea that the more the roll axis is inclined toward the front, the more load transfer there will be at the front axle, and the more understeer the vehicle will have.

But I have put into an Excel spreadsheet the formulation from Milliken and I find to my surprise that the higher the front roll center, the greater the load transfer at that end – which works against my intuition.

Can you explain this?

Short answer: higher roll center at the front implies more geometric roll resistance at the front, hence more load transfer at the front, other things being equal. So the typical nose-down roll axis inclination does not increase front load transfer.

There are cars that have a nose-up roll axis. They are all rear-engined. Probably the most extreme example is the Hillman Imp, which had a front roll center near hub height and a rear roll center near ground level.

Like many things, the subject of roll resistance and load transfer is fairly simple once you understand it, but will drive you crazy until you get to that point.

When discussing this subject, I am always quick to plug my video, *Minding Your Anti*, which covers the subject at length. It costs US$50.00, shipping included, payable by check or money order to me at 155 Wankel Dr., Kannapolis, NC 28083-8200, USA.

In steady-state cornering (constant speed, on a constant radius), on an unbanked road surface, the total load transfer from the inside wheels to the outside wheels depends entirely on the height of the whole vehicle's center of mass (center of gravity, or c.g.) and the track width at the c.g.

Suspension design and tuning have almost no effect on the magnitude of the total load transfer. What we mainly do with suspension design and tuning is control the distribution of that total, between the front and rear wheel pairs.

We customarily consider the car to be a rigid object, supported by a single compliant structure at each end. The sprung structure is the rigid object; the front and rear suspension systems are the compliant structures.

As an analogy, imagine that you and a friend are carrying a sailboard, as used for windsurfing, along the beach. Each of you is carrying one end of the sailboard. The sail is up, and there is a breeze blowing. The force of the wind on the sail tries to overturn the sailboard.

The overturning force depends entirely on the design of the sailboard and the amount of wind. The total counterforce that you and your friend together need to exert to balance this does not depend on you and your friend. However, the amount of counterforce that you individually need to exert depends on the amount exerted by your friend, and the amount of counterforce he has to exert depends on you.

You and your friend are like the front and rear suspension systems. The sailboard is like the sprung mass.

There are portions of the load transfer that come from the unsprung components, and there are portions that come from the dampers if the car is rolling upon corner entry or de-rolling on exit. However, for simplicity in answering the present question let's look just at the components of the load transfer that come from the inertia force (centrifugal force) of the sprung mass acting through the suspension, in steady-state cornering. There are only two such components: elastic load transfer and geometric load transfer. Elastic load transfer comes from elastic roll resistance: the roll resistance supplied by the springs and anti-roll bars. Geometric load transfer comes from the properties of the structural components attaching the wheels to the sprung mass, which can be arranged to generate forces opposing roll, or geometric roll resistance.

With independent suspension, these two components influence each other more than is commonly recognized. The load distribution on an independently suspended wheel pair affects how much geometric roll resistance the wheel pair has, for any given suspension geometry. To illustrate with an extreme case, if the inside wheel is off the ground, the geometry of its suspension linkage is irrelevant and only the geometry of the outside wheel has any effect on the car. My video deals with these effects in detail. For simplicity, I will ignore them here, but I do want note in passing that they exist.

When we speak of roll center height, we are speaking of an imaginary point whose height represents the amount of geometric roll resistance for the front or rear wheel pair. If this point is assigned properly, we can closely approximate the geometric load transfer at one end of the car as: roll center height times sprung mass centrifugal force at that end of the car, divided by track width at that end of the car.

When the suspension is symmetrical, the point you generally see in the chassis books – the force line intersection – is a good approximation. When the suspension is not symmetrical, using the force line intersection as the roll center can lead to major mis-predictions of car behavior. Sometimes the force lines may be parallel, in which case there is no intersection.

We may define a line connecting the front and rear roll centers, called the roll axis. The car doesn't really roll about this line, but as a crude approximation we can reasonably think of it as doing so.

If we raise the roll axis at both ends, the geometric roll resistance is greater at both ends. If we raise one end of the roll axis and lower the other, leaving its height at the c.g. unchanged, the total geometric roll resistance is unchanged, but we increase the geometric roll resistance at one end and lower it at the other. The elastic elements – the springs and anti-roll bars – are not affected by this.

So the end where we lowered the roll center has less geometric load transfer and the same elastic load transfer as before – hence less load transfer overall. This will make that tire pair grip better, because they will be sharing the work more equally. At the opposite end, the elastic component will likewise be unchanged, but the geometric component will be increased – hence more load transfer overall.

Okay, so if we want understeer for most drivers, why have a nose-down roll axis? There are a number of explanations.

The most obvious explanation is that when the car has independent suspension in front and a beam axle in back, we don't have much choice. Independent suspensions with roll centers much above four inches generally jack excessively. Front suspensions with high roll centers generate lateral contact patch motion over bumps, which creates kick at the steering wheel. It is possible to build a beam axle suspension with a roll center below any component of the suspension, but the linkage required is somewhat complex. Consequently, beam axles on cars with enough ground clearance to be practical on the street generally have roll centers at least six inches high, and usually at least ten inches. Of course, with independent rear suspension, the roll center is usually much lower, but most often still a bit above the front one.

The next most obvious reason is that passenger cars are generally too nose-heavy to have balanced handling, and the front suspension doesn't control camber when cornering nearly as well as the rear suspension. Consequently, we need to kill understeer, not increase it.

A somewhat less obvious reason has to do with driver-perceived car behavior in abrupt transient maneuvers, such as the lane-change test commonly used in passenger car testing. With a nose-down roll axis, there is a small yaw component with roll. The nose points out of the turn slightly, relative to the four contact patches. This makes the car feel steady to the driver, rather than twitchy.

Another reason sometimes cited is that when a car is abruptly steered into a turn, the geometric component of the load transfer is the first to act on the car. If this component is greater at the rear, we will momentarily have less understeer and the car will turn in more responsively. Note that this explanation is somewhat at odds with the one immediately preceding it.

There are somewhat logical variations on both of these two explanations. We could say that if the main mass of the car is yawing out of the turn relative to the four contact patches, that steers the contact patches into the turn, or steers the rear wheels out of the turn, momentarily adding oversteer!

Some people also believe that tire load sensitivity momentarily works backwards until the tires start heating. I personally don't believe this, but if so it means that if there is initially more rear load transfer, that adds understeer rather than oversteer, and makes the car feel stable.

Isn't this fun? If it weren't for vehicle dynamics, I'd have to do something sane for a living.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DIFF DIFFERENCES

Which differential would be best for road racing, in my 300hp Porsche? The question arises because a supplier suggests that a plate-type limited-slip is better suited to road racing than a worm gear style torque bias diff. I have used both and found Iliked the torque bias diff. I thought it was a better design, from what I read. The supplier states that the lsd will be better in corner entry and exit. What is your opinion?

Clearly, both worm gear and clutch pack differentials have their adherents, and both are used successfully in racing. You say you have experience with both types, and have already formed a preference. The most obvious answer would be that you've already answered your own question, and don't need advice.

However, the situation is actually a bit murkier, because the behavior of both types of differential can vary according to design and tuning details. Both types are similar in that they generate a locking torque in response to the total torque being transmitted. In both types, the locking torque depends on pressure angles. In a ZF-style clutch pack design, the angles are those of the ramps on the spider shaft and the housing halves. In a worm gear design, it's primarily the helix angle on the gear teeth, and secondarily the pressure angle of the tooth profile. Lubricant choice also influences behavior.

Consequently, all clutch pack diffs don't act alike and neither do all worm gear diffs. A lot depends on how a specific example is tuned.

That said, the clutch pack design probably offers a greater range of tuning options, and probably greater wear resistance. With the worm gear designs, we are trying to make gear teeth act as a friction device. Clutch discs are designed to be a friction device; gear teeth can be made to act as a friction device, but they are less comfortable in that role.

This affects the ability of the differential to maintain consistent properties over time, and its longevity.

The pressure angles determine how rapidly locking torque builds as transmitted torque increases. The preload in the diff determines how much locking torque there is when no torque is being transmitted. A clutch pack is easily preloaded, and it maintains its preload relatively well, especially if the preload is applied by springs or some other compliant system such as dished clutch plates. Worm gears can also be preloaded, but because they are not very compliant, the preload rapidly goes away as the teeth wear.

One limitation in worm gears is that the pressure angle is generally the same for forward torque and rearward torque (as when engine braking, or when transmitting brake torque from a single rear brake, as seen in FSAE cars). In a clutch pack diff, we can use different ramp angles for power and decel.

Another peculiarity of worm gear designs is that because power and decel apply force to opposite sides of the gear teeth, preload doesn't have identical effects in both directions. If we preload the gears in the direction they're loaded under power, what happens under decel is that we have diminishing friction with increasing reverse torque, until the preload is overcome, at which point locking torque is zero. As reverse torque increases beyond that point, locking torque builds again. With a clutch pack, preload has similar effect in both drive and decel modes.

This does mean that we can make a worm gear diff act different in drive and decel, but not in a manner that's independent of preload.

One interesting, though uncommon, trick we can use in a worm gear diff is to use plain thrust washers to absorb the thrust of the worm gears in one direction, and needle thrust bearings to absorb the forces in the other direction. This can afford us some limited measure of difference in friction depending on torque direction. Last year's North Carolina State University FSAE car had a diff like this.

It will be clear, however, that using these tricks is not nearly as straightforward as varying the ramp angles in a clutch pack diff.

Finally, neither option is ideal, because neither is speed-sensitive. Both clutch pack and worm gear diffs rely on Coulomb friction, which is largely dependent on normal force and not speed. We would rather have the locking torque vary with the speed difference between the wheels, either entirely or at least in part. This argues for either a pure viscous limited-slip, or a design that uses a pump, driven by relative ouput shaft rotation, to load a clutch pack, or a design that combines viscous effects with a clutch pack.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
April 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SPRINGS, BARS, AND LOAD TRANSFER

I've been reading around and am having a hard time comprehending a spring's effects on load "change" and how they effect understeer/oversteer balance as stated in your April/May/June 2003 newsletter.

<u>What I know (or think I know):</u>

Longitudinal Load Transfer = acceleration x [(weight x cg height)/wheelbase]

Lateral Load Transfer = (Lateral Acceleration x weight x cg) / Track width

So, by looking at the equation, spring rates are not part of the formula, therefore they play no role in the amount of load transfer. However, I do know that sway bars DO affect load transfer by unloading the inside wheel and loading the outside wheel. In Carroll Smith's "Tune to Win," he states:

"The greater the resistance of the springs, the less roll will result - but there will be no significant effect on the amount of lateral load transfer because the roll couple has not been changed and there is no physical connection between the springs on opposite sides of the car. The same cannot be said of the resistance of the anti-roll bars. In this case, because the bar is a direct physical connection between the outside wheel and inside wheel, increasing stiffness of the anti-roll bar will both decrease roll angle and increase lateral load transfer."

By looking at the above information, we would assume springs are merely there to control body roll which affects camber, toe, etc.

And from your article, I gather that the amount of load transfer between the two front wheels vs the two rear wheels indicates how much oversteer/understeer a car will have. However, the above equations show that spring rates don't affect load transfer.

<u>What I don't understand:</u>
I realize that using higher spring rates in rear vs. spring rates in front will cause a car to oversteer - given a simplified car with equal motion ratios front/rear and without using sway bars. In the real world, it proves true. However, I'm having a hard time comprehending how spring rates can affect the balance of the car when they play no role in load transfer? What do you mean when you say stiffer springs cause more "load change"?

The questioner is quite right that springs do affect car balance, and that this could not be so if they had no effect on load transfer.

With all due respect for Carroll Smith's memory and legacy, significant portions of the passage the questioner cites here are simply incorrect. Springs and anti-roll bars both affect load transfer – or more properly, load transfer distribution – in cornering or lateral acceleration, and they do so in essentially the same way. Indeed, anti-roll bars are springs, and the ride springs are connected side-to-side, by the car's frame.

The only difference between an anti-roll bar and a ride spring is that anti-roll bars act only in roll and warp, whereas ride springs act in all four modes of suspension movement: roll, pitch, heave, and warp. In the modes in which it is active, the bar is just another spring.

The equations cited are for total load transfer, in lateral or longitudinal acceleration. Taken as such, they are correct. The springs and the bars do not affect these total quantities, or at least not very much. However, the springs and bars affect the **apportionment** of that total, between the front and rear wheel pairs in lateral acceleration, or between the right and left wheel pairs in longitudinal acceleration.

Let's consider lateral acceleration first. The roll-resisting moments produced by the springs and bars are called the elastic component of roll resistance, and they produce the elastic component of load transfer at the tires. There is also a geometric component, at each end of the car, which comes from the forces in the (comparatively) rigid suspension components. These two components make up the total roll resistance for the front or rear suspension.

For purely lateral acceleration, in steady-state cornering, assuming equal track width at both ends, the total sprung mass load transfer at the front or rear is the total sprung mass load transfer for the whole car, times the percentage of the total roll resistance at that end. For example, if the front suspension's combined elastic and geometric roll resistance is twice as great as the rear suspension's combined elastic and geometric roll resistance, the front end then has 2/3 of the total roll resistance, and it will see 2/3 of the total sprung mass load transfer.

Increasing the front roll resistance or decreasing the rear will increase front sprung mass load transfer, and decrease the rear. Increasing front and rear roll resistance together, maintaining 2/3 of the total at the front, will result in no significant change in the wheel loads. It doesn't matter whether we change roll center heights (the measure of geometric roll resistance), springs, bars, or all three, as long as the totals front and rear are in the same ratio as before we made our changes.

In longitudinal acceleration, the same principles apply, except that the effects of the bars are negligible since we are dealing with pitch rather than roll, and we are concerned with the relative pitch resistances of the right and left wheel pairs, rather than the relative roll resistances of the front and rear wheel pairs.

In either case, the elastic component of the wheel pair load changes depends on the rate of both springs in the pair, and increasing the rate of either spring in the pair increases the average or total for the pair.

So, for the rear suspension in roll for example, stiffening either the inside or the outside spring adds roll resistance. On the outside wheel, a stiffer spring increases the **rate of force increase** with respect to suspension compression: more pounds of load change per inch of suspension movement. On the inside wheel, a stiffer spring increases the **rate of force decrease** with respect to suspension extension: again, more pounds of load change per inch of suspension movement.

Either way, we have more pounds-feet of roll-resisting moment per degree of sprung mass roll.

So in any situation that decreases the load on a spring, a stiffer spring gives less load on the wheel at a given suspension movement: more load change, implying less load, compared to same displacement with a softer spring.

More knowledgeable readers will note that this discussion is simplified somewhat compared to real life. Springs can somewhat affect total load transfer, chiefly because they can create ride height changes with roll and pitch. For example, if we put stiff springs in the rear and soft springs in the front, and the car has little anti-squat, then under forward acceleration the rear suspension will compress little and the front suspension will extend much. The car will therefore be sitting higher off the ground, c.g. height will be greater, and therefore rearward load transfer will be greater.

Thus, it is not strictly correct to say that springs do not affect overall total load transfer. However, for all but extreme cases, we can treat such effects as negligible.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

REAR PERCENTAGE OR POLAR MOMENT?

I am wondering about the relative merits of Front / Rear Weight Distribution versus Polar Moment of Inertia. I road-race a Mustang in a class that allows extensive engine and suspension modifications to late-model muscle cars, but the rules are strict enough to prevent extensive body, floorpan, and frame modifications. As a result of the rules, and the large V8s, the cars are inherently heavy in the front end – usually to the tune of 51-55%, depending on platform, fuel load, etc. So when deciding on the placement of heavy components (such as batteries, fuel cells, ballast, etc.) I'm wondering what the ideal position is. Obviously lower is always better, but how far back is ideal? You could place the masses as far back as possible to improve F/R balance, but at the expense of Polar Moment. You could put the mass in the center of the car, minimizing polar moment but failing to take advantage of an opportunity to improve the car's balance. Or you could put it somewhere in between.

Instinctively, my engineering training tells me that the two things have different functions. F/R distribution should make the car generally behave in a more balanced fashion, and certainly more weight on the back tires will help corner exit traction – which is a key in this type of car. Keeping polar moment low should help keep the car rotationally responsive, and more willing to respond to corrections if the tail steps out.

Several years ago I ran a simple analysis, and found that the average cornering speed of the car seemed to be a determining factor between these two trade-offs. This was not an in-depth analysis. Basically I broke cornering down into the energy it took to rotate the car (start it rotating and stop it rotating) and the energy it took to redirect the velocity vector, and compared the two. It seemed to me that the faster the mid-corner speed, the more the velocity vector change became the dominant force requirement, and therefore we might assume the F/R balance would be more important. My analysis said that, in the kind of racing I do, F/R distribution would be far more important, but I'm surprised at the number of competitors, in amateur and pro competition, who seem to disagree. Of course, my analysis was simplified to a great extent, with numerous assumptions being made that are not necessarily always valid.

So I ask – how far back should I place that ballast, or those movable components? On an average road course, where mid-corner speeds probably average from 40mph to about 100mph or so, which would you deem more important? Would you place that ballast in the back bumper, or under the floorpan or inside the inside frame rail?

And as a secondary question – what would you consider the ideal weight distribution for a car of this type. Obviously, if I can achieve whatever I decide is "ideal", then I want to work to reduce polar moment, but what is that "ideal" number – in your estimation of this type of race car? I've heard arguments for 50/50 and arguments for as much as about 58% rear weight. I tend to think the answer is closer to 55%, but I'm very curious as to your input on this matter as well.

Taking the last part first, the ideal rear percentage, or the desirable range, depends on your tire rules and the nature of the course. To a lesser extent, it is also influenced by the level of available grip, and by aerodynamics.

Suppose we have no tire rules, ample power, lots of freedom with our wings and other aero devices, and a track with serious straights and tight turns. Suppose also that the rules require rear wheel drive. In such a case, we might want as much as 65% rear, or even a bit more, and rear tires about twice the size of the fronts. We would want more than 2/3 of the downforce at the rear. We would want ample area in rear wing sideplates and/or tail fins, for yaw stability.

The reasoning here is that the car will spend a lot of its time accelerating longitudinally: accelerating forward out of the turns and down the straights, and accelerating rearward (braking) at the end of the straights. Propulsion comes only from the rear wheels, so we want as much of our weight on the rear as possible to maximize forward acceleration. Braking is shared by all four wheels, and if the car has ample braking capability, the fronts will do more than half of the braking even if they only carry a third of the car's weight at rest. This is due to the large amount of dynamic load transfer forward during hard braking. So for both forward and rearward acceleration, we want as much rear percentage as possible, at least within practical limits for a road racing car.

This would be so even if we are constrained to equal size tires front and rear. If our event is, say, a zero to 100 to zero competition (accelerate from a standstill to 100mph, then brake to a standstill, all in a straight line), we'd want to build the car so it just barely picks the front wheels off the ground at launch, or almost does so. Note that the rear percentage needed for this varies dramatically depending on the wheelbase, the c.g. height, and the amount of grip available. With a production-based car – relatively short wheelbase and high c.g. – and good drag slicks, the ideal rear percentage may be less than 50%. The grip of modern drag tires is so good that a 50/50 car can actually be limited by wheelstand rather than wheelspin. With a high c.g. and a short wheelbase, the ideal rear percentage varies dramatically with the grip level.

If we have more chassis design freedom, we'd rather have a dragster than a pro stock: long wheelbase, low c.g., lots of static rear percentage. With such a layout, rearward load transfer is less, and consequently dynamic rear percentage varies less with grip level. A single layout and setup will therefore be more nearly optimal over a wider range of conditions. With street tires or road racing tires, the ideal rear percentage will be considerably over 50%. As long as the car doesn't have to turn, we'd want to build it like a dragster, even if the front and rear tires have to be the same size.

Of course, the questioner here doesn't have such design freedom, nor is he running on a track with no turns. I'm just examining extreme cases, to illustrate some points.

The opposite extreme case would be a skidpad competition: the car just has to corner at the highest possible constant speed. Accelerations now are almost purely lateral. We might suppose that the ideal design here would have equal-size front and rear tires, and 50% static rear weight. That is indeed close to correct, although if we can use unequal tire sizes, we can get as good cornering from a tail-heavy car as from a 50/50 one. If there are no limits on tire size, two practical considerations will limit us: camber control and steering geometry. We can control camber about equally well at both ends of the car, but there is a case for using smaller tires on the front from the standpoint of steering dynamics. If we do that, we want the car tail-heavy, roughly in proportion to tire size.

Oddly enough, the radius of the skidpad affects the optimum rear percentage. This is a bit counter-intuitive, but it's true. Really small skidpads (or really tight turns) call for a more tail-heavy car than the tire sizes would suggest, and really big skidpads (or really big turns) call for a more nose-heavy car. The reasons for this are of different natures for the two cases.

On a really small skidpad, such as the 50-foot diameter one used in Formula SAE competition, the front wheels track noticeably outside the rears, even when the tires are sliding. Consequently, the drag forces of the front tires, especially the heavily loaded outside front, act at a larger radius than the radius the c.g. is following, and the propulsion forces from the rear tires act at a smaller radius than the c.g. is following. This creates a yaw moment out of the turn, and tightens the car (adds understeer). Additionally, in some cases there may be effects from a limited-slip differential or a locked rear that create understeer in small-radius turns.

When the turn radius is really large, the car will need to transmit substantial amounts of power through the rear tires just to maintain constant speed. On really fast turns, the car may actually be near full power, and not gaining any speed at all. The rear tires are transmitting hundreds of horsepower, just to overcome aerodynamic drag and the induced drag from the front tires as they run at a slip angle.

This means that the rear tires are using a substantial portion of their performance envelope, or traction circle, to propel the car, so they have less of their capability available to generate lateral force. The car is consequently subject to power oversteer. The best way to counter this is to have a disproportionate amount of aerodynamic downforce at the rear of the car. However, in many classes, including stock cars running on high-speed ovals, the rules may not allow the aerodynamic devices needed to achieve this. It then becomes desirable to make the car a bit nose-heavy to add understeer and counter the power oversteer.

I am not trying to confuse the issue. I am merely pointing out that, as the questioner has already come to appreciate, we cannot state categorically that a particular rear percentage is ideal. It depends on other factors.

That said, the questioner appears to be correct that in his class, more rear percentage is better, within the limits imposed by the rules, without going above minimum legal weight. And there is indeed a tradeoff in such a situation between getting good rear percentage and reducing yaw inertia (polar moment of inertia in yaw, commonly "polar moment" for short). I agree with the questioner's tentative conclusion that it is better to go after rear percentage and forget yaw inertia, for road course applications.

It is interesting to consider what it might take to make us go the other way, and put the ballast closer to the middle. I think we might do that if we were autocrossing, particularly if a large part of the course was made up of a long, constant-speed slalom, and if the course had no significant straights, so that the car was continually cornering and continually changing direction, and never had to spend a lot of time accelerating longitudinally – in other words, if there was a lot of yaw acceleration and relatively little longitudinal acceleration.

But no road course is like that. Almost all of them have serious straightaways, and few really abrupt transitions. A car with large yaw inertia will tend to understeer into the slower turns and oversteer coming out, but to some extent the driver can overcome the entry understeer with trailbraking, and the exit oversteer by using an "out fast" line and judicious throttle management.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
June 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

EFFECTS OF SPRING RATES AND DAMPER SETTINGS DURING CORNERING

I would be interested to hear your comments about the effects of relative spring rates (front/rear) and damper settings on steady-state cornering characteristics. I do not race, but I am interested in the effects of suspension settings on the relative under/over-steer characteristics.

I presume that (all other things being equal) any increase in spring rate at the front will increase the apparent roll stiffness and therefore make the car trend towards an increase in understeer (which I presume is the same as an increase in static directional stability) – and an increase in the rear will cause the reverse. And an increase in damper rate at one end would seem to have a similar dynamic effect during turn-in, while the load transfer is actually taking place.

I own a Morgan – which as you are aware does not have any anti-roll bars – and I am about to change the front springs to a higher rate (140 lb/in vs. 105 lb/in) and I presume that this is going to cause some change in behavior. It would be interesting to know what order of magnitude of change I could expect, though, and whether it would be worth experimenting with the (adjustable) damper rates to try to modify turn-in behavior.

If I had more extensive experience with Morgans, perhaps I could predict the change with more confidence. As things stand, I can say that you have correctly understood the effect that spring rates have on steady-state handling balance in most cases: stiffen one end, and you get more load transfer at that end and less at the other, so you reduce grip at the end you stiffened and add grip at the other end.

However, in certain cases we can add roll resistance at the front and <u>reduce</u> understeer! This is most often seen in cars with beam axles in back, and independent suspension with poor camber recovery in roll in front – most commonly small rear-drive sedans that roll a lot and have lowered MacPherson strut suspension in front. What's going on in these cases is that although the front tires are less equally loaded, the reduction in roll improves their camber so much that the camber improvement more than makes up for the more unequal loading. At the rear, the beam axle gives 100% camber recovery at any roll angle, so rear camber is unaffected.

The Morgan is a similar case in some respects. It rolls a lot less than a tall, narrow sedan, but it has no camber recovery in roll at all with that sliding-pillar front suspension. The front wheels lean the same amount as the body. So when you add roll resistance at the front, you are hurting the load distribution at the front but helping the camber. At the rear, you are helping the load distribution and leaving the camber largely unchanged.

Also complicating prediction in the case of the Morgan is that the frame is unusually flexible in torsion. That mutes the effect of relative roll stiffness changes.

Actually, all cars with independent suspension in front and beam axles in back have poor camber recovery in front compared to the rear, so they all are subject to the same conflicting effects when we add front roll stiffness. Interestingly, when we change rear roll resistance the effects on front load distribution and front camber are additive rather than subtractive, and we can predict the effect on car behavior with much better certainty. Reducing rear roll stiffness will hurt both camber and load distribution at the front, while helping load distribution and not affecting camber at the rear. We know that will add understeer. Conversely, adding rear roll stiffness will help both camber and load distribution at the front, while hurting load distribution and not affecting camber at the rear. We know that will add oversteer.

As to the effect of shocks, yes stiffening the fronts will add understeer during entry and stiffening the rears will add oversteer, provided that the road surface is smooth. This effect requires that the car have a roll velocity outward, and that this be the main source of suspension movement. When the car is cornering steady-state on a smooth surface, the roll velocity should be zero, the suspension should have displacement from static but not velocity, and shocks shouldn't matter. During exit, the car has a roll velocity inward (it's de-rolling). In this situation, the effect of the shocks reverses. Stiffening the fronts adds oversteer; stiffening the rears adds understeer.

So to add understeer or oversteer overall, we use the relative stiffness of the front and rear springs (and/or bars, if present). To change entry and exit properties in opposite directions, we use the relative stiffness of the front and rear shocks (remember, only on smooth surfaces).

I sometimes refer to damper forces as creating frictional anti-roll or pro-roll (anti-de-roll). Even forces generated purely by a liquid may be termed a form of friction if they act in opposition to motion. Speaking of friction, I believe I have observed a phenomenon watching Morgans run that may be of interest here. I think that these cars can easily experience excessive friction in the sliding pillar mechanism when subjected to the forces modern racing tires can generate. This causes understeer until the driver finally gets the car rotating, gets on the power, and starts unwinding the steering. Then the car snaps into oversteer as the front end suddenly frees up and can roll. I therefore always tell people running these cars to keep the pillars in good condition and well lubed.

Readers may be a bit baffled by the questioner's reference to "static directional stability". After all, aren't a car's static properties the ones it has when it's motionless? And any car is directionally stable when it's sitting still, right?

The questioner is in fact using terminology that is familiar to engineers. It has been traditional to analogize a car's directional stability to a statically stable stationary object, i.e. one which rights itself when disturbed a moderate amount by an outside influence or force, rather than tipping over. A directionally stable car tends to "right itself" similarly in yaw. If the car is disturbed in yaw while running straight, say by one wheel hitting a piece of debris, it will then travel down the road in a yawed condition, with all tires running at a slip angle. The car's inertia then has a car-lateral component, as in cornering. If the car understeers in gentle cornering, it is said to have greater cornering stiffness at the rear than at the front. If that is the case, it will also tend to straighten itself out when disturbed; it will tend to rotate in the direction of its own inertia rather than the other way, absent any steering input from the driver.

Usually, discussion related to this ignores aerodynamic factors in directional stability, but actually the analogy applies, and its relation to under/over-steer applies, when aerodynamic yaw moments and downforce/lift are present.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

BEAM AXLE PROS AND CONS

You mentioned in your latest newsletter [June 2006 issue] *that a beam axle recovers camber in roll (pretty much all of it) but that the usual case for an independent front end is that camber is not recovered in roll at all. So, why not put a beam axle at the front of the Morgan? I know that there is a possibility of tramp and shimmy but surely these can be avoided with appropriate engineering. Most trucks have beam axle front ends as did the Indy roadsters. There do not appear to be shimmy and tramp troubles with those vehicles.*

A while back a friend (ex-chief of Ford New Product Development) of mine told me about a '48 Zephyr he had. This had a beam axle front end. He reckoned there was no shimmy trouble at all. This was due to the transverse leaf spring system which put the attachment points of the spring well outboard. The GM cars of the time had longitudinal leaf springs which were located some distance inboard, in order to provide reasonable lock for the front wheels. They had terrible problems with shimmy and tramp. He surmises that, due to the advantageous outboard positioning of the spring attachment to the axle, in the Ford the spring had better control of the axle. Hence no shimmy or tramp.
He also thinks that the geometry of the shackles had a beneficial effect during cornering as well. When the car negotiated a corner on the tension side the shackle would rotate to form a straight line with the spring while on the compression side the shackle would move more upright. How this worked probably needs some further consideration and thought...

I've searched the literature for any information or papers about the Ford transverse leaf spring front beam axle but without luck. There is plenty about the GM or conventional longitudinal leaf spring and all its woes.

Surely people must have been aware that the transverse spring with its outboard mounting on the axle was far less susceptible to the troubles that plagued the conventional system. The fact that Indy roadsters employed a torsion bar system with outboard pickup points for the lever arms suggests it was known that outboard actuation had a beneficial effect controlling the axle *to prevent shimmy and tramp etc. I have not been able to locate anything that records or documents this knowledge in the literature though.*

Are you aware of anything written up about this? Is there any technical information with a mathematical or even a descriptive treatment of the outboard or transverse spring front beam axle?

I understand from the Milliken book about Olley that front IRS was used to replace the front beam primarily for the purpose of eliminating steering kick, shimmy and tramp. It's likely Ford followed suit for marketing and competitive reasons but they didn't need to as their axle was already under good control. So, it may well be that a modern road car (sedan) could benefit from the fitting of a front beam axle since that suspension would keep more tyre on road in most situations. Since the COG tends to be relatively high for such road cars (compared to a race car) and since it is usually not possible to arrange for enough anti-roll stiffness (without wrecking the ride and mechanical grip) surely there is a case for the re-introduction of the front beam.

Have you any comments on front beam axles for race cars as well as the potential for application to modern road cars?

I doubt that the earlier questioner is up for the kind of major surgery and engineering involved in converting an existing car to a different style of suspension. Street rodders do this, of course, but usually only in the course of a comprehensive frame-off rebuild involving many other modifications.

Regardless, it is definitely interesting to explore the pros, cons, and possibilities of beam axles, and to try to understand what is involved in optimizing them.

I have never driven one of the early GM cars with the beam axle front end, but I have driven a few cars and more trucks with beam axle front ends, and although at times you can feel the axle gallumphing around down there, shimmy and tramp are not common problems. This is true even with parallel leaf springs.

Many trucks nowadays have power steering, which tends to damp out vibrations, but I drove one Class 8 (big) truck a few times that had no power steering, parallel leaf springs, and, I'm pretty sure, no damping aside from inter-leaf friction in the springs. The steering was really heavy, but it didn't shake. It gave you a workout, but it didn't beat you up.

I Googled <"wheel tramp" or "shimmy"> just to see what would come up. Four pages into the search, I had seen a lot of mention of wheel imbalances and worn parts as causes, but absolutely no mention of beam axles as a cause, nor of greater susceptibility in beam axles. There was plenty of mention of independently-suspended cars having shakes in the front wheels.

It may be useful to briefly discuss exactly what we mean by the terms tramp and shimmy. Tramp means a linear up-and-down oscillation of the wheel: basically, the wheel hopping up and down. Shimmy means an angular oscillation of the wheel about the vertical axis (in the yaw or toe direction) and/or the horizontal axis (in the roll or camber direction).

I have read in chassis books explanations of a theory of a combined tramp and shimmy effect in beam axles which relates to gyroscopic precession on a one-wheel bump causing the wheels to want to steer in the direction that the axle tilts, while tire scrub is trying to steer them the opposite way, and somehow all of this gets into a self-exciting feedback loop and the steering shakes like crazy. Oddly, in all the beam axle vehicles I've driven, I've never seen this happen.

I have also read that steering shake didn't become a problem until people started putting front brakes on cars. This apparently caused problems for two reasons. One was that the torque from the front brakes would cause wrap-up of the parallel leaf springs. This would cause the caster to diminish under braking, or even go negative, and this would cause a loss of self-centering in the steering, reducing forces that try to hold the wheels straight or even creating forces that try to pull the wheels away from straight. Additionally, the axle could oscillate rotationally on the springs (springs cyclically wrapping and unwrapping). With a longitudinal drag link, this could also cause the wheels to steer back and forth as the axle rotated, creating all sorts of playful behavior. Height of the drag link, relative to the axle's center of rotation, would have a big influence on this.

Speaking of drag links, in the old days most of them on parallel-spring front ends didn't move in arcs that agreed very well with the motion of the axle, at least for large motions. The springs generally had single pivots in front and shackles in back, so the axle roughly moved in an arc about a point in front of the axle, while the drag link pivoted about a point behind the axle. Modern parallel-spring suspensions on trucks often have the steering box ahead of the axle, and the motions of the drag link and axle agree much better.

The Ford layout of the early '30's was better than the conventional layout of the time in this regard, and would have had less bump steer. Actually, there were two Ford steering layouts. Later models had transverse drag links. The Zephyr would have been one of those.

The Fords also reacted brake torque through radius rods rather than leaf springs. Although the radius rods were fairly thin in section, they were a lot more rigid than a leaf spring.

The other thing that changed when front brakes were added was that the unsprung mass natural frequency in roll greatly decreased. We are considering the frequency at which the axle assembly will oscillate resonantly, in a mode where one wheel moves up as the other moves down – that is, a two-wheel, 180-degree out-of-phase tramp. Other things being equal, this frequency goes down dramatically as we add mass at the outboard ends of the axle. As the frequency goes down, there is increased likelihood that road irregularities will excite the system at or near its natural frequency and cause a resonant oscillation.

Because of this, early front brakes were often smaller than the rears, to keep them lighter, and also to reduce the torque the front springs would have to react. However, it became apparent that for shortest braking distances, the front brakes should be not only as powerful as the rears, but at least twice as powerful. So the front brakes grew, and with them the antics of the front axles.

The questioner brings up the attendant issue of what is sometimes called spring base: how far outboard on the axle the springs act. If the springs act further out, the wheel rate in roll increases, and so does the natural frequency in roll, for both the sprung and unsprung masses. Assuming no anti-roll bar, the wheel rate in roll, and the angular anti-roll rate, increase with the square of spring base. The sprung and unsprung mass natural frequencies in roll increase with the square root of the wheel rate in roll, or the angular anti-roll rate. So the natural frequencies in the roll mode are directly proportional to the spring base.

It is also possible to raise the wheel rate in roll by adding an anti-roll bar. This can help a car with the springs set close together. I don't know when the first anti-roll bar was used or patented, but their widespread adoption coincided with the widespread adoption of independent suspension, and people didn't start using them on beam axles until later.

Incidentally, I have read one or two authors who say that anti-roll bars are unsuitable for beam axles. That is not so. Anti-roll bars may be less necessary with beam axles than with independent suspension because of the usually ample geometric anti-roll of a beam axle, but they do exactly the same thing in a beam axle suspension as they do in an independent suspension, and can be very useful.

Not only does it matter how far apart the springs are, but also how far apart the shocks are. Most dampers in 1930 were lever-action designs. Many were still friction shocks. Some were relatively crude hydraulics. In the Fords, the shocks were mounted with the levers extending transversely, outboard from the frame, so that they acted further out on the axle than in most of the parallel-spring designs. In a traditional parallel-spring design, the shock levers extend longitudinally, parallel to the springs and frame rails, and the shocks act on the axle near the spring mounting points. This results in a system that is not only comparatively lightly sprung in roll, but also lightly damped.

This is particularly true with dampers like the early hydraulic shocks, which were more or less fixed-orifice dampers and therefore steeply progressive. To prevent excessive harshness at high velocities, they had to be very soft at low velocities. This resulted in wallowy vehicles, and poor wheel control over low-amplitude disturbances. And with the shocks mounted well inboard, even a fairly large disturbance at the wheel looks small to the shock.

Consequently, if you were a passenger car manufacturer around 1930, seeking shorter stopping distances with large front brakes, and seeking less impact harshness by replacing friction shocks with the early hydraulics, you would end up producing front suspensions that were prone to 180-degree out-of-phase tramp. If this weren't bad enough, any attempt to improve ride by using softer springs would worsen every problem we've discussed so far.

Another factor in this would have been the roads of the day. In 1930, most roads in the US were still dirt or gravel. Until the discovery and development of the oil fields in Texas in the 1920's, asphalt was expensive. Even after it became cheaper, it took time to lay the acres of pavement we drive on today.

Dirt and gravel roads are prone to washboarding: they develop regular ripples that coincide with the unsprung mass natural frequencies of the vehicles that run on them. Nowadays, we mostly see in-phase ripples that excite rear axles in the ride mode. However, in an era where front axles were prone to excitation in the roll mode, we would have had out-of-phase ripples developing that would create resonant oscillation in that mode. I wasn't around then, but this would seem a logical expectation.

This would at least partially explain why we aren't seeing today the tramp and shimmy problems in beam axles that Olley and his associates faced.

One of the big reasons Ford abandoned beam axles was packaging. If we want to build the car low, and if we want to place any portion of the engine in the plane of the front wheel axis, we have a problem finding room for a beam axle under the engine. Since the axle has to move up and down, it needs more space than the frame crossmember we have under the engine in an independent suspension layout. If Ford had retained the beam axle, the '49 Ford would have had to be a much taller car, or a considerably longer one.

Packaging remains a dominant issue in passenger car suspension design. If a suspension will give at least decent dynamic properties, and it saves room, it is attractive for a passenger vehicle. This is a big factor in the continuing popularity of strut suspension. Any rival concept must be competitive in terms of space-saving, even if it offers superior dynamics.

For race cars, I think beam axles offer interesting possibilities. The UNC Charlotte team, which I have been advising, has considered beam axles in the past, and is taking a fresh look at them for 2007. In many classes of racing, the choice of beam axle or independent is made for us by the rules. Even where the rules do not explicitly dictate the choice, packaging considerations may virtually dictate independent suspension at one end of the car. In that case, we may want independent suspension at the other end, merely to maintain similar properties at both ends of the car.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

PANHARD BAR LONGITUDINAL LOCATION

What is the effect of locating the Panhard bar in a live axle rear suspension a) behind the rear axle, or b) just ahead of the rear axle, or c) even further forward in the car? For example, the Frankie Grill All-American Race Car chassis is now attaching the Panhard bar to the right-side rear door post. From there the bar runs across the car and about eight inches rearward, and attaches to a bracket extending about eleven inches forward from the left axle tube. These cars appear to be dominating at the current time.

These cars are a variety of Super Late Model, running on paved ovals in the northeastern US. Longitudinal axle location is by a form of 3-link system, with dual, compliant upper links. One upper link reacts tension forces occurring under power. The other reacts compression loads occurring in braking, including engine braking. The two links generally have different angles. They both have rubber or urethane biscuits in them that can be varied to change the rigidity of the link. The rules prohibit compliant lower trailing arms, which are commonly used at the right rear where legal.

The effect of locating the Panhard bar forward or back with respect to the axle depends on the layout of the rest of the rear suspension system. In some cases there is little or no effect; in other cases there can be a significant effect.

To understand this better, it is helpful to introduce the concept of the rear axle axis of rotation in roll. This is sometimes called the axle's roll axis. There is nothing wrong with calling it that, provided one understands that it is not the same as the the car's roll axis, the line connecting the front and rear roll centers.

The axle roll axis is a notional line about which the axle moves in the roll mode of suspension movement. The point where this line intercepts the rear axle plane – the vertical plane containing the rear wheel axis – is considered the rear roll center.

The axle roll axis is usually constructed as a line connecting two points: the instant center of the links or arms that locate the outer ends of the axle, and a point taken as representing the height and longitudinal location of the Panhard bar or equivalent lateral locating device.

In some cases, the assignment of these points can be rather tricky, and may call for some approximation. For example, it is quite possible that there may not really be an intersection point of the lower trailing arms in a three-link system. The link centerlines may converge a bit toward the front of the car, and may have an intersection in plan view, but they may pass over and under each other at that location, rather than truly intersecting. Or they may be parallel in plan view and therefore have no intersection, even if they lie in the same plane.

In the former case, it is reasonable to take as an assumed front point for the axle roll axis a point midway between the two link centerlines, where they pass over and under each other. In the latter case, in side view the axle roll axis is parallel to the trailing arm centerlines, or an average of their inclinations if they are not parallel in side view.

In the former case, we need a second point to determine our axle roll axis. In the second case, we know the inclination of our line, but not its height, so we need a point to establish that.

In both cases, we take for this a point representing the height of the Panhard bar or equivalent lateral locating device. In a passenger car, with a Panhard bar, the usual practice is to assume that the car has close to 50% left weight, and take the point where the Panhard bar centerline intercepts the vehicle center plane. If the Panhard bar is centered in the car, this will also be the midpoint of the Panhard bar.

Things get a bit more complex when the c.g., or the Panhard bar, or both, are offset significantly to the right or left. Here, we have a choice of two methods. We can take the point where the Panhard bar centerline intercepts the sprung mass c.g. plane (the longitudinal, vertical plane containing the sprung mass c.g.). Alternatively, we can take the Panhard bar midpoint.

When the Panhard bar is significantly off center, and significantly inclined, as in many dirt chassis these days, the heights of the c.g. plane intercept and the bar midpoint can differ by as much as two or three inches. Which method is more correct? They are both reasonably correct, provided we apply them properly. If we use the c.g. plane intercept, we do not make an additional correction for the vertical component, or jacking force, resulting from the Panhard bar inclination, when modeling roll behavior and wheel loads when cornering. If we use the bar midpoint, we have to include the jacking force in our calculations. The former method is simpler, and yields a good enough approximation for most purposes; the latter is more rigorous and accurate, but more complex.

In any case, by some rationally defensible method we choose an effective acting height for our lateral locating mechanism. We now still have to assign it a longitudinal position. If we have a Panhard bar that runs straight across the car in plan view, this is presents no difficulty. On the other hand, if the bar has significant plan view angularity, we have a bit of a puzzle. I think the right approach for assigning the x-axis, or longitudinal, coordinate is to use the x coordinate of the bar midpoint in all cases.

We now have vertical and longitudinal coordinates for a point that is a reasonable approximation of a lateral force coupling point between the sprung mass and the rear axle assembly. We can now draw our rear axle roll axis, or axis of rotation in roll, in side view. If we have an exact or approximated instant center for the longitudinal locating links, we draw our line from that point through the lateral force coupling point. If instead we know the inclination of our axis of rotation, we draw a line at that angle, passing through the lateral force coupling point.

Once we do that, we can see where this axis of rotation intercepts the axle plane, and we can take that as our rear roll center when modeling roll and wheel loads in cornering.

Now, returning to the original question, what happens to the roll center when the lateral force coupling point moves forward or back? It depends on the rear axle roll axis inclination angle.

If the axle roll axis slopes down toward the front, then moving the lateral force coupling point forward while keeping it at constant height raises the rear roll center. If the axle roll axis slopes up toward the front, the effect reverses: moving the lateral force coupling point forward while keeping it at constant height lowers the roll center. If the axle roll axis is horizontal, then we get no change in rear roll center height from moving the lateral force coupling point forward or back.

There are other effects as well, when we move the Panhard bar forward or back. If the axle rotates under power or braking, as it does when the upper link is compliant, the end of the Panhard bar that attaches to the axle rises or falls as the axle rotates. That means the roll center rises or falls with power or braking. The further the Panhard bar is from the axle centerline, the more it rises or falls, and the more the roll center rises or falls.

When the Panhard bar is far ahead of the axle, as the questioner describes, the roll center rises under power and drops under braking. That makes the car tighter (adds understeer) on entry and loosens the car (adds oversteer) on exit. I don't see how that would make a car faster, but it would make it different. There are other ways of controlling the car's balance during entry and exit, so with the right combination overall, such a car could win races.

One advantage of having the Panhard bar really far forward, if you’re going to have it ahead of the axle at all, is that it's easier to keep the Panhard bar out of the way of the driveshaft, without putting a bend in the bar. There are other packaging implications as well. It becomes harder to find room for the oil tank and the battery behind the driver. Overall, I would have to judge this idea a mixed blessing, and ascribe any success to users having the overall combination dialed in.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

PANHARD BAR JACKING FORCES

How do jacking forces (anti-roll, anti-dive, anti-lift, anti-squat) affect transient wheel loading during corner entry/exit? I realize that total lateral load transfer is purely a function of of c.g. height, track width, and lateral acceleration (neglecting the load transferred due to c.g. movement); however, I also feel that magnitude of the jacking forces should have some bearing on how the loads are transferred during the non-steady-state stages of the corner.

Taking a NASCAR rear axle, for example, we can de-rake the track bar (increase the left side track bar height, decrease the right side track bar height) such that the rear roll center height, as it is traditionally understood, remains unchanged. This causes the rear axle jacking force to increase, which will cause the rear spoiler to rise during cornering, and until the spring forces have changed enough to balance the jacking force, keep the left rear tire more heavily loaded – thus keeping the diagonal percentage higher during early corner entry. Is my thought process correct? Can we make similar conclusions regarding the front suspension jacking forces?

Jacking forces are the source of geometric roll and pitch resistance. They are present whenever the tires are generating horizontal forces (lateral or longitudinal) at the contact patches, and they influence wheel loads whenever they are present. Therefore, they influence both transient and steady-state car behavior.

The reason you will sometimes hear that jacking forces have disproportionate influence in transient handling is that when we have an abrupt control input (usually steering or brakes, or both), the forces at the contact patches build up more rapidly than the roll and pitch displacements of the sprung mass. Consequently, for a brief time the elastic components of roll and pitch resistance are smaller than in steady-state longitudinal and/or lateral acceleration, and the geometric components accordingly assume greater importance.

I am reluctant to believe that this effect significantly influences entry or exit behavior in oval track racing. The steering and braking are too gentle and prolonged. The dominant factor in oval track turn entry is the combination of fairly steady braking and turning together, over a period of roughly one to four seconds. The duration of this phase of the cornering process, the severity of braking, and the abruptness of brake and steering application and release all vary with the track, the setup, and the driver's style. However, we can say with certainty that the large radius of the turns inevitably precludes really abrupt control inputs, compared to what we see in other realms of motorsport, unless we are dealing with an unusually small oval.

On the other hand, I am willing to believe that lag in pitch and roll displacement is significant in a passenger car test track j-turn or lane-change test, in a chicane or street intersection turn in road racing, or in the sort of tight turns we encounter in American autocross.

In a NASCAR oval track chassis, a Panhard bar that slopes down from its attachment on the left axle tube to its frame attachment on the right does create a force trying to lift the rear of the car. This force is present through the entire turn, not just during entry. This force does not just load the left rear tire. It does pull down on the axle on the left, but it also lifts up on the frame on the right. Its effect is most commonly modeled as a force spreading the axle and frame apart, acting at the midpoint of the bar's span, approximately in the middle of the car.

If the car has little or no rear spring split, a force in the middle of the car, lifting the frame away from the axle, gets the rear spoiler up in the air but does not significantly change wheel loads, except by aerodynamic effects. However, current NASCAR setups use considerably stiffer springs at the right rear than at the left rear, so there is some increase in left rear load, and diagonal percentage, because of that. If the car has a left-stiff rear spring combination, the effect reverses, and the jacking force actually increases right rear tire loading and reduces diagonal percentage.

Again, these effects persist through the entire turn, and only go away when the rear tires cease making lateral force.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortiz@vnet.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL CENTER IN TRAILING ARM FRONT SUSPENSION

How does one determine the roll center of a trailing arm front suspension as in the VW beam front end? I have a Formula V with link pin trailing arms front suspension with 2 degree negative camber offset bushings. How can the roll center be raised or lowered in this type of front end? I'm thinking it is similar to a straight axle, because the intersection of the lines drawn through the two trailing arms is at infinity. Please correct me if my premise is wrong.

The short answer to the first part is that the roll center is at ground level, and you cannot adjust it or move it. As to the second point, the suspension behaves like a beam axle in ride, but it is very different in roll. A beam axle has a roll center well above ground level.

That's the short answer, and for practical purposes it's fairly close to correct. However, while both of the above statements are close to correct, they aren't quite perfectly accurate. A beam axle suspension can theoretically have a roll center at or even below ground level, but the linkage required to do this is unusual, and I have never seen an actual beam axle suspension with a roll center that low. And when the VW trailing arm suspension is in a rolled condition, the roll center, properly assigned, isn't exactly at ground level. That is, in the real world, the suspension actually generates a small geometric anti-roll moment when cornering. When the car has two or three degrees of roll, the anti-roll moment from the front suspension is sufficient to produce a significant modeling error if we imagine the roll center to be at ground level.

In any independent suspension, the front (or rear) view force line is an instantaneous perpendicular to the path that the contact patch center travels as the suspension moves. In the VW front suspension, the contact patch always moves straight up and down in front view, relative to the sprung mass. This is also true of many (though not all) beam axle suspensions, in ride, but not in roll. That makes the VW front suspension similar to a beam axle in ride, but not in roll.

When the contact patch moves straight up and down in front view, relative to the sprung mass, the force line, being perpendicular to the line of travel, is always horizontal relative to the sprung mass.

We can also view the trailing arm suspension in terms of short-and-long-arm (SLA) suspension principles. Looking at it this way, the front view projected control arms are: parallel to each other; horizontal relative to the sprung mass; and infinitely long. There is no instant center, or we might say that the front view virtual swing arm is infinitely long. The force line, which would normally run through the contact patch center and the instant center, becomes parallel to the front view projected control arms, and therefore horizontal – again, relative to the sprung mass. It remains horizontal at all ride heights.

The two front view force lines for the two wheels lie right on top of each other: they coincide. They can thus be said to have an infinite number of intersections, or to have no single definable intersection. If we try to define the roll center as the force line intersection, and if we believe that the lateral position of that intersection is significant, this is a very problematic case. The traditional expedient is to make a special rule for this case and say the roll center is on the centerline – but that's obviously arbitrary.

If that weren't enough, suppose we consider that nothing in the real world is made to zero tolerance, so there is a good chance that a real VW's force lines won't be exactly horizontal or parallel. That means they probably do have an intersection, somewhere – and it could be miles from the car, and way above or below the ground. Horrors! The car's behavior is completely unpredictable!

But it's not, of course, and VW front ends are no less tolerant of minor production and setup variations than any others. We merely have a case here that shows up the analytical deficiencies of taking the force line intersection as the roll center. If, instead, we look at the system in terms of individual wheel anti-roll geometry, and assign a roll center – a notional coupling point between suspension system and sprung structure, for lateral forces only – things start to make sense.

The height of each force line's vehicle longitudinal centerplane intercept remains at ground level for any pure ride displacement. The suspension thus has what I call a Mitchell index (and Bill Mitchell calls an incline ratio) of zero, throughout its travel. This means that the suspension's anti-roll geometry does not change in ride, but it changes in roll. In roll, the outside or loaded wheel gains anti-roll, and the inside wheel loses anti-roll.

In the pure trailing arm suspension, the anti-roll is zero on both wheels in the unrolled condition. So as the car rolls, the outside wheel assumes some positive anti-roll, and the inside wheel assumes some negative anti-roll, or pro-roll. The force lines remain parallel to each other (disregarding real-world variabilities), but they are no longer parallel to the ground. They both are horizontal relative to the sprung mass, but the sprung mass is leaning with respect to the ground, and therefore the force lines also both lean with respect to the ground, a similar amount. With respect to the ground, the outside wheel's force line slopes up toward the vehicle centerline, and the inside wheel's force line slopes down toward the vehicle centerline.

Okay, the inside wheel has pro-roll, the outside wheel has anti-roll, the force lines have the same absolute angle relative to the ground, and they intercept the vehicle centerplane equal distances above and below ground – so can't we just average the height of those centerplane intercepts and say the roll center is still at ground level and there is no net anti-roll or pro-roll?

No. The reason we can't do that is that although the force line angles are equal, the forces at the contact patches, that are acting along those force lines, are not equal in magnitude. In most cases, the more heavily loaded tire is making more lateral force than the unloaded one, so the vertical component of the resultant along the force line is correspondingly greater, if the force line angles are equal. We have more cornering force on the outside wheel than the inside wheel, and therefore greater anti-roll force than pro-roll. Net effect: net anti-roll; roll center height above ground.

How far above ground? It depends on two things: the amount of roll, and the amount of front load transfer.

My video, "Minding Your Anti", provides an illustrated explanation of how to assign a roll center height for known force lines and a known or assumed tire-pair lateral force distribution. I like to think I can paint pictures in a reader's mind with words, but explaining this method without illustrations is beyond my powers. Bill Mitchell's method for finding the force-based roll center agrees pretty well with mine.

Anyway, let's take one case as an example: a VW front end with a 52-inch track, rolled 3 degrees, with the inside front wheel at the point of impending lift – all load on the outside tire. Using my method, the roll center height is 2.72 inches. That's significantly different from ground level.

Returning to the original question, you can't adjust the roll center, but it does vary somewhat depending on how much of your total load transfer occurs at the front end and how much the car rolls.

Just as an aside, toward the very end of the era of trailing-arm front suspensions, Porsche experimented with inclined upper-arm pivot axes, to get more anti-roll and camber recovery. That's a story for another time, and not applicable to Formula V, but it is interesting to note that there are ways to make a trailing-arm front end that does not have quite the properties we usually think of.

Also speaking of Porsche – but now referring to Dr. Ferdinand Porsche himself, rather than the modern company – I often wonder what he thought he was getting in terms of roll center with that suspension. It may be that he thought the roll center was at the height of the springs, as with a beam axle. His choice would have made more sense on that basis.

One other note: although a VW front end does generate some anti-roll in real-world cornering conditions, most of the front roll resistance in a Formula V is elastic. The rear roll resistance, assuming we have one of the so-called "zero roll" suspensions, is entirely geometric, if we disregard frictional and unsprung components. The car thus provides an interesting case demonstrating the difference between elastic and geometric roll resistance, and also demonstrating that roll resistance is roll resistance, and you can get roll resistance in different ways, but the tires don't know how you did it. They respond to the amount of load they're carrying, but they can't tell where it came from.

MORE ON ROLL AXIS AND RELATED ANALYTICAL CONCEPTS

In your article in RCE April '06 [drawn from the February '06 newsletter] *you state that "So the end where we lowered the roll centre has less geometric load transfer and the same elastic load transfer as before - hence less load transfer overall."*

This I do not understand: assuming the CG remains in place, lowering the RC means that the distance between RC and CG increases, and the roll moment will increase. Lowering the RC will decrease the geometric load transfer indeed, but it will increase the roll-moment and hence increase the elastic load transfer in my opinion. In the end the total load transfer remains the same, but the way it is divided in elastic- and geometric load transfer differs.

Do you agree, or where am I wrong?

This is the passage in question, from the earlier newsletter:

If we raise one end of the roll axis and lower the other, leaving its height at the c.g. unchanged, the total geometric roll resistance is unchanged, but we increase the geometric roll resistance at one end and lower it at the other. The elastic elements – the springs and anti-roll bars – are not affected by this.

So the end where we lowered the roll center has less geometric load transfer and the same elastic load transfer as before – hence less load transfer overall. This will make that tire pair grip better, because they will be sharing the work more equally. At the opposite end, the elastic component will likewise be unchanged, but the geometric component will be increased – hence more load transfer overall.

The context of my statement is important. I am referring to a case where we lower one roll center and raise the other, so that the roll axis height at the c.g. remains similar. In that case, the total geometric and elastic weight transfers remain the same, and so does the roll angle. The elastic weight transfer at each end of the car remains unchanged. The geometric weight transfer decreases at the end where we lower the roll center; increases at the end where we raise the roll center; and remains unchanged in total. Therefore, the total weight transfer for the end where we lower the roll center decreases; the total weight transfer for the end where we raise the roll center increases; and the total weight transfer for the whole car is unchanged.

Remember, the sprung structure is one essentially rigid mass. It has one c.g., and it is the moment arm of this single c.g. about the roll axis that determines the overall elastic weight transfer. Treating the car as if it had a swivel in the middle will lead you to all sorts of erroneous conclusions.

Addition:

In the same article in RCE you gave a number of reasons why many cars have a nose down roll axis. Another reason I once heard was: to avoid torsion in the chassis you want the same roll moment in both the front axle and rear axle. As the weights front and rear differ, the distance from CG to RC have to differ to get the same roll-moment front and rear.

Torsion costs energy, and since there is only one source of energy - the engine - chassis torsion costs power and consequently speed. So to save speed you want as little torsion as possible in the chassis, and a (more or less) similar torsion-moment in the front and rear will help.

Do you agree with this vision?

No, I disagree. That is, I agree that similar roll resistance front and rear reduces torsional loading on the frame, or at least tends to in hard cornering, but I disagree that this is necessarily what it takes to make the car go fast. In certain cases it coincidentally is. In other cases, such a roll resistance distribution is disastrous.

Again, the sprung mass does not have a front and a rear piece that can act independently. The car is not a tractor and trailer, a locomotive and tender, or two men in a horse suit. The sprung structure is made up of many pieces, all of them having weight and inertia. Each of the components has its own center of mass, and they attach to the frame in various places. All of them apply inertia loads to the frame as the car undergoes accelerations. These loads stress the frame in all sorts of ways, including torsionally. No part of the sprung mass is truly rigid, nor is any other object, but we customarily ignore the deflections within the sprung mass for simplicity. We generally get away with this, because the sprung mass is relatively rigid, compared to the suspension and tires, in most cases. To the extent that the sprung structure is rigid, it is properly treated as one mass, with one center of gravity – most definitely not two independent masses with two centers of gravity.

If the car has 50/50 front/rear weight distribution and identical tires front and rear, and if the car has good aerodynamic balance, and speed is moderate, meaning the rear tires are not using most of their traction to propel the car, then equal overall roll resistance front and rear will get us close to a balance that will please most drivers. If the situation is identical except the car has 55% rear, and we make the roll resistances equal, we will have godawful oversteer. If we make the front and rear roll resistances proportional to the front and rear weight percentages, i.e. 45/55 front/rear, we will have even worse oversteer. Now, if we make the roll resistance inversely proportional to the weight distribution – 55/45 – we'll be closer to right, but there is no guarantee that that will be enough front roll resistance. Probably it will not be.

If we now put smaller tires on the front or larger ones on the back, things change again. Now, the car will not need as much front roll stiffness compared to rear as with equal tires.

If the car is at Daytona or Talladega, running near its top speed, with full power going to the rear tires just to maintain speed, it will need more front roll resistance than it would at lower speeds.

If the car is running American autocross, and the turns are tight, the front wheels will track significantly outside the rears even when the tires are sliding, and the car will need more rear roll stiffness to prevent understeer. In really tight turns, as in Formula SAE/Formula Student, the magnitude of this effect is quite startling.

My point is this: overall roll resistance should not necessarily be equal at both ends, nor should it be proportional to weight distribution. In fact, given equal tires front and rear, roll resistance generally needs to vary inversely with weight distribution: the light end needs more, not less. Even this is a highly non-linear relationship, and only applies if we are comparing otherwise similar cars and conditions. Across a broad range of tire, track, speed, and aerodynamic conditions, required roll resistance distribution has no simple relationship to weight distribution at all!

Now as to the idea that loading the frame in torsion due to unequal front and rear roll resistance absorbs energy and slows the car down, that's mostly nonsense. Yes, the frame does absorb a small amount of energy as it twists and then straightens out again, but the amount is negligible compared to tire and aerodynamic drag.

The frame also twists a bit as we go over bumps, both while cornering and while running straight. However, since the frame is essentially undamped, most of the energy we put in when we deflect it is returned as soon as the load is removed and the frame springs back. Movement of the suspension, on the other hand, does absorb energy, because the dampers oppose both a deflection and the recovery afterwards. We do save energy by softening the damping, but that doesn't necessarily improve handling, and of course controlling the car takes priority over this relatively small energy saving in most applications.

It is also true that at least in some cases we can make the car lighter if we can reduce torsional loadings, and that can raise cornering speed and also save energy. However, usually we are not at liberty to make this a dominant design priority. In passenger cars, there is a case for a live rear axle with a high rear roll center instead of an independent rear with more elastic roll resistance, from the standpoint of weight reduction. With ample elastic roll resistance at both ends, the car needs more torsional rigidity to avoid shakes and creaks over bumps, and this is one reason why cars with independent rear suspension tend to weigh more than cars with beam axle rear suspension.

Really, though, that is an argument for high roll axes rather than sloping ones, and for suspension with a soft wheel rate in warp, rather than an argument for equal overall roll resistance front and rear or for roll resistance proportional to weight distribution.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TWIST BEAMS AND JOINTED DE DION TUBES

Continuing on the topic of beam axles, there was an interesting de Dion axle design produced by the Rover Car Company. It appears it was originally invented by Stewert Tresilian for Armstrong Siddeley. The design featured a pair of fixed length half-shafts taking drive from the differential to the hubs. The half-shafts located the hubs transversely and were responsible for a small amount of track width change as they moved through their fixed radius arcs.

The hubs were joined by a beam axle that passed to the rear of the differential, behind the rear wheel axis. The beam axle was split into two portions in such a way that each portion could rotate relative to each other. A joint between the two portions could accommodate this and was also able to accommodate some plunge (transverse movement) so that the beam axle did not fight the half-shafts. There were trailing arms from the car body to the hubs.

Where is the roll centre? It appears that the roll centre is located at the centre of the differential although it can move about a little. Is this correct? If so, I think this position is somewhat higher than ideal.

Loading the differential with lateral forces is not necessarily ideal either (although many manufacturers have successfully accomplished it). The approach requires large stiff bearings in the differential housing and the method of locating the differential to the vehicle body becomes an interesting design issue (usually needing a sub-frame). All told this is an expensive solution.

Despite the roll centre being non-ideal (a bit of a pain to locate on a drawing board for some wheel articulations and difficult to place in the most advantageous location) there is an interesting and desirable effect with this suspension system. In roll it appears that the wheels do not lose camber as the vehicle body rolls outward. They do not adopt disadvantageous camber angles but relative to the vehicle body they gain camber with roll so that they stay perpendicular to the road surface. This seems to be because the split beam axle can twist at the sliding joint.

Even a normal beam axle allows the wheels to tilt outwards with roll. There is a contribution from tyre deflection and axle roll (tyres and axle are a spring mass system and there is deflection). With the Rover this can be allowed for and tuned out I think. It seems this is due to the beam axle passing behind the rear wheel axis. Am I correct?

What is the relationship between the trailing arm length, wheel camber in roll and beam axle longitudinal position?

How can I lower the roll centre with this suspension? It would be better to have some freedom about where the roll centre can be located. What is wanted is a way to take the lateral loading through a link or linkage instead of the half-shafts and differential unit. If this could be done, then the issues pertaining to mounting and locating the differential with precision would be eliminated. Also there wouldn't be such a need for such large bearings etc. in the differential housing.

Do you have any comments on this?

In related investigation I've looked at the torsion beam or twist beam axle used at the rear of many fwd cars. At first they seem to be a bit of a cheap mess but there is more to it than that rather cursory view. They are a clever and subtle design when considered in depth. And there are plenty of new applications yet to be addressed. For example, the torsion beam lends itself to being used as a variant of the de Dion system and being employed at the rear of a rwd car. As far as I know this has not been done. Are you aware of any?

There are many fwd cars with the torsion beam axle but in every case the beam axle is placed ahead of the rear wheel axis line. This means the wheels will roll outwards from a turn in the same sense as the vehicle body does. So they lose camber with body roll. Surely it would be beneficial to place the torsion beam well behind the wheel axis? If this were accomplished wouldn't it achieve the same effect as the Rover two piece axle? That is, matters could be arranged so the wheels remained perpendicular to the road regardless of body roll.

Certainly it is true with both beam axles and independent suspension that there is roll due to tire deflection, and that the camber recovery with this component is zero. This also happens with no suspension, as on a go-kart. The magnitude of the effect varies with tire selection and pressure. Consequently, we cannot predict it just by knowing the suspension geometry. Merely to keep the complexity of the discussion manageable, I generally discuss camber recovery as if the tires were rigid. That way, we can discuss properties of the linkage in isolation.

But of course, ignoring tire deflection doesn't make it any less real. Not only would we like to be able to compensate for camber change due to tire deflection, we would like to tilt the wheels into the turn slightly with roll, because they make greatest cornering force that way.

There is no free lunch here, however. There is no way for a passive suspension system to distinguish between suspension roll (oppositional motion within a front or rear wheel pair) due to sprung mass roll, and suspension roll due to road irregularities. Reducing camber change in the former case inescapably increases it in the latter case. With independent suspension, to obtain better camber recovery in sprung mass roll, we must accept more camber change both in ride and in suspension roll motion resulting from road irregularities. With a beam axle or de Dion system, we do not have the penalty in ride, but we still have it in roll. That is, to get more than 100% camber recovery in roll, we must increase the already considerable camber change over one-wheel bumps. That is what the Tresilian design does.

The twist beams used in front-drive cars typically play the tradeoff the other way, and accept poorer camber properties in sprung mass roll in exchange for better camber properties over bumps. Also, when the twist beam is further forward, the suspension's properties change less when the inside rear wheel lifts off the ground, as it commonly does in front-drive cars during limit cornering.

The idea of allowing the de Dion tube to twist in roll is not unique to the Tresilian design. The Mercedes-Benz W125 also had a de Dion tube that could twist freely, with the twist joint behind the differential, and simple trailing arms running forward from the hubs. The Mercedes design differed from the Rover in that the de Dion tube was not allowed to telescope. The halfshafts accommodated plunge instead. Lateral location was provided by a roller on the de Dion tube, running in a vertical slot cast and machined into the rear of the differential housing. Since this was a race car, and noise isolation was not a concern, the diff was mounted solidly to the frame. Lateral force was reacted through the differential housing, but not through the differential bearings.

We should note that thrust loads due to cornering force always have to be reacted through a bearing, one way or another. If we do not put the load through the diff bearings, the wheel bearings have to absorb it instead, and the wheel bearings are unsprung. Therefore, there may be a rational case for making the diff bearings a bit heavier if that allows us to reduce weight at the outboard ends of the de Dion tube.

Now, where is the roll center in a twist beam system? It is not at the height of whatever provides lateral location, except in the case where the twist beam is exactly in the axle plane (the vertical, transverse plane containing the wheel axis).

To locate the roll center on a drawing board, here's what we do: in a side view of the system, draw a vertical line from the center of the trailing arm pivot down to the ground. Insert or establish a point where this line meets the ground. We're going to call this point A.

Draw another vertical line from the wheel center to the ground, and also upward a ways from the wheel center. This represents the axle plane, in edge view.

Draw a third vertical line through the twist axis of the beam or deDion tube. This represents the twist axis plane (the vertical, transverse plane containing the twist axis), in edge view.

Now find the lateral force coupling point between the twist beam assembly and the sprung structure. Note that this is usually not the roll center, but it is a definable point. In a VW Rabbit/Golf rear suspension, it is the trailing arm pivot. In the Rover system, it is the diff axis, or more precisely the inboard U-joint center. In the Mercedes W125 system, it is the center of the roller. Draw a horizontal line from this coupling point to the axle plane. Insert a point at that intercept. We're going to call this point B.

Draw a line passing through points A and B, long enough so that it passes through the twist axis plane. Find the intersection of this line and the twist axis plane, and insert or define a point there. We're going to call this point C.

The height of point C is the roll center height. In the Mercedes or the Rover system, point C will be above point B. In the VW system, point C will be below point B.

The roll center is a modeling abstraction. It is the effective force coupling point for lateral forces passing between the two-wheel suspension system and the sprung structure. It is always considered to be in the axle plane. So, to complete our work on the drawing board, we construct a horizontal line from point C to the axle plane, and that intercept is our roll center.

If we want to be really rigorous, in a rolled condition the Rover design has slightly different anti-roll properties for the inside and outside wheels, and overall anti-roll depends on tire lateral force distribution, as with independent suspension. For most practical purposes, though, I think we can safely ignore this.

It would be possible to put the twist beam or de Dion tube ahead of the diff in a front-engine, rear-drive car, but there are some packaging issues. The beam or tube has to have enough room to go up and down without hitting the drive shaft.

It is also possible to have a lower roll center than in the Mercedes or Rover designs, while still having the tube or beam behind the diff, and still having greater than 100% camber recovery for the suspension component in cornering roll. All that is needed is a lower point B. This can be obtained with any of the known lateral locating devices. The roll center will still be somewhat above point B, assuming point B is above ground level. With a Mumford linkage, point B could be below any point on the linkage, and the roll center could then be as low as you'd probably want it.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
November 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

STIFF REBOUND DAMPING

I am working for a dirt late model stock car in the midwest. Recently I have been listening to some pretty good racers outside our area and they started "tieing down" their front end and right rear with excessive rebound. They are talking in the range of 200 - 300 pounds at 6 in/sec (depending on track condition and spring rate in that corner). One of the drivers I talked to said it gives the car a very comfortable, predictable feel. I was always taught that excessive rebound can take grip out of the tire. I know many NASCAR teams are now coil binding and tieing down their front end, mainly I think for aero advantage. Is there a better path to follow to get the front end down besides excessive rebound or is the advantage of this worth the loss in grip through excessive damping? I can see there is a trade off dilemma here. I would also like to here your opinion on linear vs. digressive shock valving; I listened to both sides of this and I tend to think digressive seems to make more sense.

We could say that most of the time, any damping at all takes grip out of the tire, and so does any spring rate greater than zero. That is, ideally we'd like the tire to go over bumps with no change in load whatsoever. This is not possible in the real world, for a variety of reasons, but the basic idea is that we'd like the suspension to be as compliant as possible, from the standpoint of keeping the tires in best contact with the road.

On the other hand, from the standpoint of aerodynamic control and camber control, we want a go-kart: no compliance at all.

All suspension settings are compromises between these considerations.

Additionally, in some racing classes, people do funny things to work the rules. In stock car racing, including all NASCAR classes, there is a ground clearance rule. We want the car to have enough static ride height to pass tech, and still go through the turns with the front valance barely clear of the ground. Once the valance is down there, we'd like it to go up and down as little as possible, so it can stay as close to the road as possible without scraping. That means we want the front suspension to compress easily, especially when cornering, and then go solid when it reaches the desired ride height. This is terrible for riding bumps, but if the track is smooth and air speed is high, it works.

Now, does this work when air speed is lower, and the track is rough? Probably not. Of course, this depends on just how much slower and rougher the conditions are. Dirt Late Models do go fast enough to generate serious aero forces. However, even with the setups we have come to think of as conventional for these cars, the lower edges of the body, along the right side and the right front, commonly dig into the track on bumps and are designed as sacrificial parts: we expect them to get torn up. We make them out of tough plastic, and make them easy to replace.

I question whether it makes any sense to try to get the car to ride lower when it's already tearing up bodywork against the track.

For stock car shocks, 200 pounds at 6 in/sec in extension is not really extreme. 300 is definitely stiff. Whether either of these implies a hold-down valving depends on the corresponding value in compression. The relationship between anti-extension (extension or rebound damping) force and anti-compression (compression or bump damping) force, at a given absolute velocity (e.g. 6 in/sec) is often described as the *control ratio* at that velocity. Conventionally, the anti-extension or rebound damping force is taken as the numerator. A control ratio of 1 is a true 50/50 shock. 1.5 is a popular control ratio for most purposes. Anything over 3.0 is considered a definite hold-down valving. Anything less than 1.0 is an easy-up or hold-up valving.

Extension damping unloads the tire whenever the suspension has an extension velocity. It is widely recognized that this reduces grip on the downhill side of a convex bump. It is less widely recognized that compression damping also does this. This is so because when the suspension resists compression more forcibly, the sprung mass acquires a greater upward velocity on the uphill side of the bump. When the wheel passes over the crest of the bump, the sprung mass is still moving upward, and accelerating downward. The downward acceleration effectively reduces the weight of the sprung mass, giving the spring less to react against as it tries to push the wheel down. The upward displacement of the sprung mass translates to less force at the spring, also reducing the force pushing the wheel down.

A high control ratio does tend to improve ride quality. Or, more properly, soft compression damping improves ride quality more than soft extension damping does. Soft compression damping also tends to extend the life of the suspension components, and indeed many other components. Theoretically, this might let us build things lighter.

For readers less familiar with shock terminology, "digressive" valving properly means that the force increases with absolute velocity at a decreasing rate. Stated another way, the force vs. velocity curve is concave toward the velocity axis of the graph (conventionally the horizontal or x axis): in a conventional plot, the compression traces are concave down and the extension traces are concave up.

Some shocks sold as "digressive" actually are only digressive in the upper velocity ranges. In Bilsteins, "digressive" valvings are ones using the "digressive" piston design. The main feature of this design is that it accommodates discs with notches at their periphery, which create bleeds. The bleeds affect the entire curve, but most of all they affect the very lowest velocities. At these low velocities, they actually give the curve a progressive characteristic: concave away from the velocity axis. The force vs. absolute velocity curve has a pointy "nose" on it. The damping is soft at very low velocities. The greater the bleed area, and the greater the preload on the stack, the further up the velocity range this region of progressive damping extends.

Conversely, in Bilsteins at least, the valvings with "linear" pistons are either near-linear or digressive throughout the velocity range. So the terminology can be confusing.

Anyway, is soft damping at very low velocities desirable? It depends to some extent on the driver, and to some extent on how much sliding or mechanical friction there is in the suspension. Mechanical friction, also called Coulomb friction, adds *Coulomb damping* to the system: damping that is largely independent of velocity, and therefore highly digressive at low velocities. Having the hydraulic damping progressive at low velocities tends to compensate for this. Racing suspensions, with lots of fresh, tight rod ends and spherical bearings, have considerable Coulomb damping. Struts have considerable friction as well. Rubber bushings have little or no Coulomb damping.

And is true digressiveness at higher velocities desirable? Most experience suggests so. We can say for sure that mid-to-high-velocity digressiveness allows a more controlled feel with less harshness when hitting curbs or big bumps.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
December 2006

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

COIL-BINDING SETUPS IN STOCK CARS – FRIEND OR FOE?

In reading the November newsletter, a couple of thoughts came to mind. My application is pavement late model racing (ASA, CRA, etc.).
1) To remedy coil binding, what about using torsion bars instead of coil springs? What I am thinking of is mounting the bars parallel to the front-to-rear vehicle centerline, connecting to the lower control arms (Chrysler-style). This would get some weight lower in the vehicle as well.

Rules permitting, certainly you can use longitudinal torsion bars. However, the coil binding currently found in upper-division NASCAR racing is not an unintended problem; it's a deliberate strategy, aimed at working the ride height and spring rules, and better controlling the car's aerodynamics. The idea is to have the car pass ride height inspection, yet come down as readily as possible to a ride height where the valance just clears the track surface, and then go solid, so the valance stays at that height and doesn't rub or lift. This is awful for riding bumps, but if the track is smooth you can accept that in return for good control of the aerodynamics.

Perhaps a better compromise is to have the suspension go very stiff, rather than solid. This used to be achieved by using bump rubbers instead of having the springs hit coil bind. But NASCAR outlawed the bump rubbers. At about the same time, they instituted a minimum spring rate of 400 lb/in.

To make the spring coil bind, you either shorten its length or increase the wire diameter and number of coils, leaving the length the same. To make the spring operate in coil bind as much of the time as possible, you minimize the spring-to-wheel motion ratio, and create as much downward jacking with the suspension geometry as you can.

One insider I talked to recently tells me that crew chiefs are now using so much pro-roll on the right front that the cars are now sitting considerably higher at static than necessary to pass ride height inspection. Apparently, some people believe that if you have more "drop" to coil bind, that somehow adds load to the tires.

It doesn't.

It occurred to me that maybe having the nose of the car ride higher down the straights might reduce drag, perhaps due to the roof masking the rear spoiler more. I asked Gary Eaker at AeroDyn about that possibility. He says no, having the nose higher is bad in terms of both drag and lift. The increased air under the car adds more drag than you save by masking the rear spoiler.

FALLING-FORCE SHOCKS

[Continued from previous question]
2) For shock damping on bumpy tracks, it seems that if damping forces at high shaft speeds could be significantly reduced, grip could be enhanced. I don't mean simply digressive valving, but significantly dropping damping force, say above shaft speeds of maybe 5in/sec. I don't know if this is even possible?

What you are describing is a shock with a negative damping coefficient over part of the velocity range.

I'm not sure it's impossible, but it's tricky to create such an effect in a package that would look like the shocks we use today. It is possible to create a falling-rate leaf spring or Belleville washer (and a shim in a damper is essentially a round leaf spring, or diaphragm spring), but usually we can't actually obtain a negative rate, or a force that decreases with increasing displacement. There is one exception to this, which I will discuss below.

Although it is normally not possible to create a negative-rate or falling-force spring, it is possible to create a springing device where the force diminishes or even reverses as displacement increases, by combining a spring with other mechanical elements. One example would be the over-center spring found in some clutch linkages. These are coil springs arranged so that after about half pedal travel, the spring tries to pull the pedal to the floor, rather than pull it up. These forces act in parallel with the springs in the clutch, which resist pedal force at a steadily increasing rate with respect to pedal displacement. The net result is a falling-rate resistance at the pedal, but not necessarily a falling-force resistance. However, if the over-center spring is strong enough, it is possible to create a falling-force characteristic. On some cars, the over-center spring tension can be adjusted to create the desired pedal feel.

To produce the over-center action, the coil spring needs to have a lever or rocker to act on. I can't think of a way to incorporate that into a shock.

Ordinarily, the over-center spring is used to make a Borg & Beck clutch feel similar to a diaphragm clutch. The spring in a diaphragm clutch is an interesting case. It can actually have a falling-force characteristic – not over its entire deflection range, but over a portion of it. I have not actually tested clutch diaphragms, but I have an old textbook that shows a force-vs.-displacement curve for a typical clutch diaphragm (Herbert Ellinger, *Automechanics*, Prentice-Hall, 1972, p. 329). The curve is S-shaped. It starts out fairly linear, then becomes concave-down, and continues to bend downward until the slope becomes negative. Then it becomes concave-up, and continues to bend upward so that the slope becomes positive again. The force never becomes negative, but the rate does, over an interval.

A clutch diaphragm is somewhat like an ordinary Belleville washer, but with some differences. It has substantially more dish, and it has radial slits running from its inner diameter out more than half way to its rim. In other words, rather than being a dished continuous disc, it is a ring with fingers extending inward. Additionally, a clutch diaphragm has a perimeter outboard of the ring it bears against, that moves oppositely to the motion at the inner ends of the fingers. This feature takes the place of the release levers in other clutch designs, and retracts the pressure plate.

It would be possible to make a diaphragm spring like that, minus the reverse-motion portion of the rim, and use it as part of the stack in a deflective-disc shock. It would then be possible to have at least a portion of the valve stack have a falling-force characteristic. That would still not necessarily create a damping force that falls as absolute velocity rises. The stack force does not directly determine the piston force. Rather, it determines the flow characteristics of the orifices in the piston which the shims mask. The piston force then depends on the resistance to oil flow through these partially masked orifices.

A falling-force stack could, however, produce a more highly digressive shock than a stack with a modest rate and high preload, which is the usual approach. It might even be possible to have a true falling-force or negative-damping-coefficient shock, but a falling-force stack would not necessarily produce this.

It would be possible to mount the shock so that its force would diminish with increasing displacement, but that is not the same thing as having the force diminish as velocity increases.

It would be possible to deliberately make the shock acceleration-sensitive. (Velocity is the rate and direction of displacement change; acceleration is the rate and direction of velocity change.) Edelbrock has been selling acceleration-sensitive shocks for some years now. Acceleration-sensitive shocks generally use a weight of some kind, working against a spring. The weight can either open an orifice or move a needle to change the area of an orifice, or it can vary the effective preload on a stack.

Finally, there is one type of shock that produces a negative damping coefficient over a small portion of its velocity range, and a damping coefficient of zero over most of its velocity range. Furthermore, it is easy to make the damping force adjustable, while retaining the zero coefficient for all adjustment settings. What kind of advanced design might that be? Surprise: a friction shock! A friction shock makes the same force regardless of velocity, except when it's just starting to move. Then it has a "stiction" zone, where the friction force is greater than the force we get once it's moving. As the stiction goes away, the velocity is increasing; the damping force is decreasing; the damping coefficient is negative.

Nobody uses friction shocks anymore, but all suspensions have some sliding contacts in them. These generate what we call Coulomb damping: friction that is independent of velocity, except for the stiction effect.

Some cars still use leaf springs. Leaf springs have a lot more Coulomb friction than coils do, because of inter-leaf friction. In the days when hydraulic shocks were primitive, this was actually an advantage, and potentially it still could be today, on dirt. Much of this depends on rules, but in some dirt classes, leaf springs are actually faster than coils. Unfortunately, they are also a pain to run, because to get the soft rates needed for dirt, the leaf springs have relatively few leaves, and stresses are high. Consequently, the springs fatigue and lose ride height rapidly, and you have to replace them all the time.

The normal assumption is that friction is bad, so it's smart to minimize the number of leaves. But if Coulomb friction is actually okay, because it gives us damping that's stiff at low velocities yet no stiffer at high velocities, maybe there's a future in leaf springs for dirt that use a larger number of thinner leaves than those we see today. That would increase inter-leaf friction, and also increase weight, but result in lower stresses, longer life, and better ride height retention.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TALL VS. SHORT SIDEWALLS

I am a Racecar Engineering subscriber and have enjoyed your later articles regarding tyre width and grip.

I seem to recall another one regarding wheel and tyre profile effects as well, but can't find it. Was it written by you? Do you remember in which issue it was published?

Basically, there have been some discussions about why the current trend to big wheels/small tyre profiles and why some racecars like F1 don't use a small-profile tyre. One of the opinions is that F1 is limited to 13" wheels (to limit brake size), but if it wasn't, bigger wheels would be of benefit. Personally I think the bigger wheels/smaller profiles are (beyond a certain point) just a styling and marketing exercise.

Any insight on this point?

I think I can recall having discussions about this, but I don't think I've addressed the subject in the newsletter or column before. Perhaps I have done so on a forum at some time, or in private correspondence. Checking my newsletter back issues list, I don't see any such article listed – only the one on why wide tires are better than narrow ones.

Are big wheels with low-profile tires just a styling fad or marketing gimmick? In some market sectors they have become that, but this is a styling fad with some basis in engineering.

Actually, the trend to very low tire aspect ratios began in racing. Low-profile tires are still used in racing, where they are allowed: Trans-Am cars, touring sedans and other production-based road racing cars, autocross cars, sports racing cars. Tall sidewalls are only seen where the rules limit wheel diameter, and on the driven wheels in drag racing.

Thirty-five years ago, there were no street tires below a 60% aspect ratio. If you wanted the low-profile look that the racing cars had, you had to buy racing tires. People sometimes ran racing tires on the street, although of course that was illegal. But that was the only way to get racy-looking tires with stiff sidewalls and sticky tread rubber. Car shows were full of display-only vehicles with race tires. The reason we now see low profiles on the street, and tall profiles in F1, is that street tire technology has belatedly caught up with discoveries originally made in racing, while racing rules have restricted F1 cars, and some other classes as well, to anachronistic wheel sizes.

Is this a bad thing? Not necessarily, although it is a bit odd. If you are running a racing series, it makes economic sense to restrict any technological progress that would require competitors to replace equipment or require manufacturers to retool. If you are trying to keep speeds down and keep fields full, why would you permit anything that raises both speeds and costs? As long as the cars go fast enough and make enough noise to put on a good show, that's enough, isn't it?

Certainly it is if you are openly promoting budget racing, with technologically restricted cars. But if you are billing your series as the premier class in motorsports; if you are charging an arm and a leg for tickets; if everybody knows the expense to compete is ridiculous but this is part of the draw – then it becomes harder to defend draconian restrictions on tires and wheels. And if you are still trying to justify the show as an exercise in "improving the breed", that really does get rather awkward. Wide, low-profile tires are the main street-applicable technological advancement that racing can claim to have originated during the last forty years.

Are they really an advancement? The questioner appears to have some doubt.

I say yes, they are an advancement, although they are something of a mixed blessing and have become a customizing fad.

Making the sidewall shorter and the wheel diameter bigger has two, or perhaps three, advantages. First, it makes the sidewall stiffer. Second, it makes room for bigger brakes. And, if we don't use all the extra room for bigger brakes, we can get more air through the wheel. If we shape the fenders properly, that lets us extract more air from under the car through the wheel wells. This not only helps brake cooling, but also aids lift reduction/downforce creation. This assumes, of course, that the car has fenders.

Are stiffer sidewalls always better? I think we can say that for racing and for high performance applications, we want as much lateral stiffness as we can get. There is some penalty in directional stability, because the car will have more tendency to "tramline", or follow edges in the road surface that nearly parallel the vehicle's direction of travel. But a performance-oriented driver will generally put up with this to get more responsiveness and greater cornering power. Greater lateral stiffness helps keep the tread flat to the road and prevent the tire from rolling under and concentrating load on the outside shoulder of the tread.

Greater vertical stiffness is more of a mixed blessing, and a more complex issue. The tire is to some extent a secondary suspension system, acting in series with the main suspension system. In stiffly sprung winged formula and sports racing cars, tire compliance may be as much as half of the total: the suspension may be as stiff as the tire sidewalls.

Considered as a suspension system, a set of tires is very good in some respects, and horrible in others. For unsprung weight, it's unbeatable. The only unsprung components are the contact patch and some material near it. It has no camber change in ride. On the other hand, it has no camber recovery in roll, and it is seriously underdamped.

We might be tempted to decide that we could accept a very high vertical stiffness from the tire, i.e. a very high tire spring rate, and get our compliance from the suspension proper, where we can get camber recovery in roll and control the damping properties.

However, there is one other factor: the tire's vertical spring rate is inextricably related to the contact patch size, and the contact patch size is related to the amount of grip we have. As the tire spring rate approaches infinity, the contact patch length and area approach zero. If the spring rate were truly infinitely large, the contact patch would be a line of zero width front to back, and the contact area would be zero.

The original objective of radial tires was to have greater vertical compliance and a longer contact patch, while still keeping the contact patch flat to the road in cornering.

Theoretical considerations aside, for most purposes practical considerations limit our sidewall height. We have to have enough distance between the rim and the ground so that we don't damage the rim on pavement slab edges and potholes. One of the features that has made today's short sidewalls possible has been the development of sidewall designs with a meaty region near the rim flange that protects the rim aginst both curb scuffs and pothole damage. This is combined with a flexible zone closer to the tread shoulder that provides vertical compliance, albeit over a smaller range than with traditional construction.

It is important to note that sidewall height is not the sole determinant of sidewall properties. It is quite possible to put a big, strong bead stiffener in a tall sidewall, and make it act like a short sidewall. It will weigh a bit more, and the brake size will suffer, but the car will corner about the same as would with a short sidewall. I am told that most "radial" racing tires are not actually radials at all anyway. The main plies are actually not as close to 90 degrees as physically possible. They are more like 70 or 80 degrees – closer to radial than a true bias-ply or the bias-belted street tires we saw in the '70's, but not truly radial. In other words, relatively stiff construction is still favored where cornering performance is paramount, and therefore the trend toward short sidewalls for performance tires is fundamentally sound.

BRAKING DISTANCE, BRAKE TORQUE, AND TIRE TRACTION

I recently saw a quote by John Barnard about the early days of carbon fiber, where he said that Niki Lauda was braking for a certain corner at 100m but when carbon brakes were tried he could brake from 60m.

I have always been under the impression that if the brakes on a car were powerful enough to lock up the wheels, then the shortest braking distance would be dictated by the load on the tires and the coefficient of friction between the tires on the road. Given the same tires and road surface, how can one type of brake stop a car faster than another one?

Perhaps John Barnard reads *Racecar Engineering*, and perhaps he will see this when it appears in my column, and perhaps he will favor us with his own explanation. Absent this, I will speculate as best I can.

In theory at least, braking is primarily limited by tire traction, provided the brakes can lock the wheels. The brakes can only stop the tires. The tires then have to stop the car. However, real world braking at the end of a straightaway is constrained by some additional factors besides sheer braking power.

First of all, the driver is not locking the wheels. The driver must avoid locking the wheels, in order to maintain directional control and not flat-spot the tires. That means it is crucial that the brakes exert a predictable and consistent torque, and that the front-rear balance be appropriate.

It is important for the brakes to come up to desired torque promptly: they must have good initial "bite". They must apply equally on both sides of the car throughout the braking event, so that no large yaw moments result. Any problems in these areas will require the driver to brake earlier to compensate.

Brake release is more important than many people realize. If the brakes continue to drag after the pedal is released, not only does that heat the brakes unnecessarily, it saps cornering power from the tires. That lowers cornering speed, again requiring earlier braking.

Finally, brakes relate to driving technique, particularly as regards the ability to trail brake. It is commonly believed that trail braking was adopted primarily because certain drivers, such as Mark Donohue, recognized that it could improve lap time by allowing braking to be delayed. This is true, but it is also true that trail braking as we know it was not possible until the advent of brakes with good directional stability and consistency, and controllable release properties. In the days of drum brakes, drivers had to do their braking in a straight line because brakes were not controllable enough to allow the driver to turn and brake at the same time, or to controllably reduce braking while feeding in steering.

Any combination of these factors could allow one set of brakes to outperform another approaching a particular turn, even if both designs can lock the wheels.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MODERN DRIVESHAFT JOINTS VS. RUBBER DOUGHNUTS

I am running a Historic Formula 2 March 712 in Europe. I am working hard to improve the handling. To make the car lighter, first I removed the heavy doughnut with heavy strong driveshafts and I saved a total of 7 Kg after the modification to modern CV joints. Can I expect a improvement on handling? I think the rubber doughnuts are likely working as a spring and are not so nice to drive on corner entry and exit and I am losing power on acceleration (and the polar moment is much higher; that means the acceleration is different.) Can you give me more information about that?

Handling should be helped a little. The difference may be too small to feel or even measure, but anything that reduces weight should help handling, once the car is optimized for the new weight and weight distribution. In this case, roughly 50% of the mass in question is unsprung, and that improves the payoff, at least when there are any bumps.

Rubber doughnut joints are compliant in torsion, and this can create surging: an oscillatory longitudinal acceleration of the car, caused by the joints wrapping up and unwrapping. Any jerky application, release, or reversal of torque can provoke this. This is not strictly a handling issue – more of a driveability issue. However, we do partly control the behavior of the chassis with the application of torque to the wheels, and anything that reduces our control of torque to the wheels inevitably hurts our ability to control the car.

As for the addition to the wheel rate from the doughnuts, no doubt there is some. However, we need wheel rate; we aren't necessarily seeking to minimize it. We can compensate for any contribution from the joints by running slightly softer springs.

We do want to know wheel rate as accurately as we can, and the addition from the joints complicates this a bit. If we are really fastidious, we can measure the joints' contribution fairly accurately by putting the car on the scales, removing the springs, jacking the sprung mass up and down, and noting changes in the scale readings. Probably the doughnuts have a somewhat different rate when everything is in motion, but measuring their rate statically should be at least as valid as measuring a tire's rate statically. Of course, all this is moot if the rubber doughnuts aren't there anymore.

Regarding improved acceleration due to reduction of moment of inertia, that's also good, but I expect the effect here will also be small. It will be nowhere near as big as a similar moment reduction at the flywheel or clutch, because the drive shafts turn much slower. If the engine turns, say, four times as fast as the shafts, then any given saving at the flywheel is worth sixteen times as much as the same reduction at the shafts. The engine has to accelerate the flywheel four times as fast as the shafts, and any given inertia torque at the shafts looks only one-fourth as great to the engine because it is reduced through the gears. Still, even if the effect is small, it is surely helpful.

ZERO-DROOP SETUPS

[Continued from previous question]
The next question I have is about suspension droop travel limit fwd and aft. I have seen that modern single seaters have no droop at all. Now I am not sure about a historic racecar like my March F2 (year 1971).

The shock stroke (fwd and aft) during hard driving is about 20mm. I know aft shocks should have more travel than fwd shocks but how much?

Corner exit on acceleration the car is pitching fwd up and with no droop can help to get more weight on the car fwd because the unsprung weight will help to hold the nose down. Until now I made the shock adjustment with more rebound but I think that's the wrong way to fix the problem.

The only reason it makes any sense to not let the suspension move freely in droop is to control ground-clearance-sensitive aerodynamic elements. From every other perspective, making the suspension top out prematurely is a bad thing.

For best mechanical grip, we want the suspension to extend freely until the springs reach zero load, and then stop. If the suspension extends further, so that the springs hang loose, that doesn't hurt grip but it can cause the springs to beat up the shocks, or the spring retainers or adjusting collars, or other pieces, if it happens very often.

If the suspension tops out before the springs unload, that abruptly unloads the tires, but it also keeps the ground clearance from growing, at least as long as we don't pull the wheels off the ground. That's bad for mechanical grip, but it's good for aerodynamics if we've got a floor, a valance, a splitter, or a wing that has to be near the ground to work well.

From what I can tell by pictures on-line, the March 712 has a fairly broad chisel-shaped nose, with two small nose wings on the sides of the chisel. The radiator is in the nose, fed by an intake below the leading edge of the chisel, and exhausting through an outlet on the top of the chisel. The little wings are up fairly high compared to later cars. It doesn't look like the car would be highly sensitive to ground clearance at the front. The car pre-dates tunnels and diffusers; the floor isn't designed to make downforce. It has a rear wing, but this would not be highly sensitive to ride height either. So I don't think the car would realize the same advantages as a more modern car from zero-droop suspension.

There is no hard and fast rule for the relationship between front and rear shock travel, as shown by travel indicators on the shock shafts. In tail-heavy, rear-engined cars, it is common for the front suspension to have a higher natural frequency and smaller static deflection than the rear. That is, the front suspension is stiffer in ride than the rear, relative to its sprung mass. This normally results in more suspension travel at the rear than at the front. Also, it is normal to have more aerodynamic downforce at the rear than at the front. This will result in the shock travel indicators showing more travel at the rear than at the front, if the motion ratios are similar. Looking at photos, it appears that the motion ratios at the front and rear of the March 712 are fairly similar. So if the travels are similar, it may be that the rear springs are a bit stiff relative to the front – but not necessarily.

Keeping the front end from lifting under power will only add a little bit of load to the front tires – and concomitantly reduce rear tire loading a little. The amount of rearward load transfer, for a given forward acceleration, depends entirely on the height of the center of gravity and the wheelbase. How much the nose lifts wouldn't matter at all, except that more nose lift does result in a slightly higher c.g. If the front of the car seems to lift excessively under power, stiffer front springs will reduce that. You may want to combine the stiffer springs with more rear anti-roll bar or less front anti-roll bar.

Some readers may find it surprising that the front end lifts less with stiffer springs, but that is indeed the case. A stiffer spring doesn't mean more upward force. It means less travel for a given load change – hence less extension travel for a given load decrease.

Ordinarily, we don't want to keep load on the front tires under power; we want load to transfer to the rear so we can put power to the ground. This is true unless the car understeers excessively on corner exit. This is not uncommon in rear-engined cars, but a more common complaint in racing generally is that the car spins the wheels and/or oversteers on exit. A car that is too tight (understeers too much) power-on is a problem that racers in many classes would kill to have. If the back tires stick too well, just feed them more power, goes the reasoning.

This doesn't always work, however. There may be no more power available; the driver may have the throttle wide open already. This is often encountered in Formula Fords. Or the understeer may simply increase until power breaks the rear tires loose completely, whereupon the rear end suddenly snaps out.

The March appears to have very short control arms, especially the uppers. That means that either the camber recovery goes away markedly as the suspension extends, or the roll center rises as the suspension extends, or both. Having that happen at the front would worsen a power push, and stiffer front springs would reduce the lift. Restricting droop travel would keep the nose down too, but the effect would be abrupt, and the roll resistance would abruptly increase, worsening the push. More spring and less bar keeps the front end down more, without that problem.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DIAMOND LAYOUT

Attached are illustrations of a proposed car for F-1 racing, suggested during the early 1990s. [For newsletter recipients, one of these illustrations is included as a separate attachment.] *The car has a single steered but unpowered front wheel. A pair of wheels is located further aft located on an axle passing close to the COG or possibly ahead of it. These wheels are non-steered but are powered. Finally there is a powered and steered single wheel right aft. The car is a diamond in plan view. I understand the idea was to provide a chassis with better aerodynamic downforce and less drag than a conventional layout. In addition the car featured better traction by being able to drive three wheel drive (the rear three). Only the front wheel was not powered. This nicely got around the rules banning four-wheel drive in F-1!*

How would such a car handle? Assuming the front wheel and the rear wheel accomplish the steering task (that is the non-paired wheels steer whereas the "axle" pair remain fixed with respect to steering) what would happen as the car reached maximum slip angle and the driver decided to increase the radius of cornering path (that is, reduce lateral acceleration)? In this case he'd be adding to the rear wheel slip angle while simultaneously reducing the front wheel slip angle. What would occur if the rear wheel lost static traction because of this? Would the effect of the single rear wheel losing grip dominate? Would the increase in slip angle at the rear tend to over steer the car putting it into a slide or a spin? Or would the combination of the three other wheels prevent oversteer and slides?

An alternative arrangement would be to keep the single rear wheel fixed while steering the other three. There would need to be some sort of linkage to do this- a sort of modified Ackerman geometry needing to be implemented. Would this be a better arrangement? The designer must have thought about this possibility and thought not. It would be interesting to understand why.

An interesting feature is that it would seem that only a single differential unit would need to be used since the distance of the path traversed by the rear wheel is exactly equal to the mean of the wheel *paths of the middle wheels. The rear wheel could be driven from a shaft taken from the crown wheel.*

Comments?

The first attempt at a diamond layout that I ever saw was a model that somebody submitted to the Fisher Body Craftsman's Guild styling competition back in the 1960's. To my knowledge, nobody has ever tried to actually build a race car this way.

I love this kind of outside-the-box thinking. To give a brief answer first, I think the diamond layout holds more promise for straight-line cars such as dragsters or speed record cars than for road racing cars. There is, at least potentially, an aerodynamic benefit. There is also, unfortunately, a handling penalty.

There would be an advantage in forward traction with three wheels driven. Whether this would actually be allowed would depend on the exact wording of the rules, and their interpretation by the officials. Even if three-wheel drive were ruled legal, the rules could be changed to prohibit it if the FIA saw fit. Any dramatic innovation that obsoletes existing cars in any class faces the same problem. Even if it is clearly legal when introduced, it can still be prohibited by a subsequent rule change if those in charge deem this desirable.

Whether the diamond layout stands to win road races or not, it is fascinating to consider what its pros and cons are, and what would be involved in optimizing it.

Perhaps the most obvious problem with the diamond layout is that roll is only resisted by a single wheel pair, at least up to the point where one wheel lifts. In low-speed turns, with minimal aerodynamic downforce, and with racing tires, the inside middle wheel will lift before the tires will slide, just as we know one wheel would lift in a conventional layout, with all the roll resistance at one end. Beyond that point, the vehicle is a tricycle, or maybe a sidehack, and any further overturning moment is resisted by the front and rear wheels and the outside middle wheel, acting on half the track width. The vehicle is unlikely to flip, as long as it doesn't snag an edge while sliding, but it is not making very good use of its tires. If we are driving the middle wheels, the diff had better be able to lock with one wheel in the air.

In a road car, we would have much less grip, although we'd also have a higher c.g. Maybe a road car could be made to keep the inside middle wheel on the ground at the limit of adhesion. Even then, however, the middle suspension would need to be very stiff in roll to keep the roll angle within reason.

In a conventional layout, we can control the car's oversteer/understeer balance by juggling the relative roll resistance for the front and rear wheel pairs. With the diamond layout, we can no longer do that.

Another problem is that we cannot achieve camber recovery in roll on the front and rear wheels. Short of the point of inside middle wheel lift, we have no suspension displacement to work with on the front and rear wheels. After the point of inside middle wheel lift, the front and rear wheel suspensions extend whether the car is rolling to the right or the left. So we really have no choice but to make the front and rear wheels move without camber change in ride and lean with the sprung mass in roll – and, as we have noted, the sprung mass is apt to lean considerably.

As with three-wheeled vehicles, the worst case for overturning is a combination of longitudinal and lateral acceleration. If we compare conventional and diamond layouts, with identical wheelbase and track dimensions, with the diamond layout the center of mass is much closer in plan view to the nearest line connecting two contact patches. That means that the diamond layout will bicycle more easily in some combination of lateral and longitudinal acceleration than the conventional layout will in its worst case, which normally is pure lateral acceleration.

Short of the point of bicycling, a layout with poorer overturning resistance will experience greater load transfer. It will load its tires less equally, and will consequently use them less effectively.

One might argue that a fairer basis of comparison would be a case where the diamond layout has similar worst-case overturning resistance to the conventional layout. That would imply a longer wheelbase and wider track for the diamond layout. Ordinarily, our vehicles are constrained by the width of the lanes on our roads, and the length and width of our parking spaces. The envelopes thus defined are rectangular rather than diamond-shaped, and therefore a conventional rectangular vehicle fills them more effectively, with its wheels spread further from its center of mass, than a diamond-shaped vehicle. So a diamond-layout car needs to take up more room on the road, if it is to have comparable overturning resistance and comparable load transfer.

In racing, we usually have an overall width limit, an overall length limit, and sometimes a minimum and/or maximum for wheelbase and/or track. We also have some advantage from a narrow car, in the line we can use, especially through doglegs and slaloms or esses. All these factors favor the conventional layout over the diamond.

For the diamond layout, many of the steering, differential, and suspension design considerations are common to more conventional three-axle vehicles, and in some instances even two-axle vehicles.

Even when the two rear axles are quite close to each other, as in heavy trucks, in tight turns at low speeds there is a significant difference in the speed of the rearmost axle and the middle one. Heavy trucks consequently have three differentials, two of them housed in the forward drive axle. Many trucks have a driver-controlled lock for the center diff. There is a noticeable increase in tire scuffing in low-speed maneuvers with the center diff locked. I am not sure a second diff would be necessary in a diamond-layout race car, but the mean speed of the middle wheels would not be identical to the speed of the rear wheel in all situations.

As for steering, in multi-axle vehicles it is most common to steer the front axles and not the rear ones. This gives the most predictable high-speed behavior, at the expense of turning circle. A neighbor down the road from me operates a fleet of concrete pumping trucks. The largest of these are non-articulated trucks with seven axles. The front three steer. The linkage is designed so the front axle steers the most, the second axle less, and the third axle still less. The rearmost axle is a non-driven tag axle that is lifted off the ground in low-speed maneuvering.

Looking at the drawings by Mr. Scalabroni, it doesn't look like the front wheel could steer very much without hitting the leading arm that locates it. A different design could be used, of course, but as drawn, the car would need to have the rear wheel steer just to get around the tighter turns. If the rear wheel is driven by a longitudinal shaft, there will be problems making it steer much either.

Making a single front wheel steer a reasonable number of degrees presents some packaging challenges. There has to be some fairly bulky structure either alongside the wheel or above it. We can't have the steering axis behind the wheel; we would have negative trail.

Older readers will recall the four-wheel-steer cars sold by Honda and Mitsubishi during the 1980's. These were originally designed to cope with cramped roads and parking areas in Japan. The rear wheels steered out of the turn, adding yaw moment, in some situations, and steered into the turn, reducing yaw moment, in other situations. The Honda system was purely mechanical. It was ingeniously arranged to make the rear wheels steer into the turn at small handwheel (steering wheel) displacements, and steer out of the turn at large handwheel displacements. The Mitsubishi system, called HICAS, was electronic and computer-controlled. It steered the rear wheels out of the turn at low vehicle speeds and into the turn at high vehicle speeds. That is really what we want. The Honda system was intended to approximate this with a simpler, passive system, the thinking being that large handwheel inputs mainly occur at low speeds. In both of these systems, the rear wheels steered much less than the fronts.

There was a four-wheeled truck made in the 1930's that steered the front and rear wheels equally and oppositely. The idea there was to allow full-time four wheel drive with no center differential, and have a good turning circle without excessive U-joint angle at the outboard end of the axle. Apparently, high-speed operation was not contemplated.

In the Scalabroni, would there be a possibility of the rear tire going to a slip angle past peak lateral force when catching an oversteer slide? I think so. In fact, this is fairly common with unsteered rear wheels. It is also common for a driver to push the front wheels past optimum slip in an understeering car, and actually have the front end unexpectedly grab, or bite better, as the steering is unwound on exit. Stock car drivers sometimes call this a "push loose" condition. So, yes, it could happen, but the problem is not unique to the Scalabroni design.

The reader will note that we have so far assumed purely passive suspension. Current F1 rules do require this. However, active suspension could very significantly enhance the properties of the diamond layout. I doubt if it would equal a conventional layout even then, but again the possibilities are interesting to contemplate.

One possibility with active suspension would be to make the car lean into the turns, essentially making it a large motorcycle with training wheels. Very good computer control would be needed, along with a lot of travel for the middle wheel suspensions. Why not just have a two-wheeler instead? Well, the diamond-layout car could be larger, and would be less prone to falling over to the inside when traction is suddenly lost, and less prone to high-siding when traction is suddenly regained. A disadvantage versus a two-wheeler would be that the operator would not have the advantage of being able to straighten out the turns as much as with a single-track vehicle; the line could not be as good, due to the width of the vehicle.

It appears that the designer of the Scalabroni was thinking more of aerodynamics than mechanical dynamics. I think that is where the concept's advantages lie, and therefore I think there are some really interesting possibilities for diamond-layout speed record cars. By putting the middle wheels about a third of the way back, it would be possible to have a nice teardrop shape in plan view. Large amounts of steering lock would not be needed. Probably steering just the front wheel would be fine. All four wheels could be driven. If the rules required, the front and rear wheels could be slightly offset, rather than directly in line.

In any case, it is heartening to see this kind of original thinking.

ORIGINAL NOTE ARE IN ITALIAN; SORRY FOR MY NOT-TECHNICAL TRANSLATION!

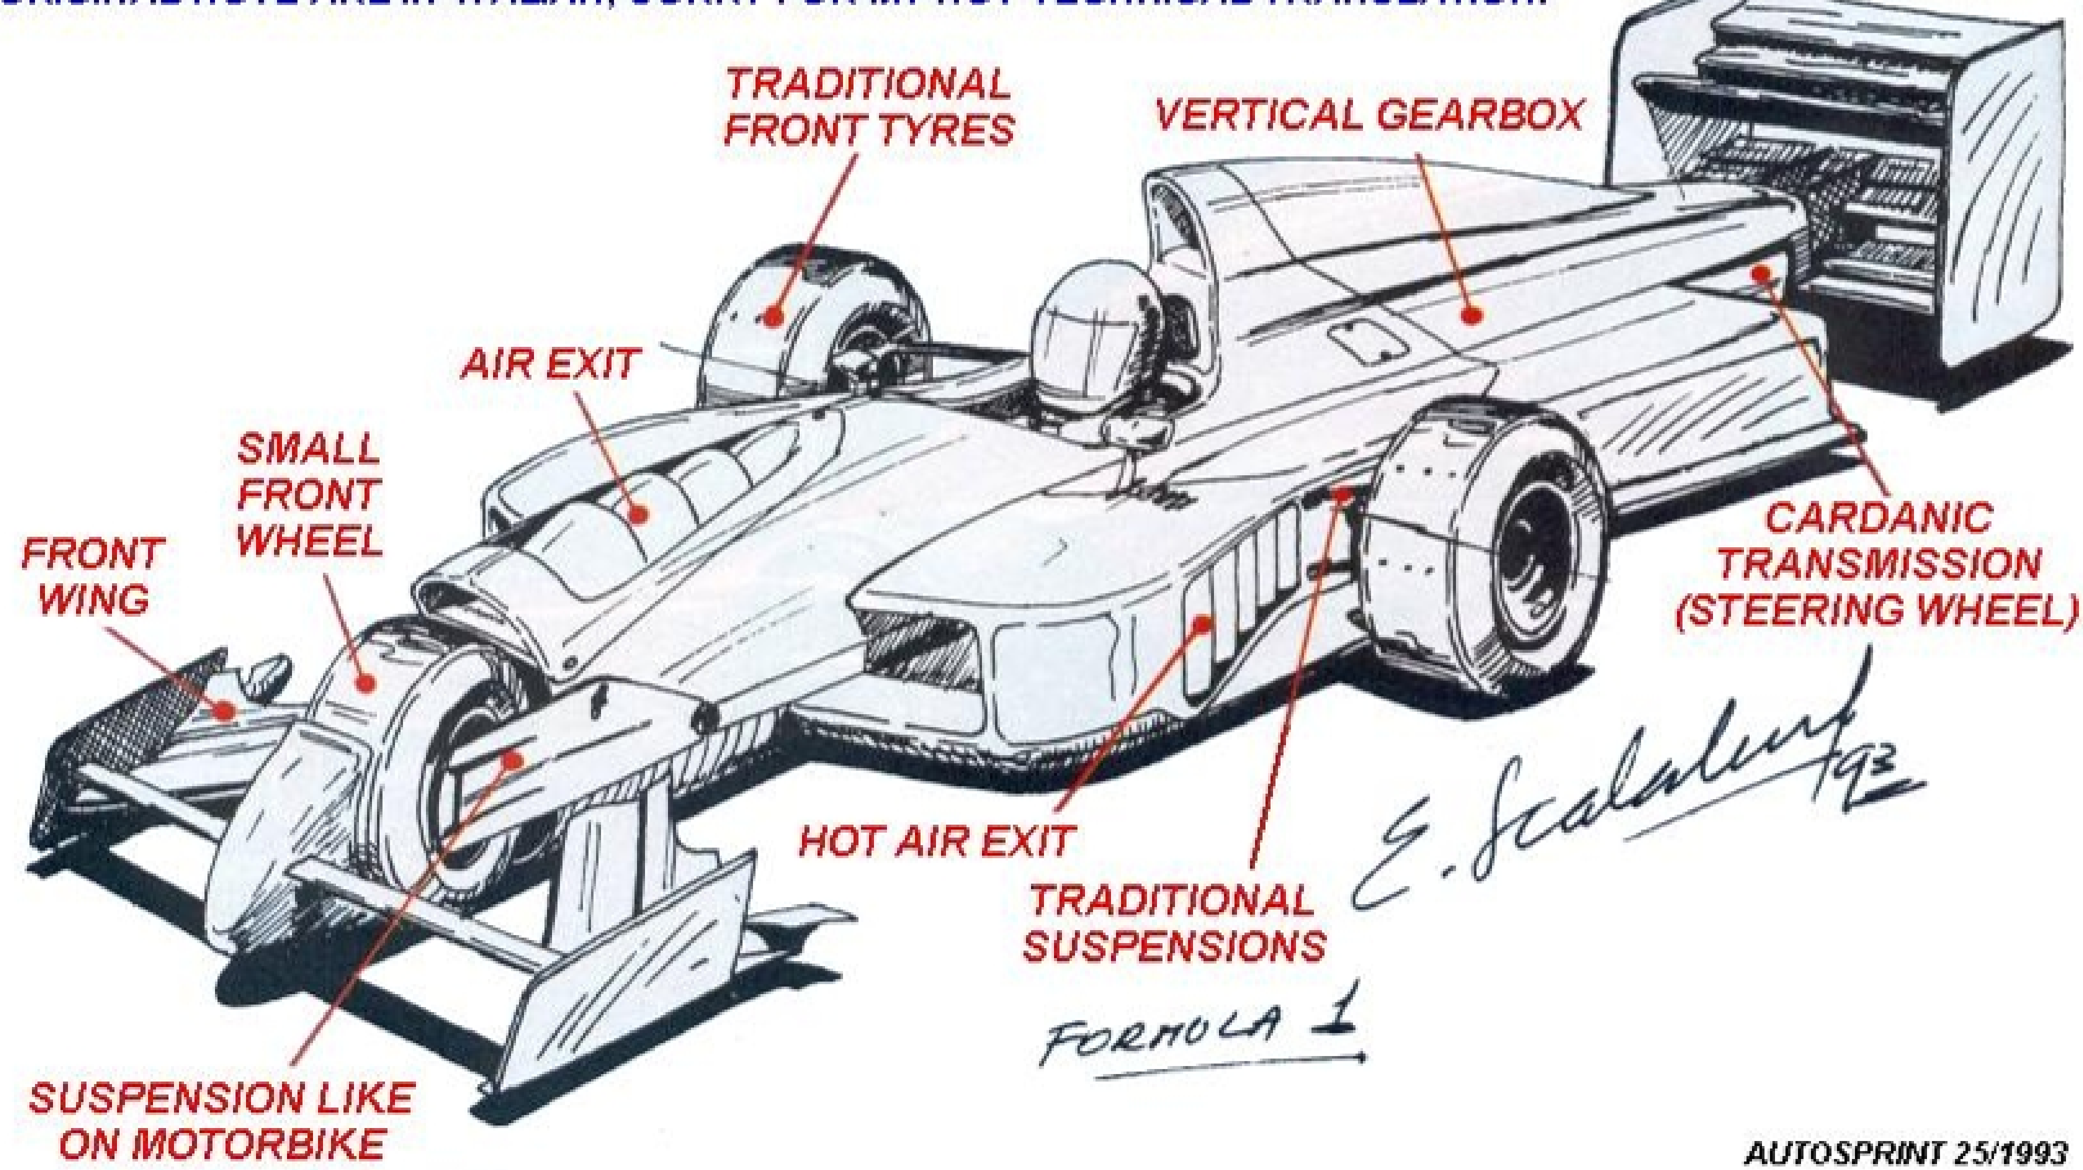

AUTOSPRINT 25/1993

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2007

Revised 4/5/07

This document has been revised to correct an error it contained as originally published.

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL STEER AND ANTI-SQUAT/ANTI-LIFT IN STRUT SUSPENSION

As a reader of Racecar Engineering, I would like to know your opinion about a concern I have at the moment. I hope you will consider it and will be able to include it in a future issue of RE, because I think it is quite interesting.

I am studying the suspensions of my rally car at the moment. Our rally car is a WRC car that we use for both gravel and tarmac events. There are some small differences between tarmac and garvel spec suspensions but the main layout is the same. It is a 4WD, with an active hydraulic center differential and a 50/50 torque split, most of the time.

I have a lot of toe variations during full travel at the ***rear*** *suspension (pseudo MacPherson type, see attached file).* [Newsletter subscribers will receive the questioner's illustration as a separate attachment.] *In order to solve the problem, I have been testing (real tests and computer simulations) various configurations, with varying mounting point positions.*

The results give me the same direction. The toe variations during travel decrease when points B and G are at the same height (vertical axis). We have always used our car with setups where B was around 5 to 10 mm lower than G. Knowing the results from my tests, I would logically change this setup to a line BG nearly horizontal.

However, some people are saying that it will greatly alter the anti-lift/anti-squat feature of the rear suspension. I have the feeling that it's not a problem but I am unsure about that. The problem is that I don't really know how to simply calculate the "anti" amounts of a MacPherson suspension in general and I can't find anything in many books. Could you give me some indications about this?

Readers are referred to the questioner's illustration. We are discussing the rear suspension, which is shown in the right half of the illustration. The suspension has a telescoping strut located by a single pivot at its top end, and by three links at its bottom end. Links CB and FG run transversely and provide lateral location, controlling camber and toe. Link AM runs longitudinally and provides longitudinal location and reacts brake torque. Link FG is partly hidden by the anti-roll bar. Point C is hidden behind a donut-shaped object whose function is not obvious to me. Point B is adjustable for height.

Any rear suspension that has a strut of this type, with the lower end of the strut located by some system of arms or links, is sometimes called a MacPherson strut suspension, or sometimes a Chapman strut suspension. MacPherson was an engineer for Ford in England, and he is credited with inventing this type of suspension. However, he and Ford only applied it to the front end of the car. Lotus founder Colin Chapman was the first to apply it at the rear, although arguably his idea is a fairly obvious application of MacPherson's concept. So there is a rational case for either name when referring to a rear suspension, and I consider them legitimate synonyms. (I would not, however, consider it correct to call a strut-type front suspension a front Chapman strut.)

To understand the behavior of this suspension, and indeed any suspension, we need to define two more points: the wheel center and the contact patch center or load centroid. These points are not in the illustration, but they are easily described verbally. The wheel center is the intersection of the wheel centerplane (the plane midway between the rim flanges) and the wheel's axis of rotation. I will call this point J. For the contact patch centroid, I like to use a point directly below point J, even when the wheel has some camber. Let's call this point K. Let's also suppose that the wheel has no toe-in or toe-out at the condition we're evaluating. The plane containing points J and K, and exactly longitudinal to the car, forms a good effective wheel plane to which we can project the suspension's side-view geometry. Being a vertical, longitudinal plane, it is in true shape in a side view of the car.

All of these points have three coordinates in the car's axis system: x (longitudinal, per SAE convention), y (transverse), and z (vertical).

The distance from point K up to point J is the tire's loaded radius r_L.

All we really need to know about any suspension's geometry to determine the suspension's "anti" properties is the motion path of point K. More precisely, we need to know the instantaneous side-view and end-view (front- or rear-view) slope of that motion path at the suspension position we're evaluating, i.e. the instantaneous rate of change of the point's x or y coordinate with respect to rate of change in the z coordinate – in calculus notation, dx_K/dz_K for anti-squat and anti-lift, and dy_K/dz_K for anti-roll.

With good enough modeling software, it may be possible to simply move the suspension a very small increment each side of the position being evaluated, and get a $\Delta x_K/\Delta z_K$ that very closely approximates a true dx_K/dz_K.

Note that there are two assumptions we can make when doing this: we can assume that the wheel is free-rolling, meaning that K always stays directly below J, and the motion path of K is J's motion path transposed to the ground; or we can assume that the wheel is locked to the upright, meaning that any change in the side view angle of the upright as the suspension moves creates a difference between the motion path shapes for J and K. With independent rear suspension and outboard brakes, we use the former method for anti-squat (under power), and the latter method for anti-lift (under braking).

We need to know three more things to get a percent anti-squat or anti-lift value: we need to know the percentage P_r of total x-axis force at the contact patch (or rear contact patch pair, if we are assuming the suspension is symmetrical and we are seeking a percent anti for the rear wheel pair); we need to have a sprung mass center of gravity height H (some estimating is normally required to arrive at a reasonable number); and we need to know the car's wheelbase L.

Then, for anti-squat, percent anti-squat = $(dx_K/dz_K)P_rL/H$. For the situation where all four wheels are driven and the torque split is 50/50 front/rear, percent anti-squat = $(dx_K/dz_K).50L/H$. Remember that for independent suspension the value for dx_K/dz_K is taken assuming the wheel to be free-rolling, so that $dx_K/dz_K = dx_J/dz_J$.

For anti-lift, we need to know the percentage of brake retardation force, at the ground, that the rear wheels provide. We use this for P_r . We use the dx_K/dz_K value for the condition where the wheel is assumed to be locked to the upright. That is, $dx_K/dz_K = dx_J/dz_J$ only if the upright does not rotate at all in side view as the suspension moves.

Many of us do not have computerized animation of our suspension layouts. We are used to finding the side view instant center, on the drawing board or the CAD tube, and constructing a side-view force line from the instant center to the contact patch center or the wheel center. To do the job this way with a strut suspension, we need to find the virtual upper and lower control arm planes, and find their intersections with the wheel plane. The intersection of two planes is a line, so we are talking about two lines, lying in the wheel plane. These lines are our side-view projected control arms. Their intersection is our side-view instant center.

A strut-type suspension has a virtual upper control arm plane, which contains the strut's top pivot point (point E, in the questioner's suspension), and is perpendicular to the strut tube axis.

This is not necessarily the same as the effective steering axis. When the strut has a single ball joint at its lower end, the effective steering axis is a line through the ball joint pivot point and the upper pivot point.

What to use for a lower control arm plane is more of a conundrum. I think I would start by finding the points where lines BC and FG intercept the axle plane: the transverse, vertical plane containing the wheel centers. Let's call these intercepts S and T. A line connecting these, line ST, would be the end-view projected lower control arm.

Once we have line ST, we can establish a plane parallel to that line, and containing line AM. Once we have that, we can find the intersection of that plane with the wheel plane, and use the line so defined as our side-view projected lower control arm.

I am suggesting using side-view and end-view projected control arms that do not actually lie in a common plane. However, all things considered, I think this is preferable to any possible alternatives. The suspension, so modeled, is probably not exactly equivalent to reality, but it is a reasonable approximation.

Once we have the upper and lower side-view projected control arms, we can find their intersection, which is our side-view instant center. We then draw in the force line for braking, from the contact patch center to the side-view instant center. We locate a point at the sprung mass center of gravity at height H. We draw a vertical line in our side view at a distance P_r/L forward of the rear wheel center. I call this the resolution line for braking.

The force line will intercept the resolution line at some height h_1. The percent anti-lift = h_1/H. Note that if the force line intercepts the resolution line below ground, h_1 becomes negative, and we have negative anti-lift, also called pro-lift. If $h_1 > H$, we have more than 100% anti-lift, meaning the rear suspension will compress in braking. As we add more rear brake, we get more rear anti-lift (or more pro-lift, if the geometry provides pro-lift), without changing anything else.

For the anti-squat, the side-view projected control arms and instant center are the same as for anti-lift, but the resolution line and force line are different. The force line is different because with independent suspension, drive torque does not act through the suspension linkage. The resolution line is different because the rear wheels do not generate the same percentage of the longitudinal force under power that they do in braking.

There are two variations on the basic technique. Done correctly, they both give the same answer. The variation I prefer involves constructing a force line that originates from the contact patch center (point K), and using the same equation or rule for calculating the percent anti that we use when the suspension linkage reacts the torque, as with a live axle under power, or with outboard brakes in braking. The other variation involves constructing a force line that originates at the wheel center (point J), and using a slightly different equation or rule.

In the first method, we construct a line from the wheel center J to the side-view instant center. Then we construct a parallel line to that one, originating at the contact patch center K. This is the force line when the car is under power. We construct the vertical resolution line .50L forward of points J and K for the questioner's example, or at the front axle for a purely rear-drive car. The force line intercepts the resolution line at a height h_2. The percent anti-squat = h_2/H.

In the second method, we do the same thing, except we omit the second line, and just use the line from point J to the instant center as the force line. The force line then intercepts the resolution line at a height h_3, and $h_3 = h_2 + r_L$. The percent anti-squat = $(h_3 – r_L)/H$. That is, of course, the same as h_2/H.

Adherents of the second method sometimes insist that this is the right way because the force acts on the car at hub height, and that any force line must pass through the instant center. I say the first method makes more sense because any force exerted on the car by the road must act where the road touches the car, and the wheel is part of the car, not a separate entity that acts on the car. The difference between independent and live axle suspension, or inboard and outboard brakes, is not where the force acts on the car, but how the torque is reacted within the car – i.e. through the suspension linkage or not. I also think it makes sense to construct force lines so that a common rule applies once the force line is constructed, rather than having two different rules. However, I don't think the whole debate is of much consequence, since both methods give identical answers. The difference really has more to do with visualization and semantics than actual natural laws. It is, however, important to understand both methods, so that one doesn't apply the rule or equation from one method when using the other.

The questioner wants to know what effect bump steer adjustment will have on the anti properties. It is customary to model our antis on the assumption that there is no bump steer, and treat bump steer as a separate issue, a steering geometry issue rather than a suspension geometry issue. However, if we really want to be rigorous, in fact there is some effect, but it is small, unless there are really large amounts of bump steer, in which case the car will be so undriveable that optimizing the antis will not be our main concern. Unless bump steer causes the wheel to steer about an axis precisely in the wheel plane, there will be a small component in the side-view motion paths of points J and K resulting from the bump steer.

If the questioner's suspension is configured for minimum bump steer, point B will lie in plane CFG. If B lies below plane CFG, the wheel will toe in as the suspension compresses, and the car will have roll understeer. If B is above plane CFG, the opposite will occur, and the car will have roll oversteer. Assuming that the wheel centerplane is outboard of the effective strut axis, roll understeer will increase anti-squat and anti-lift slightly, and roll oversteer will decrease anti-lift and anti-squat.

Note that this is opposite to the effect we would predict if we supposed that the suspension's side-view properties depended on the side view inclination of line BG, as some older suspension texts would suggest. Raising point B decreases the anti-squat and anti-lift. This example demonstrates the importance of looking at what happens at the wheel plane, rather than looking at the side-view inclination of the inner pivot axes of the control arms.

One final, somewhat self-serving note: when discussing suspension antis, I generally try to get in a plug for my video "Minding Your Anti", which includes more discussion of the topic, with more pictures. They are available in US standard VHS cassette only, for US$50.00, payable by check or money order to Mark Ortiz, 155 Wankel Dr., Kannapolis, NC 28083-8200, USA. Price includes shipping and handling, worldwide.

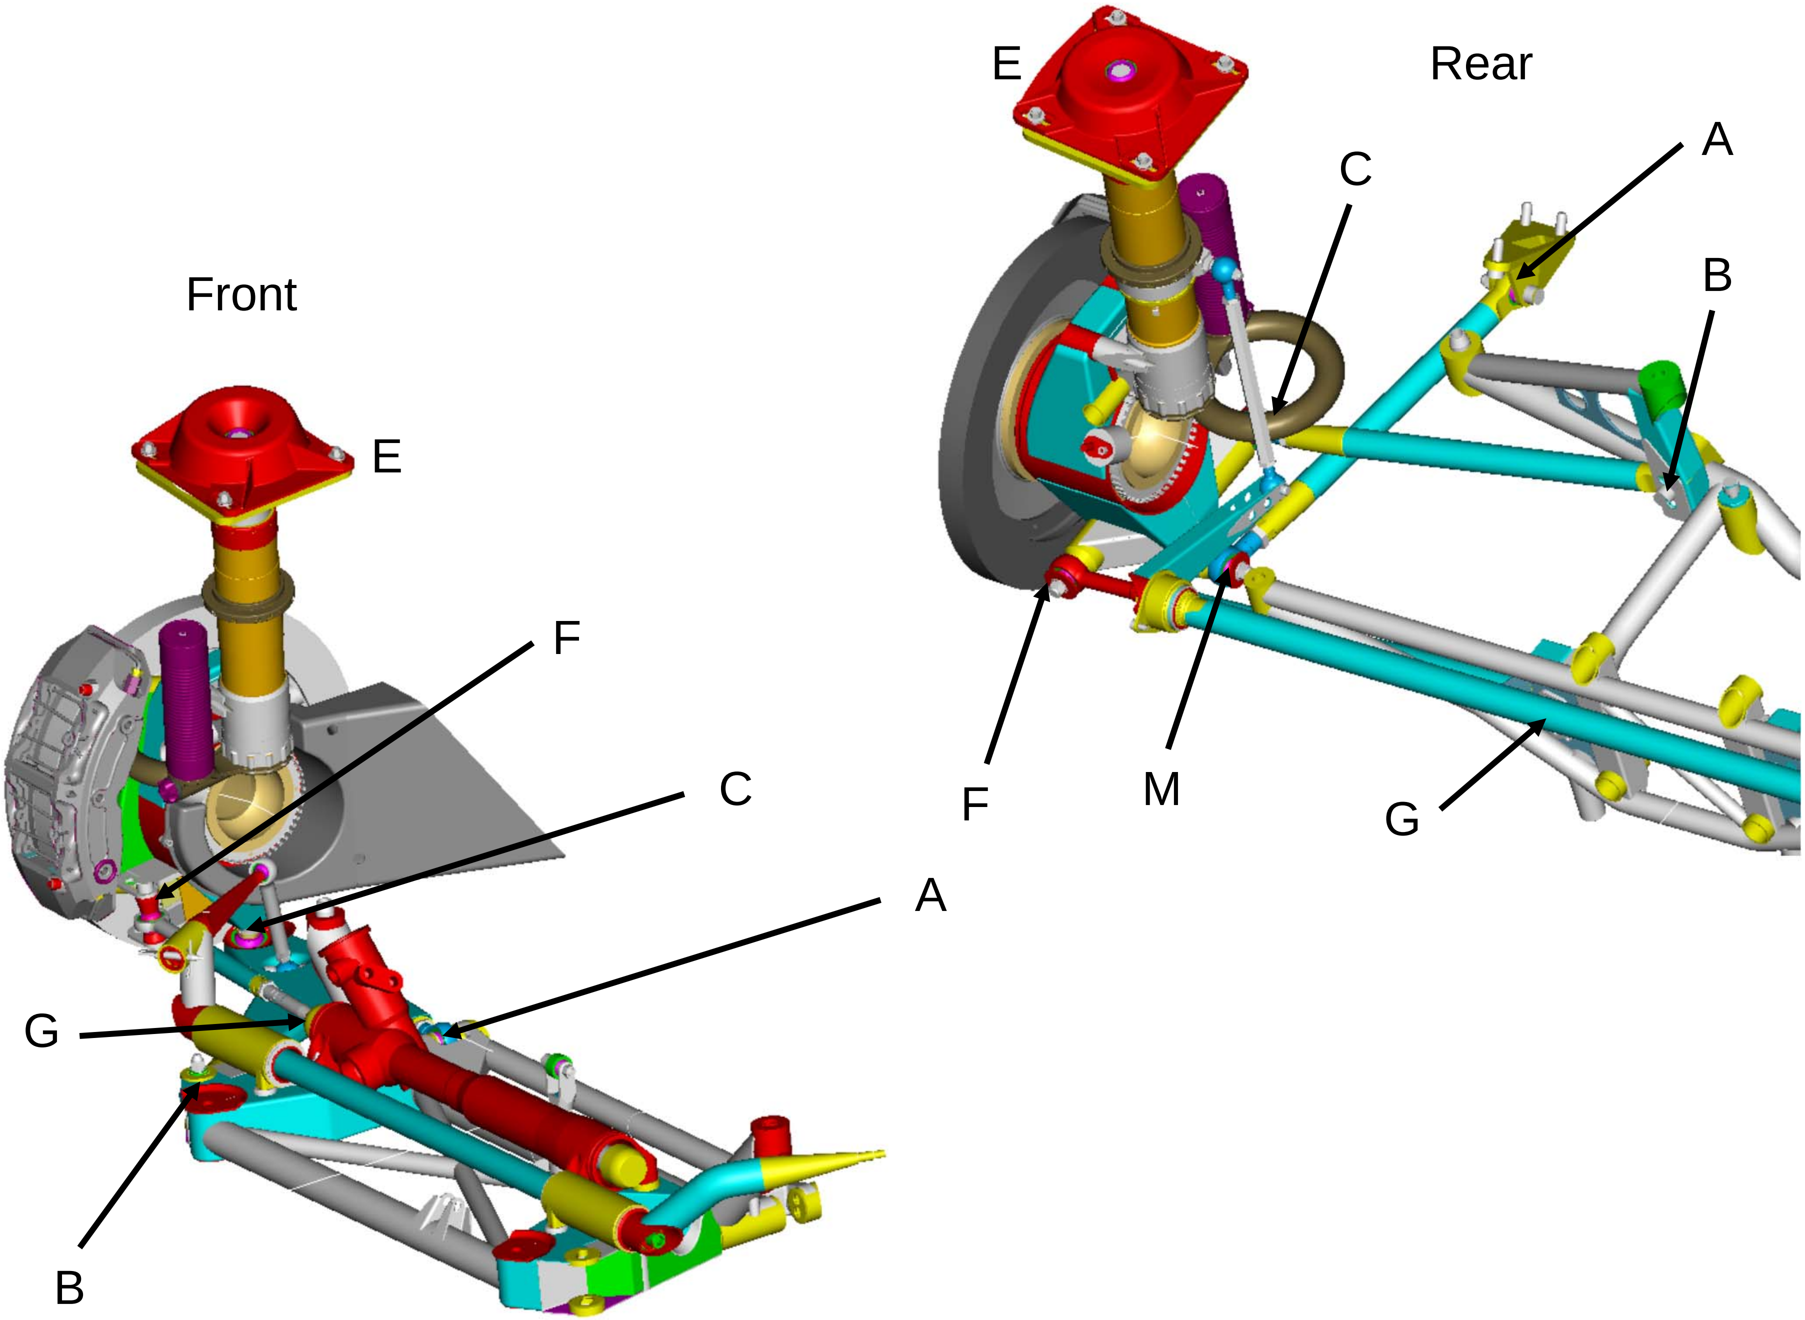

Front
E
F
C
A
G
B
Rear
E
C
A
B
F
M
G

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL AXIS FOR LIVE-REAR-AXLE ROAD/HILLCLIMB CAR

I've been reading your column in Racecar Engineering for several years now and I'm gradually increasing my knowledge of the dynamics of a suspension system, slowly but surely.

One thing I haven't been able to find much about anywhere is front roll centre height on a car with live axle rear suspension and independent front. The car in question uses a panhard rod and 4 parallel leading links to locate the axle. With the current wheel / tyre combination the centre of the diff /roll centre height is approx. 11.9 inches. The front suspension is double wishbone, and I have complete freedom to do whatever I want with the front geometry.

Most things I been told, or read, suggest that the front roll centre should be close to the ground, but surely on a car with a live axle rear this will cause a steep inclination of the roll axis.
The weight of the car is around 2100lbs, distributed approx. 49/51% F/R, car is front engined.

Can you offer any advice as to a suitable starting point from which to base my design?
The car is to be used for hillclimb and sprints, but will still be a road legal car so ground clearance is an issue (otherwise I'd be using a Woblink at the rear).

It is true that the roll axis will be steeply inclined if the front roll center is low and the rear roll center is high, and this is indeed the norm in cars with independent front suspension and beam axle rear suspension. It is usual to make up for the low front roll center by adding front roll resistance with an anti-roll bar.

In a car with close to equal front and rear weight distribution, and equal tires sizes front and rear, we will need close to equal total lateral load transfer front and rear. That implies somewhat less total roll resistance in the rear suspension than the front, because at the rear there is a substantial portion of the load transfer that does not act through the suspension. I am referring to the unsprung component of the load transfer, the portion that comes from the mass of the axle exerting a centrifugal inertia force at a height above the ground. There is also an unsprung load transfer at the front, but it is much smaller because there is much less unsprung mass, and even the unsprung mass that is there partially acts as sprung mass for purposes of lateral load transfer, because it moves laterally with the sprung mass in roll.

In addition to the unsprung load transfer, there are three other components to load transfer: elastic, geometric, and frictional. All of these result from the action of the suspension, and may be said to act through the suspension.

The elastic component comes from the anti-roll moment generated by the springs and anti-roll bars. The geometric component comes from the anti-roll moment generated by the rigid members of the suspension: the links, uprights, and axles. The frictional component comes from all the frictional forces in the suspension system. Frictional forces include the mostly unintentional ones in the springs and pivots, and the mostly intentional ones in the dampers.

Cars with independent front suspension and beam axle rear suspension generally have small geometric components and large elastic components in the front load transfer, and large geometric components and small elastic components in the rear load transfer.

Independent suspensions cannot work well when they have large amounts of geometric roll resistance (i.e. a high roll center or large amounts of lateral anti). If we build an independent suspension that way, it will jack up when cornering. The reason for this is that geometric roll resistance in independent suspension consists of linkage forces that try to extend the suspension of the outside wheel (net upward jacking force at the outside wheel) and compress the suspension of the inside wheel (net downward jacking force at the inside wheel). The outside wheel carries the greater load, so the outside tire has more grip and generates greater lateral force. When the lateral force is greater, the jacking force is greater. Thus, the upward jacking force generally will exceed the downward jacking force, and the net jacking force for the wheel pair will be upward.

To avoid this, the roll center of an independent suspension system must be low. There is no hard limit, but as a general rule anything above four inches is too high, and two or three is more prudent, because the anti-roll generally increases as the suspension extends, and we know that the car will sometimes be going over crests.

We also need to design our suspension so that we do not get any large camber changes. With independent suspension, we cannot reduce camber change in roll without increasing it in ride, so the best we can do is avoid any huge amount of change in either ride or roll.

To get good camber properties and an appropriate amount of geometric anti-roll in a front suspension for a car with a beam axle rear end, I generally recommend a front view instant center between 55 and 85 inches from the wheel horizontally, and between five and ten percent of that distance above the ground. Make the control arms as long as possible, consistent with packaging and bump steer constraints, with the front view projected upper control arm around 2/3 the length of the lower one.

At the rear, there is no need to have the Panhard bar at axle height. If we are dealing with an existing rear suspension, and we don't want to change it, we can make the car work reasonably well with any Panhard bar height by choosing appropriate spring and anti-roll bar rates.

However, there is a case for having the Panhard bar as low as ground clearance and mounting constraints will permit, particularly for a road car or a road racing car, and accepting the need to increase the elastic roll resistance accordingly. The primary reason this strategy works is that it reduces torque roll and torque wedge: body roll and wheel load change due to driveshaft torque reacting through the suspension. With a live axle, driveshaft torque rolls the sprung mass to the right, unloads the right rear and left front tires, and correspondingly adds load to the left rear and right front tires. In oval track terminology, driveshaft torque adds wedge to to the chassis. In my application of the terminology, it wedges the car for left turns, and de-wedges it for right turns. That makes the car looser (adds oversteer) exiting right turns, and tightens the car (adds understeer) exiting left turns. It also makes the right rear tire spin prematurely under power if the differential does not lock, or makes the car try to turn right if the diff is locked.

If we add elastic roll resistance at both ends of the car, we reduce torque roll but not torque wedge. If we increase the rear elastic roll resistance, leaving the front end unchanged, we reduce both torque roll and torque wedge. However, we can only do this if we reduce the rear geometric roll resistance. For this reason, the fastest live-axle road racing cars place the rear roll center as low as packaging constraints allow, and use rear anti-roll bars.

To some extent, we can compensate for torque wedge. We can set the car up with less than 50% static diagonal (RF+LR) percentage, use somewhat stiffer springs on the right side than on the left, and use a bit more anti-dive in the right front suspension than at the left front. If we apply these crutches in a suitable combination, we can get reasonably good behavior in right and left turns, and under power and braking. However, we can get away with applying these crutches in smaller measure, or get better car behavior without them, if there is less torque wedge to begin with.

The questioner mentions the WOB link (from Watt-Olley-Bastow) suspension. This refers to a mechanism that approximates straight-line motion for small displacements. It consists of a rocker and two links, like a Watt linkage. It differs from a Watt linkage in having the pivot on the rocker outside the attachment points for the links, rather than between them, with both links extending in the same direction away from the rocker, rather than opposite directions. It can produce approximately straight-line motion at the pivot if the links have unequal length, in the correct proportion to each other. It provides a way to get a low roll center when the only frame members we have available to anchor to run above the axle. It has the disadvantage that the link loads tend to be very high. We are still unable to get a roll center lower than the lowest part of the mechanism.

With a Mumford linkage, it is possible to get a roll center somewhat below the lowest part of the mechanism. In fact, it is possible to get a roll center dramatically lower, except that then the roll center then moves up and down a great deal as the suspension moves. A Mumford linkage has two rockers and three links.

It is possible to make a Watt linkage approximate the behavior of a Mumford linkage. We can put the rocker under the differential, lying flat, and at an angle when seen from above. We run the links upward toward the frame from the rocker. The roll center is approximately at the height where the two link centerlines intersect in front or rear view. Careful detail design is necessary to avoid running the rod ends on the links out of travel.

ROLL AXIS OF THE AXLE

I have read about how to find a rear axle's axis of rotation in roll. On a triangulated four-link system, for example, you find the intersection of the lower link centerlines (usually ahead of the axle and above the links), then find the intersection of the upper link centerlines (usually behind the axle and above the lower link intersection), and then connect those two points with a line. The axle moves about this line in roll. The point where this line intercepts the rear axle plane is the rear roll center.

It would seem obvious that the inclination of this line must determine the rear axle's roll steer properties also. If the axis slopes down toward the front, that would seem to imply roll understeer. Yet when I apply the suspension geometry program on my computer to such a layout, it tells me the system has roll oversteer. Am I crazy? Can the computer be wrong?

Note that the axis the questioner is referring to here is different from the roll axis of the car as a whole, which was the subject of the previous discussion. This is the axis about which the axle rotates with respect to the sprung mass.

Yes, the computer can be wrong, and I have seen similarly anomalous outputs from suspension geometry programs regarding bump steer and roll steer properties in independent systems.

For example, consider the case of a front independent suspension with the steering rack ahead of the axle line. If you raise the rack, leaving all else unchanged, you know that the wheels will toe out more, or toe in less, in bump as a result of this change. It always works this way, for intuitively obvious reasons. Yet I have had the same computer program the questioner says he's using tell me the opposite. I am leaving the name of the program out of the discussion, as the person who sells it is a friend of mine. But I will definitely say that having a computer is no substitute for understanding things with your own brain, and that computer outputs definitely can be wrong. (The computer can also be right when you think it's wrong. The computer is not infallible, nor are the people who write the programs, nor are you, nor am I. We all do the best we can.)

Returning to the triangulated four-bar, yes it generally does have roll understeer, despite the up-at-the-front slope of the lower control arms in side view. The unloaded or inside wheel moves rearward in roll, and the outside or loaded wheel moves forward. This is due to the arms' in-at-the-front angularity in plan view, combined with the lateral motion at their rear ends in roll, which results from their being well below the roll center.

If the lower arms were parallel with the vehicle centerline in plan view, and sloped up at the front in side view, then the system would have roll oversteer. It would also have a rear axle roll axis sloping upward toward the front, not downward. The axis would pass through the upper arm intersection, as before, and it would be parallel to the lower arms in side view, since they are parallel to each other, meaning they have no intersection, or, to be slightly incorrect, an intersection at an infinite distance.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANOTHER TRIANGULATED FOUR-LINK AXLE

I am using a very light (aluminum) rear beam axle on a front wheel drive autocross car. I have read your discusssion of live rear axle location by Watts, Mumford and others but I am using a four link location system without a Panhard or Watts type system. I have the two lower links wide based at the front and meeting at the middle of the axle. They are about 3" off the ground and basically horizontal, thus lateral location of the axle is about 3 inches off the ground. My question is where is the roll center? I once thought it had to be at the level of the lower links, but perhaps not.

The upper links go straight back about 9 inches above the lower links and have provision to move the forward attachment points up and down to induce a bit of rear wheel steer in cornering (not currently doing this).

In this case, the rear axle roll axis, or rear axle axis of rotation in roll, is a line passing through the lower link centerline intersection, and parallel to the upper links. The roll center is the axle plane intercept of this line, as with any rear axle roll axis. If the lower links met at the exact center of the axle, the upper link angle would not change the roll center at all. If their attachment points to the axle are very close, so that their intersection is a short distance behind the axle, then upper link angle affects the roll center, but only a little. As the lower link centerline intersection gets further behind the axle, the upper link angle has more influence on the roll center height.

In all cases, if all the links are horizontal, the roll center is indeed at the height of the lower links, and the rear axle axis of rotation in roll is horizontal. True to this model, there is no roll steer. If we raise the front of the upper links, the rear axle axis of rotation slopes upward toward the front. Again true to this model, we have roll oversteer.

Anti-squat (in forward acceleration) is zero in all cases, because the rear wheels are not driven and can make no forward force. Anti-lift in braking is affected by the side-view link angles, however. If the lower links are horizontal and the upper links slope upward at the front, the rear suspension has pro-lift in braking. This is not necessarily a serious problem, but we can have some anti-lift, and still have any desired roll steer and geometric anti-roll (roll center height), if we make the lower links slope upward toward the front.

LARGE VS. SMALL TIRE AND WHEEL DIAMETER

I thoroughly enjoyed your article "Tall vs Short sidewalls" in the April 2007 issue of Racecar Engineering magazine. In that article you discussed about the benefits of short sidewall tyres fitted on larger diameter rims. So, you explained how beneficial (or not) would be a different configuration of two wheels with exactly the same overall diameter. I would like here to enter another parameter to the discussion. What about two wheels with exactly the same tire sidewall characteristics but with different overall diameter? For example, let's say we have a 15" overall diameter (not rim diameter) wheel (rim + tire) and a 17" overall diameter wheel. And now let's assume that the sidewall characteristics of the two tires are exactly the same. Also the widths of the two tires are the same.

Which one do you think will be more beneficial to overall handling? It is quite sure that we will have two footprints of different shape (I am not sure if they will have a different area though).

I think that with the low wheel we will have a more round footprint shape, while with the tall wheel the print will be more long and narrow. Do you think this will have any effect to the centroid of the lateral force with consequences to pneumatic trail and aligning torque? What do you think?

15" and 17" are pretty small for tire outside diameters, but the question is really the same regardless of what hypothetical sizes we might posit.

As with tire width, we need to define our basis for comparison, and this is not as straightforward as one might suppose. Do we hold inflation pressure constant? Static deflection or tire vertical spring rate? Do we assume identical tread compound? Do we compare fully optimized cases, which probably implies different inflation pressures, static deflections, and tread compounds? One can make a rational case for any of these approaches. In an actual design or modification situation, availability of tires and compounds may constrain us. Since we are examining the principles here, we will look at the question from a variety of angles.

What happens if we hold tire pressure constant? In theory, the contact patch should stay the same area at a given load as the tire diameter increases, just as it should with increasing tire width. However, with a wider tire, the contact patch gets wider and shorter. With a larger diameter, the contact patch stays the same length and width. In both cases, the tire's static deflection decreases, and its vertical spring rate increases, as we add size. In both cases, tensile loads in the carcass increase. A tire is approximately round in section, so when its sectional circumference increases, the hoop stresses acting transversely and radially increase. That increases cord tension when tire width increases. A tire is also round in side view, so it also has hoop stresses acting circumferentially in side view. As we increase pressure, we can measure the increase in diameter and circumference.

This comes from the circumferential strain that goes with the circumferential stress. Tire unloaded circumference is a measurement we commonly take, especially in oval track racing, to control tire stagger. Road racers often ignore this measurement, but they shouldn't. You don't necessarily want stagger on a road course, but you don't want to have it unintentionally. It is a bit harder to compare tires of different dimensions at the same pressure, but we know that a larger-diameter tire has more surface area for its pressure to act upon, and it therefore has to have greater cord tension at a given pressure.

This means that the larger-diameter tire acts stiffer, at a given pressure, just like a wider tire does. Therefore, we will probably end up running lower pressure in the larger-diameter tire. Then the contact patch will get longer.

What happens if we set the tire pressure such that the larger tire and the smaller one both have the same static deflection – say, ½"? Taking the questioner's 17" and 15" unloaded diameters as our basis, for a 17" diameter a ½" static deflection theoretically gives us a contact patch 5.74" long. For a 15" diameter, ½" static deflection theoretically equates to a 5.38" contact patch length. That's a contact patch length ratio of 1.067, compared to a diameter ratio of 1.133. The contact patch length does not quite go up proportionally to the diameter; it goes up by roughly half the percentage. ½" is a fairly large deflection for a tire that small. The contact patch length in both cases is more than 1/3 of the tire diameter.

It turns out that for all reasonable diameters and static deflections, the amount of non-linearity does not change much. If we doubled both the diameters, and calculated theoretical contact patch length at ½" deflection, we still get a contact patch length ratio of 1.065 – not much change from 1.067. Cut the deflection to ¼", and the length ratio becomes 1.064 – still about the same.

My conclusion is that within the range of diameters and static deflections we are likely to consider for a particular car, contact patch length for a given static deflection has a surprisingly linear relationship to outside tire diameter: the contact patch length ratio stays at about .94 times the tire diameter ratio. The percentile increment of contact patch length gain remains roughly half of the percentile increase in diameter.

Using the simplifying assumption that, for constant contact patch width, contact patch length is simply inversely proportional to inflation pressure, the 17" tire should need only 1/1.067 times as much pressure and should have about 1.067 times the contact patch area, for the same static deflection and vertical spring rate. At this somewhat lower pressure, it will have around 1.062 times the side-view circumferential hoop stress and around 1/1.067 times the sectional circumferential hoop stress – meaning that overall cord tension and tire rigidity should be similar despite the reduced pressure.

More contact patch, similar rigidity; that looks good. Of course, there is a weight penalty, just as there is if we widen the tire. But overall, the bargain looks surprisingly similar.

One might think there would be a disproportionate penalty in rotational inertia, since the mass of a larger-diameter tire not only is greater, but also acts on a larger radius of gyration. However, a larger-diameter tire also turns at a lower rotational speed for a given road speed, and rotationally accelerates at a lower rate for a given longitudinal acceleration. The linear speed at the tread of any tire has to be approximately equal to the road speed of the car. Correspondingly, the linear acceleration of the tread has to be approximately equal to the linear acceleration of the car. So the rotational inertia penalty is roughly proportional to the weight gain, but not significantly greater.

There is a further weight penalty in the wheel. This is somewhat variable depending on wheel design, but in most cases it's safe to say that we add more weight to the wheel center and rim by increasing the wheel diameter than we add to the rim by increasing the width.

As a broad generalization, it is probably true that we can increase the contact patch area more weight-efficiently by adding width than by adding diameter. However, the difference is not dramatic, and there are advantages to using diameter. One big advantage is that when we add contact patch area by adding diameter, the tire does not become more camber-sensitive, as it does when we add width. There is generally an an advantage in resistance to aquaplaning when we go tall instead of wide, and better performance in snow.

Depending on how we play other tradeoffs, with a taller tire we may very well be able to run a softer and/or hotter-running tread compound.

As for self-aligning torque and pneumatic trail, if the contact patch is longer, those will increase. The centroid of lateral force should move rearward in the contact patch, at least for moderate slip angles. It should also move forward more as the tire approaches the limit of adhesion, increasing self-aligning torque falloff as the limit of adhesion draws near. Whether that improves steering feel or not is a matter of driver preference and the design of the rest of the car.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE THOUGHTS ON ZERO-DROOP SETUPS, AND RELATED OVAL TRACK SPRING SPLIT ISSUES

In the February 2007 newsletter, I responded to a question regarding zero-droop setups. These are setups where the suspension is not allowed to have any droop travel from static position. I said that these were bad for everything except controlling ground clearance in order to have consistent aerodynamic ground effects. As a result of further reflection, and conversation with a number of people, I have concluded that there is another benefit.

That benefit is that when the car is in a rolled condition, the inside wheel pair has greater pitch resistance than the outside wheel pair. This has effects similar to left-stiff spring splits on an oval-track car. The car de-wedges when accelerating rearward while also accelerating laterally, and it gains wedge when accelerating forward while also accelerating laterally. That helps the car get itself rotating in yaw on entry, provided the driver is decelerating or braking while entering, and helps it put power down on exit.

This would apply to a car with zero-droop setup at both front and rear. If only the front is zero-droop, effects would be confined to corner exit.

This benefit does come at a price in the wheels' ability to follow bumps, but if the surface is smooth enough, it may be worth it.

A similar effect, but more subdued, can be had by making the wheel rate increase in droop, either using rocker geometry or using snubbers. Note that this does not mean that the spring load or force increases in droop; it means that the force decreases at a greater rate with respect to droop displacement, as droop displacement increases.

A third way of creating a rising rate in droop is used on dirt Late Models, and could also be applied to road racing cars, although I have yet to hear of it being tried for road racing. This third way is to have two springs stacked on top of each other on a coilover, separated by a slider. The slider is arranged to top out against an adjustable collar on the coilover threads, at some point as the suspension extends. This adjustable collar has a smooth sleeve above the shoulder that the slider seats against. This provides a smooth surface for the slider to ride on, and protects the threads. These devices are usually used on the left front in oval racing.

The rate increase occurs in a single step. It would be possible to create two smaller steps by using two sliders and three springs, but I have never seen this.

Elsewhere in oval racing, one still often encounters the erroneous belief that the way to tighten the car (add understeer/reduce oversteer) on exit is to use a right-stiff front spring split. This belief stems from misunderstanding of the relationship between spring rate and spring force, in a situation that causes an extension displacement of the spring.

To understand this sometimes confusing concept, imagine a car on the drag strip, with the suspension on the left front wheel locked solid (spring rate effectively nearly infinitely large). When the car launches, the locked suspension will be unable to extend, and it will be quite easy to lift that wheel off the ground, while the right wheel, whose suspension can extend, will stay on the ground. All load remaining on the front wheels will be on the right front, and the car will have more than 50% of its weight on the right front and left rear.

The more we increase the spring rate, the more closely the suspension approximates a locked condition. The greater the spring rate, the more the load increases with compression, and the more the load **decreases** with extension. Consequently, a stiffer inside front spring increases load on the outside front and inside rear tires, and tightens the car (adds understeer) on corner exit.

HIGH OR LOW PANHARD BAR?

Can you discuss the advantages and disadvantages for various heights of Panhard bar? For example, on my stock car, I could run the Panhard bar as high as 12½", or I could run it as low as 10¾". In either case, I could get the car balanced by changing the combinations of springs, crossweight, etc.

In many cases, there is an additional variable: rear anti-roll bar stiffness. In many stock car classes, rear anti-roll bars are now prohibited, but they are still used in the top NASCAR divisions, and in many stock-car-related road racing classes. So we have at least four variables that play off against each other: Panhard bar height, springs, diagonal percentage, and anti-roll bar if allowed.

To start off, let's take a simple case: we are considering whether to use a high Panhard bar on an oval, or stiff rear springs. The high Panhard bar with softer springs will ride two-wheel bumps better. It may or may not ride one-wheel bumps better. That depends on how flexible the tire sidewalls are. The higher Panhard bar creates more lateral tire scrub on one-wheel bumps. If the tires have a lot of lateral flexibility, (e.g. dirt tires with tall sidewalls, running low pressures) that may not matter a great deal. If the tires are stiff laterally (e.g. low-profile road racing slicks, as on a Trans Am car), it matters more.

On the other hand, if we are seeking rear downforce, we may opt for stiffer springs, and accordingly a lower Panhard bar. The stiffer springs will keep the rear end of the car higher, resisting the influences of banked turns and aerodynamic downforce. The higher dynamic rear ride height will let air out from under the car better, and get the spoiler and rear deck up into the airstream more.

The higher rear ride height will also add some drag. If we're running Daytona or Talladega, that may matter more than downforce, and we may want the rear to squash down as much as possible. In that case, we want softer springs, and a concomitantly higher Panhard bar.

Another difference relates to torque wedge: the tendency of driveshaft torque to load the right front and left rear tires under power. A higher Panhard bar, with softer rear springs, gives us more torque wedge. A lower Panhard bar, with stiffer springs and/or anti-roll bar, reduces the effect.

In general, torque wedge is our friend in a left turn, and hurts us in a right turn. Therefore, for road course work, we want to minimize the effect. Especially with low-profile tires that have stiff sidewalls, we will want a really low rear roll center, and plenty of rear anti-roll bar.

As for diagonal percentage (right or outside front plus left or inside rear tire load, as a percentage of the four-wheel total), for oval track applications this plays off against both geometric roll resistance (Panhard bar height) and elastic roll resistance (springs and anti-roll bars). We can run more diagonal percentage, and balance this with increased rear roll resistance (or decreased front roll resistance) from geometric or elastic sources, or we may run less diagonal percentage, and correspondingly less rear roll resistance relative to front roll resistance.

The key to finding a good combination here is getting the car consistent as track conditions vary. A setup with relatively little diagonal percentage tends to go loose on slick surfaces and get tighter as grip improves. This is the most common pattern. A car with high diagonal percentage may go tight on slick instead. In my experience, it takes considerably more than 50% diagonal to get a stock car to do this. Somewhere between these extremes is a combination that changes its handling balance relatively little as grip levels vary.

Additionally, a car with ample static diagonal percentage and relatively great rear roll resistance tends to be tighter on both entry and exit, and freer in the middle of the turn, than one with a combination using less static diagonal.

The reason for both of these effects is that static diagonal is a starting point for load transfer and does not change with lateral acceleration, while roll resistance balance affects the way dynamic diagonal percentage varies with lateral acceleration. A setup with ample static diagonal and ample rear roll resistance has relatively great dynamic diagonal when cornering force is moderate (low grip; entry and exit) and relatively little dynamic diagonal when cornering force is great (high grip;

mid-turn), compared to a setup with less static diagonal and a less rear-stiff roll resistance distribution.

It will be apparent that this is a complex game, and there is no one-size-fits-all solution, especially with the wide variations in driver preference. But when we understand how the variables interact, we improve our ability to tailor the setup to the driver and the track, and to keep the car more consistent as conditions change.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

IS IT BAD TO LIFT A WHEEL?

After watching the new NASCAR Car Of Tomorrow (COT) run on some short tracks and then at the Sonoma road race it certainly appears that the NASCAR teams have some real development to do. The Sonoma road race certainly illustrated what I would consider a real problem with the COT and that is the very radical lifting of the inside front wheel during heavy cornering.

The three wheeling was not incurred by going over curbs, it was the classic three wheeling very similar to a dirt car with lots of cross weight and plenty of traction and torque. The lifting of the inside wheel was especially evident in the 8-9 turn complex where the car was carrying good speed and especially coming out of turn nine, which is a flat sweeping left hander, many cars would carry the inside front, left wheel, to the point that they had to brake for turn 10. Some cars would carry the wheel 4-6 inches off the ground. Of course once they braked for 10 the left wheel, which is now going to be the outside wheel in 10, would regain contact with the track and of course it is not turning so there would be a pretty good "unsettling" of the car before it was set into turn 10. I think that the drivers are using third gear for this series of turns so there is pretty good torque available as they exit turn 9. Having run sports cars at Sonoma back in the '80s, I remember that we did not do anything related to side to side weight distribution and we ran our cars with pretty much 50/50 side to side weight. Although Sonoma is a clockwise track and is predominantly right hand turns we found that trying to bias weight to the right side of the car would adversely affect the car in the left corners and we ran the "carousel" at that time which was a very long, high speed, down hill sweeping left turn onto a good straight so the right side weight was not the thing to have in this corner. NASCAR does not run the carousel so it is highly possible that right hand weight would be a advantage in the track configuration that NASCAR runs. I am not sure if it is legal to run lots of inside weight on the COT.

At the short tracks that they have been running these cars on, you can also see the cars lift the inside front wheel as the power is rolled on exiting a turn, which I am sure is very much assisted by cross weight jacking trying to keep the inside rear planted for a good drive off the corner, but at Sonoma the cars could and would carry the inside wheel in both right and left hand corners.

I have always believed that poor chassis stiffness in torsion is a major contributor to "three wheeling" and looking at the COT chassis it certainly doesn't appear to be well braced for torsional loads. What are your thoughts on the COT as to its attribute of going onto three wheels while corning. Causes, and possible fixes???

It is not necessarily bad for a car to lift the inside undriven wheel.

Ideally, we would like to use all four tires equally at all times. This would involve not only having all four wheels driven, but also having the c.g. at ground level. That's impossible, of course.

With only two wheels driven, and the inevitability of lateral load transfer, we have to make some compromises. The nature of these compromises varies depending on the rules, the track, and our overall design and setup strategy.

If the rules impose the same tire size limit on all four wheels, cornering speed is moderate, and aerodynamic lift or downforce is negligible, we get best steady-state cornering with around 50% static rear weight, and similar overall roll resistance at both ends of the car. In this situation, we should not lift a wheel. To win a moderate-speed skidpad competition, we might want such a setup.

However, a road race is not a skidpad test. On most road courses, most of the turns are of modest duration, and are separated by significant straightaways. In this situation, when cars are nearly equal in power, the race is won on the straights, but the straights are won in the turns. That is, it becomes very important to have good turn entry and exit speed, and late braking points. The brakes also have to last through the race, which is a big factor in Cup cars on a road course. To make the car brake as well as possible, and put power down as well as possible, we need as much static rear percentage as possible. This costs us some steady-state lateral acceleration, but it gains us longitudinal acceleration, both forward and rearward. It also saves the front brakes, which are normally the ones that give out first, because optimum rearward acceleration will be achieved with less front brake bias than we would otherwise use.

To make a tail-heavy car corner neutrally with equal-size tires at both ends, the front end has to absorb the greater part of the lateral load transfer. When such a car is accelerating hard both laterally and forward, it my very well lift the inside front wheel. This does not necessarily mean the setup is bad. It means that some lateral acceleration has been sacrificed to increase forward acceleration.

The COT has front and rear clips that are designed to be more deformable in a crash than was the case in the old cars. That probably does cost some torsional rigidity as well. However, if anything, low torsional rigidity makes a car less prone to wheel lifting, at least for a given suspension setup. This is largely academic, because with a less rigid frame we will normally run more front roll resistance; we will have to, to get the same wheel loads. If we are comparing flexible versus stiff frames, with equal dynamic wheel loads, there will be little or no difference in tendency to lift a wheel.

One thing that does affect the tendency to lift a wheel is the c.g. height, and here there is a substantial difference between the old car and the COT. I hear that the COT has a center of gravity fully two inches higher than the old car. I have a hard time understanding where that much difference could come from, but certainly the roof is higher by that much, and I understand the frame rails are at least somewhat higher.

Once the front end reaches 100% lateral load transfer, any further increase in roll moment can only be reacted at the rear. Consequently, the car has much less angular roll resistance beyond the point of wheel lift, and any further increase in lateral acceleration produces a relatively large increase in roll angle, with a correspondingly greater amount of daylight visible under the inside front tire. It often becomes quite difficult to carry the wheel just a little.

I understand that at Sonoma, a car spends about three times as many seconds accelerating rightward as it does accelerating leftward. That would mean a right-heavy weight bias would be advantageous. However, I doubt that it's legal.

It's not easy for people not on a team to get NASCAR rule books. They are normally only sent out to people getting a NASCAR license, and the rules are subject to revision and interpretation in mid-season. I do know a person who works on a Busch team that started out as a Cup team, and he has the Cup rules as they existed at the beginning of the season. At that time, there was a minimum left side weight for the old cars on clockwise road courses, but none for the COT's. As of the start of the season, the old cars had the same minimum weight for the left side on road courses as for the right side on ovals: 1625 lb., without driver, out of a minimum total of 3400, without driver. COT's have to have 1650 right side out of 3400, and there is no rule specified for road courses.

Evidently, the plan at that point was to have the old cars run the road races, and that got changed. With 200 lb. of driver weight, distributed 50 lb. right/150 left, an old car would have 49.3% left for a road course. If the COT is required to have 1650 left for a road course, it would be exactly 50% left with driver, based on the same assumption for driver weight. Actually, the driver sits slightly closer to center in the COT, but there wouldn't be any possibility to make the car markedly right-heavy.

One thing that would make the car carry the inside front wheel more readily, and higher, in left turns is torque wedge: the effect of driveshaft torque on the chassis. This tries to roll the car to the right, unload the left front and right rear, and load the right front and left rear.

HOLLOW VS. SOLID ANTI-ROLL BARS

I have a question about sway bars. I'm looking to upgrade the OEM bar set on my 2001 Pontiac Trans Am, and with all the different manufacturers out there that produce many different diameters/grades, there also is the issue of a solid vs. hollow sway bar. Do you have any recommendations on this? I'm sure weight is definitely a reason for hollow but my question is, is the performance of a hollow bar close to that of a solid one?

For the same outside diameter, a hollow bar is softer than a solid one. A hollow bar can provide the same stiffness as a solid one, with less weight, but the outside diameter has to be bigger.

Other things being equal, a hollow bar has higher stresses than a solid bar. If the bar is short, or the arms are short, in some cases the bar needs to be solid to avoid stress levels that would cause the bar to take a permanent set or fatigue prematurely. Short of this point, there is some reduction of longevity with a hollow bar. The bars on a Trans Am go all the way across the car, and have long arms, so hollow bars should work fine.

To give you some idea of what diameter you'd need with a hollow bar to equal a solid bar, if you had a factory rear bar 5/8" in diameter, a ¾" O.D., .060" wall hollow bar would be about the same stiffness. A ¾" O.D., .090" wall bar would be about 30% stiffer.

If you have a 1 1/8" solid front bar, then a 1 ¼" O.D., .156" wall hollow bar, or a 1 5/16" O.D., .120" wall hollow bar, would be about the same stiffness. A 1 3/8" O.D., .120" wall bar would be about 20% stiffer.

If you buy bars by an advertised rate in lb/in at the arm end, be aware that there are two ways of expressing this rate, and not all manufacturers use the same convention. The more common method is to rate the bar like a ride torsion bar. That is, one end is moved a known linear amount and the force per inch is computed. This gives you rate in pounds per inch per end pair: the force when each end moves half an inch, meaning there's an inch difference between the two ends. Some manufacturers prefer to double this figure to get the cataloged value. This method gives you the rate in pounds per inch per end: the force when each end moves an inch, meaning there is 2" difference. This method has the advantage of being easier to equate to a change in ride spring rate. Neither method is more correct than the other, but you do need to know which method a manufacturer uses, if you want to make comparisons based on rate.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DROOP-LIMITED REAR SUSPENSION

I have a question for you, and it's related to something I recently discovered on my race car, and something that you've been talking about a lot lately - droop limiting.

My car is a Mustang with a strut front, 3-link in the rear, 3100#, 400+rwhp, 275 DOT race radials all around, etc.

In a search for lots of bump travel I made some changes to the upper (chassis side) shock mounts in the rear. When my shocks arrived, I never noticed that they were shorter than I thought, and I ended up with a car that has less than 1" of droop travel in the rear. For reference, the car is relatively softly sprung (400# wheel rate in front, 300# wheel rate in the rear) and uses lots of travel with a relatively high CG (about 17" or so).

I was wondering how you thought this lack of droop travel might affect handling? Under braking? Into and through a corner?

I can say this with certainty: when the inside rear suspension tops out, the rear roll resistance increases dramatically, and that makes the car looser (produces oversteer). I cannot say with any certainty when this is happening with a particular car. Data acquisition is very useful for determining this. (The questioner here is an engineer with a NASCAR team, but the question concerns his own personal race car. Presumably he will not be allowed to borrow team data acquisition gear for his own races. He didn't mention whether he has his own.)

If the 3-link rear is set up for moderate anti-squat and anti-lift, and the springing is fairly soft, it is highly probable that the car tops out the inside rear shock when trail-braking. I would also expect that the shock would top out in steady-state cornering.

It is possible to have so much anti-squat in a beam axle rear suspension that the suspension tops out under power, even when the car is running straight. It is not uncommon to see this in dirt cars, even with as much as four inches of droop travel. I do not recommend this, but I see it when I go to races.

It is also possible to have so much anti-lift that the rear suspension compresses rather than extending in braking.

With a simple three-link rear, ample anti-squat implies fairly ample anti-lift. This does not mean that the car necessarily compresses the rear suspension in braking if it extends it under power. The percent anti-lift in braking will generally be considerably less than the percent anti-squat under power, because the rear tires have to provide all of the propulsion force, but only about a quarter of the retarding force. If the questioner's Mustang were set up to actually lift the rear under power, the rear suspension would very probably still extend under braking. It would be possible to have a situation where the only time the rear suspension wasn't topped out would be when there was little forward, rearward, or lateral acceleration.

This is probably academic, because in general road racing cars are not set up with extremely severe anti-squat.

With the more complex suspensions sometimes used in dirt cars, it is possible to have a lot of anti-squat and also have severe pro-lift in braking. I really do see dirt modifieds and Late Models that appear to lift the left rear to the droop limit most of the way around the track. The left rear only comes down momentarily when the driver is transitioning from power to braking at the end of the straights.

SATCHELL LINK REAR SUSPENSION

How do you calculate the roll center for a Satchell link rear end?

For readers unfamiliar with the Satchell link suspension, it is a form of triangulated four-bar linkage for a beam axle. This family of systems use diagonal semi-trailing links to provide lateral location of the axle, rather than purely lateral links such as a Panhard bar.

The most common version of triangulated four-bar, sometimes called the Chevelle style, has two upper links that are angled around 45 degrees in plan (top) view, and two lower links that are closer to parallel in plan view. The upper links are further apart at the front than at the back. Their centerlines converge at a point above and behind the axle center section. The lower links converge to a point well ahead of the axle. The two convergence points, or instant centers, define an instantaneous axis of rotation, about which the axle moves in roll. The roll center height is the height at which this axle roll axis intercepts the axle plane: the vertical, transverse plane containing the axle centerline.

Note that this is not the midpoint height of the axle roll axis, or the point midway between the instant centers, which at least one author has suggested should be taken as the roll center. Indeed, as we will discuss, there are cases where no such midpoint can be defined.

It is also fairly common to see this system used upside down. When a low roll center is desired, it is possible to have the lower links converge toward the axle center section rather than the upper links.

The Satchell link system is basically a Chevelle system, turned upside down and back-front, or turned upside down by rotating it about a transverse axis rather than a longitudinal one. That is, it has lower links that attach to the axle near its outboard ends and come close together near their front ends. The upper links also attach near the outboard ends of the axle, and are either parallel or close to parallel in plan view.

Finding the roll center is the same as with other triangulated four-bar systems. Find the instant centers for the lower and upper link pairs, define a line connecting these, and find the height at which this line intercepts the axle plane.

There are cases where the link centerlines will have no intersection. Even if they do at static position, they won't in a rolled condition. Generally, the link centerlines will pass over and under each other when the suspension has some roll displacement. In such cases, the best approximation is to find where the lines cross in plan view, and then find a point midway between them in height. In other words, find where the lines have the same x (longitudinal) and y (transverse) coordinates, and average their z (vertical) coordinates for those points.

It is also possible for one pair of the links to be parallel in plan view. They then will not cross at all in plan view. The axle roll axis then passes through the instant center of the non-parallel links (lowers in a Satchell link), and has a side view inclination parallel to the side view inclination of the links that are parallel in plan view (uppers in a Satchell link). If those links do not have the same inclination, we take an average of the two. Note that in this case we have only one link instant center to work with, and it becomes impossible to define a midpoint between two instant centers.

If both upper links in a Satchell system are horizontal and parallel, the roll center is at the same height as the lower link instant center. However, if the uppers are not horizontal and parallel, the roll center height is different from the lower link instant center. As the car moves in ride, the roll center stays about the same height from the ground if the upper and lower links have similar side view projected length. This is important, because the front roll center in most independent front ends moves up and down with the sprung mass, though not necessarily the same amount as the sprung mass. If we supposed that the roll center is at the lower link instant center in the Satchell system, that might lead us to suppose that we'd have better compatibility with an independent front suspension than we have with more conventional triangulated four-bar systems. Unfortunately, that turns out not to be so. However, the system is no worse than conventional triangulated four-bar designs in this regard.

We also do not escape the interrelatedness of roll steer, anti-squat, and geometric anti-roll, as some have suggested. But again, this aspect is no better or worse than in other triangulated four-bar systems.

Advantages and disadvantages of the Satchell system mainly come down to packaging and load paths, both in the axle and in the frame. Whether we gain or lose in those regards will depend on the particular installation.

A WORD FROM THE INVENTOR

After I originally mailed this newsletter, I received a note from Terry Satchell, who originated the suspension bearing his name. He contributes some interesting background on the history of the design and his views regarding its advantages. He writes:

…I had done a Lotus Super 7 type rear for a Trans Am car where the lower A-arm has a pivot under the axle with two upper longitudinal arms. After running it for a year we kind of got the impression that the rear roll center was too low, it being at the pivot of the A-arm to axle joint. Since I had designed four bar link rear axle suspensions for General Motors for several years, I knew how to analyze them and create what I wanted. I wanted a geometry with a roll center basically midway between the bottom of the axle housing and drive axle centerline. By reversing the lower arms to converge to the center ahead of the axle I was able to achieve a good geometry. It worked well on the track, and in fact one of the neat parameters was that the anti-squat increased on the inboard wheel with roll giving a tightening effect on power that helps corner exit.

I did a version of this geometry for Herb Adams who did all of Walker Evans truck suspensions and they were very happy with it. They in particular liked the anti-squat difference across the back with roll on throttle for their Stadium racing trucks.

Since then I have done several versions for various disciplines and now there is a company in Pennsylvania that provides aftermarket conversions for Mustangs, Camaros and Firebirds to get rid of the leaf spring suspension and use this four bar link with coilovers.

I have also had success with several of the "Locost" builders using it. [The Locost is a kit car similar to a Lotus 7.] *One in particular used it to control his deDion beam.*

Just thought you might like to know a little more of the story.

By the way, I never intended to put my name to this version of a four link. Herb Adams named it in one of his books and that is how it got started being call the Satchell Link.

I'd like to thank Terry for his contribution, and also mention that he was the person who taught me the right way to determine the front-view and side-view projected control arm geometry in independent suspensions.

I concur about the roll center height, and the desirability of having more anti-squat on the inside rear wheel than on the outside rear.

In fact, this latter property is a characteristic of most trailing arm rear suspensions, including independent ones that use trailing arms. Strictly speaking, in the case of independent systems, we usually have less pro-squat on the inside wheel with roll, and more pro-squat on the outside wheel, rather than anti-squat as such. However, the effect on roll and on wheel loads is similar: the difference in longitudinal "anti" produces an outward roll moment, or pro-roll moment, in the rear suspension that adds wedge to the car – adds load to the inside rear and outside front wheels.

It is a bit risky to generalize about this subject purely on the basis of general suspension type. Exceptions can be found to most such generalizations, and the dynamics of any actual car will depend on the specific design of that car.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
October 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHAT EXACTLY ARE SHAKER RIGS FOR?

Over the past few years I have seen and heard about all of the developments with 4-post, 5-post, and 7-post shaker rigs. However, it seems that most articles focus on the technology employed by the latest and greatest rig rather than what type of results the rig is going to yield. My question is this, when teams test on a shaker rig, what exactly are they going to get in return in terms of results (i.e. lateral and longitudinal G forces or something else)? Also, maybe even a better question, what is it they are testing for (i.e. is there a way these rigs measure "grip" in particular) and what type of testing regiment is going to be employed? For instance if you are a team that is on a limited budget and can only afford to put your car on the rig once per year, do you put the car through a generic set of accelerations and simulated corners or do you try and tune for something more specific or maybe even a particular track? And lastly, at the end of the day how do you know what was the better setup (if there isn't a way to measure "grip" specifically)?

Shaker rigs are an example of street driving improving the breed for racing. They originated in passenger car development, and were later applied to race cars. The original objective was to explore the car's response to excitation at the wheels, in a more controlled and observable situation than could be achieved on a test track. Improving handling wasn't the main goal. Engineers were more concerned with isolating frequency-sensitive effects that would impact durability and noise, vibration, and harshness (NVH). Does a portion of the roof or the floor pan resonate at a certain frequency? Are there brackets or ducts that buzz or might break? Does anything rattle? At what speeds and frequencies do the wheels "dance", and does this agree with calculated natural frequencies?

Early shaker rigs had posts only under the wheels, and the sprung mass was allowed to do whatever it wanted to do in response to excitation at the contact patch. It soon became apparent that it would be more desirable to allow the contact patches to float both laterally and longitudinally, so that the track and wheelbase could vary with suspension movement, without having tire scrub fight the movements. A fifth post attached to the sprung mass allowed the car to be anchored horizontally while all four contact patches could float. The car still needed either an anti-rotation feature in the fifth post, or some other mechanism, to constrain the car in yaw.

This was good for replicating highway driving, but it was impossible to explore suspension behavior in conditions of sprung mass displacement caused by aerodynamic loads, banked turns, or pitch and roll due to longitudinal and lateral acceleration. To reproduce these displacements, three posts were attached to the sprung mass. Conveniently, three posts also constrain the car in yaw, without any additional devices. With this addition, it is possible to reproduce any combination of suspension displacements that the data acquisition system may have recorded on-track, or that an engineer might imagine.

Importantly, this does not mean that the rig reproduces all the forces acting on the suspension. The posts only move vertically. The rig cannot apply or reproduce any horizontal forces. It also cannot measure any horizontal forces. Therefore, we cannot measure grip at all. We also cannot reproduce loads and frictional influences in the suspension that result from horizontal forces at the contact patches. We do not even know vertical wheel loads with any accuracy, because horizontal forces at the contact patches affect these.

This last phenomenon is particularly significant in stock car rear suspensions, and in any suspension with large geometric "anti" effects. (Stock car rear suspensions, and most beam axle suspensions, have ample geometric anti-roll.) With formula cars, where there is modest anti-roll and anti-pitch, the wheel load values on the 7-post are closer to reality, but still not highly accurate.

I am referring here to individual wheel loads. On the rig, we do get reasonably accurate total loads for all four wheels, and for the left, right, front, and rear wheel pairs. It's the loads for individual wheels and for diagonally opposite pairs that are inaccurate.

Even though we do not read accurate individual wheel loads on the shaker rig, we can get a reasonable comparative evaluation of how much these loads vary dynamically in the excitation conditions that we test. What we are after is minimum load variation. We cannot measure grip itself, but we do know that variation of normal load is bad for grip, and minimizing such variation will improve grip.

This is true for two reasons. First, if the tires unload for any significant length of time, the car is limited by whatever grip the tires have at that load. If the grip limit at that load is exceeded, the car will break traction, and once traction is broken it is hard to regain. A slide, once initiated, tends to persist.

Second, even if we are dealing with load variations over small time spans, so that vehicle inertia masks the intervals of low load and grip, we do not fully recover during the highly loaded periods what we lose in the minimally loaded periods. Why? It's our old friend, load sensitivity of the coefficient of friction. Adding 100 pounds of load doesn't help as much as losing 100 pounds of load hurts.

So we try different combinations of shock valving and springing, and we try to minimize load variation at the wheels.

I was asked recently how you valve shocks based on suspension displacement traces provided by data acquisition. I had to reply that I don't know of any way to do that, and I don't know anybody who claims to be able to do that. However, with the use of a 7-post rig, we do at least have a means to make gains by trial and error.

There is another kind of test rig that does let us look at the effects of horizontal forces at the contact patches. It's called a kinematics and compliance tester, or K&C rig. Until recently, only passenger car manufacturers had these, and nobody was offering testing by the day or hour to racers. There is now a K&C facility available to the motorsports community, just up the road from me in Salisbury, NC. It's called Morse Measurements. The people who run this facility are friends of mine, so I'll give them a plug. They are at www.morsemeasurements.com or 704-638-6515.

What the K&C rig does is grab hold of the frame of the car, and then apply horizontal loads at the tire contact patches. The sprung mass can be kept from moving, and the changes in wheel loads measured that way. Alternatively, the sprung mass motion can be controlled to keep the total normal force at the contact patches constant and we can measure the displacements and load changes that result. Roll and pitch moments are applied to the frame, based on an input sprung mass c.g. height. Individual wheel horizontal forces are applied based on estimates, which in turn are based on tested or estimated tire properties and calculated wheel loads.

We then see from the test results whether our estimated wheel loads approximate those produced in the test, and we may want to adjust our contact patch forces and do another iteration on the rig.

There are other tests on the K&C rig as well, including simply cycling the car slowly in heave through its suspension travel and measuring the camber changes, contact patch scrubs, and wheel loads, with the contact patches allowed to float horizontally so that all horizontal forces at the contact patches are eliminated. This test can be done on either a K&C rig or a 7-post, so it can be used to compare the agreement of the two types of rig.

What the K&C rig does not do is move the chassis fast. The 7-post can load and move the chassis at racing speed, but only up and down. The K&C can load and move the chassis in all directions, but only slowly. At this writing, there is no test rig that moves and loads and measures a car in all directions at once, at realistic on-track speeds. And even the low-speed testing done on the K&C rig uses estimated or assumed grip levels at the contact patches.

We do have a test that does measure grip: the skid pad. It measures average grip over a lap, by stop watch, or over smaller intervals, using accelerometers or other sensors. Of course, we have difficulty observing the car during testing, because the car is in motion rather than on a test rig. And the grip measured is only on a particular surface, on a particular radius.

We thus have three useful methods of testing, all of which provide useful comparative measurements, but none of which allows us to fully replicate track events in a controlled environment. I could even say we have four methods, if we include wind tunnels, or a much larger number if we include component testing devices such as tire testing machines, shock dynos, and engine dynos. But so far we have no device that really puts the car through its paces in a manner that replicates in full what happens on the track, on a stationary platform where we can safely watch and measure it. The best we can do is look through a number of windows, each of which affords a view from a different angle, and use our mind to try to put these glimpses of reality together.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2007

Reproduction for free use permitted and encouraged.
Reproduction for sale subject to restrictions. Please inquire for details.

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ON THE MARCH AGAIN

A few months ago I was asking you two questions, and you have been very helpful! New question: with my Formula 2 March 712 at Dijon (France) I have on 2 corners at the beginning and end of main straight in 3rd gear. In these turns I have UNDERSTEER and I feel I have to try first with stronger aft spring or with stronger aft ARB, not only to worsen the grip in the rear but also to improve the grip in the front.

I am also considering wedge adjustment. Is it also useable on road circuit like we have in Europe or just for ovals? How about if I put more corner weight on the LH aft and or RH fwd wheel – or is this a bad compromise?

Before I get into answering the main question here, I thought I'd pass along some information submitted by reader Robert Koch.

You were helping a person who drives a March 712. In 1972 und 1973 I worked on such a car in Switzerland. This car was very softly sprung. It had a very soft chassis which meant that increasing the anti-roll bar or the springs the chassis just flexed. Now the reason why I am writing you: the March 712 racecar had two weak points. One of them was the inserts in the back wall of the monocoque where you fasten the tubular frame with the engine and the transmission. They are not such that they live a long time. One has to replace them quite often. The problem is, you do not see that they are loose! Point 2 is: the upper squared-tube where the differential and/or the gearbox is mounted cracks on the underside, close to the two points where the springs are mounted. This is also something you do not see just like that.

It is important to make sure you eliminate any problems like these, rather than tune around them. I would also make sure the rear wheel pair is not aimed to the left or the right. Both rear wheels should generally have the same toe-in relative to the vehicle centerline. Also, make sure that both rear tires are the same size. Measure their circumference with a tape measure.

Looking at a track map of Dijon (see http://en.wikipedia.org/wiki/Image:Circuit_Dijon-Prenois.png), I see that the first of these turns, called Courbe de Pouas, or Turn 8, is the largest-radius one on the track, and therefore probably the fastest. The one at the end of the straight, Double-droite de Villerois, or Turn 1, is a slower one, but not really slow. Both appear to be double-apex turns. Turn 1, as its name suggests, is really a double-right, with two fairly small-radius segments joined by a markedly larger-radius segment. Turn 8 has a short, somewhat tight-radius segment at its entry, followed by a large constant radius over most of its length. The circuit is run clockwise, so most of the turns are right-handers. There are only two left-handers of any significant duration.

The main straight, called Ligne Droite de la Fouine, connecting turns 8 and 1, is long. From the map, its length must be well over half a mile – perhaps a kilometer, perhaps three quarters of a mile. The back section of the circuit is almost all turns. There are four straightaways, but none of them is anywhere near the length of the front straight. So the questioner is correct that speed in these two turns, especially exit speed from turn 8, is important to lap time.

For a course like this, I don't generally advise a wedged setup, at least not as a starting point. I do advise a right-heavy setup, for cars that have ballast, which the questioner's car probably does not.

Cars with flexible frames or sprung structures do respond to suspension adjustment. They just respond less to a given adjustment than cars with good torsional stiffness. You need more change in the springs and bars to get a given change in dynamic diagonal percentage or in cornering balance. You need to put more turns into the spring seats to get a given change in static diagonal percentage.

When attacking a handling problem that occurs only in certain turns, I always advise looking for characteristics in those turns that are not shared with the turns where the problem doesn't show up.

If the problem shows up in fast turns but not slow ones, that points to a problem with aerodynamic balance. If the car has understeer only at high speeds, that is best cured by adding front downforce, or reducing rear downforce. In some cases the rules will prevent us from making the changes required to achieve this, but it is the best approach when we can do it.

If the car has understeer only in right turns, and the car is to all appearances symmetrical, that points to an alignment problem, or something loose or flexing that only affects turns in one direction. A car designed for road racing shouldn't act different in right and left turns, if the wheel alignment is symmetrical, the right and left springs are identical, the right and left tire sizes are the same, and there are no other asymmetries. If it does act different in right and left turns, we can crutch that by running more or less than 50% static diagonal, but chances are that we are tuning around some other problem that we have yet to discover.

I do recommend running some negative wedge (usually meaning less than 50% on the right front and left rear) on live-axle road racing cars, especially for courses with lots of right turns. This compensates for the effect of driveshaft torque on wheel loads. However, this would not apply to a rear-engined formula car with independent rear suspension.

RADIAL TIRE CAMBER

I have just started racing an MG Midget. The weight is about 50/50 front/rear and at present it runs 8in crossply slicks all round. Front is basically double wishbone and the rear is a live axle with Panhard rod.

I have considered running radial tyres but have been told that I would need a lot more negative camber (at present about -1 deg).

As far as I can tell most modern (up to F1) cars seem to run quite a lot of negative on the front but little if any on the rear.

So, do radial slicks need more negative camber than crossplies, and if so, why? If it is needed on the front why not the rear (though difficult with a live axle!)? With a live rear axle and a lot of negative on the front, wouldn't the rear be relatively lacking in grip?

It is not necessarily true that radials always want more negative camber than bias-plies, but it is valid as a generalization. Many racing radials are not truly radials, because the sidewall plies do cross each other, although not at as great an angle as in true bias-plies.

It is very common to see more than a degree of static negative camber on independent systems of all kinds, on bias-plies or radials.

Why do radials tend to want more camber? Mainly, I think because the sidewalls are more flexible. This makes them tolerate more camber, and want more.

If the tires want more negative camber, and you can only get this on the front, will the car tend more toward oversteer? Maybe. Remember, your independent front suspension allows the front camber to change with roll, while the rear beam axle produces no camber change with roll (except a bit due to tire deflection). Most independent suspensions provide some camber recovery in roll, but generally not 100%, because this would create excessive camber change in ride. You may currently have a degree of negative at the front statically, but that would mean you barely are keeping the outside tire upright, if that, once the car has reached full roll in a turn.

In any case, if you find you have more front grip compared to rear grip with any new tire or front camber setting, but the new tire or setting gives better overall grip, stay with the new tire or setting and work with the roll resistance distribution to get the balance back. If the car oversteers, the proper fix is to add more roll resistance at the front, or take some away at the rear, using springs and anti-roll bars.

Why do you see so much negative camber at the front on F1 cars, and so little on the rear, even though the car has independent suspension at both ends? I don't know. I can say with certainty that this helps propulsion and hurts braking. It may be that the engineers have decided that forward acceleration out of the turns is worth more than rearward acceleration at the ends of the straights. Rules permitting, I would be tempted to ask the tire supplier for bigger front tires, and run them at less camber, or add more camber at the rear.

A RADICAL RADIAL NOTION

Instead of relying on low profile tyres with short sidewalls for handling response, why not utilize high profile types (for good surface following capability) and employ internal bracing to resist lateral loads introduced by cornering? One scheme would be to have tension cables or "strings" running from the bead to the intersection of the tread and sidewall on the opposite side of the tyre. That is, tension members would be diagonally disposed relative to tread plane. The system would look analogous to a wire wheel but be inside the tyre itself. Such a tyre would have a high sidewall for good ride and terrain following capability but would also be stiff laterally. Would this work?

I do think this would increase lateral stiffness, and maybe increase response, while maintaining vertical softness. My worry would be that the tire would act more like a bias-ply with respect to the tendency to lift or unload the inside portion of the contact patch. In a conventional radial, the tread can move laterally relatively freely, yet stay planted relatively flat while doing so. I would think the cables or strings would tend to lift the inside shoulder instead.

Another concern might be damage to the cables when mounting the tire. Two-piece rims might be necessary to overcome that.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2007

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

INS AND OUTS OF TOE AND ACKERMANN

Your comments regarding Ackermann, or anti-Ackermann, or perhaps why bother with Ackermann, would be appreciated. It occurs to me that the more heavily loaded outside front tire must operate at a larger slip angle than is possible for the opposite unloaded front while cornering. The result is that the lightly loaded tire is pulled sliding across the pavement not doing much work. I assume it is reasonable to distinguish between the slip angle of a tire which is being deflected by cornering forces and one which is sliding. Would it be more efficient to use front roll stiffness to completely unload the inside front during hard cornering? Does Ackermann aid turn in? Would anti-Ackermann provide an advantage? Is it possible that toe out in combination with anti-Ackermann might be effective? Is the turning angle of the front wheels in most hard cornering racing situations so small that Ackermann is not a factor? What are your thoughts?

For the benefit of newbies, Ackermann effect is a property of steering geometry that causes the front wheels to toe out as steering angle increases. If the front wheels toe in as steering angle increases, that is called negative Ackermann or anti-Ackermann. If the toe angle does not vary with steering input, that is zero Ackermann, or parallel steer.

Ackermann effect must not be considered in isolation. The tires do not know what kind of Ackermann properties the steering system has. They only know how much they are toed in or out, at a particular instant. How much the wheels are toed in or out at a particular instant depends on a combination of Ackermann effect and static toe setting.

Additionally, there can be toe changes due to bump steer or compliance steer. For simplicity, we will disregard these effects here.

The static setting provides a starting point when the steering is centered, and Ackermann effect adds toe-out from there, in a fixed relationship to handwheel (steering wheel) angle.

Trouble is, the optimal toe angle in terms of tire performance is not a constant, nor does it have a fixed relationship to handwheel input. Unless we are prepared to engineer some sort of elaborately programmed steer-by-wire system that controls the right and left front wheels independently, we cannot obtain optimal geometry for all possible situations. We are stuck with striking a compromise for a particular set of conditions.

The nature of that compromise depends in part on how much extra slip angle we wish to give the outside front wheel. One can reasonably argue that the more heavily loaded wheel reaches peak cornering force at a greater slip angle than a more lightly loaded one, so the front tires achieve the greatest peak cornering force when the outside tire has a greater slip angle than the inner one.

The questioner asks whether there is a difference between slip angle of a tire that is sliding and slip angle of a tire that is not sliding. More precisely, we might talk about a tire that is only partially sliding, in the rear portion of the contact patch, and one where sliding is occurring in the entire contact patch. There is no difference in the definition of slip angle; it is simply the angular difference between the tire's instantaneous direction of travel and its instantaneous direction of aim: the difference between its bearing and its heading. However, there is a difference in the effect of adding slip angle in the two cases. If the tire is below the slip angle where its lateral force peaks, adding slip angle adds lateral force and also adds drag. If the tire is above the slip angle where its lateral force peaks, adding slip angle does not add cornering force and indeed probably reduces it. However, up to slip angles associated with total loss of control, drag continues to increase as we add slip angle.

One thing that makes all this a bit complex is that when a tire is near its peak cornering force, lateral force greatly exceeds drag force, yet moderate slip angle changes have a fairly small effect on lateral force, but a relatively large effect on drag force. This makes it difficult to evaluate the effects of toe changes on cornering capability, purely by observing changes in car balance or amount of understeer.

Drag forces at the tires do not turn the car in the sense of accelerating it laterally, or centripetally (toward the center of the turn), but they do tend to steer the car: accelerate it in yaw, or rotate it.

This means that it is possible to have a case where we are adding moderate amounts of cornering force at the front by increasing outside tire slip angle, yet the steering trace and the driver feedback may show increased understeer, and the car may be slower! In such a case, we can, at least in theory, dial a bit of oversteer in by juggling tire load distribution, and then we may have a slightly faster car than we started with. The only way to know is to try this rather than immediately backing up on the toe or Ackermann change.

We have so far been assuming that what we're after is the highest peak cornering force. We get that if both tires are at their optimum slip angle for lateral force at the same time. However, one could also make a case for having the tires not peak together, to make breakaway gentler and make the car more driver-friendly. This is somewhat analogous to the question of whether to tune the engine's exhaust and induction systems for the same rpm, to get the highest peak power, or tune them for different speeds, to spread the power band.

Is this complicated enough yet? We're just getting started.

Suppose we have sufficient information to decide what slip angles we want on the two front tires, or what difference we want in their slip angles. Does that allow us to say what our toe-out or toe-in should be at a particular instant? Nope. Without knowing the geometry of the car and the turn, and without knowing what the rear wheels are doing, we can't even get close.

For simplicity, let's suppose we don't want any difference in the inside and outside tires' slip angles. Let's take a look at what it would take to get that, in various situations.

Some readers will be familiar with the concept of a turn center. This is the point about which the car's center of mass or c.g. (sometimes approximated as the midpoint of the car's centerline in plan view) is instantaneously revolving, as it negotiates the turn. If the car has a constant attitude angle – that is, if it is drifting or sliding a steady amount – all points on the car are moving about the turn center.

At any given instant, the car has an instantaneous direction of travel, which is always a tangent to its path of motion. If the car is traveling in a curved path, that path has an instantaneous radius r. This radius is equal to square of the car's instantaneous speed along its instantaneous direction of travel, divided by its centripetal acceleration: $r = v^2/a$. In a totally unbanked turn, with the tires sliding only a little, these two quantities are approximately equal to the car's speed as read by a speedometer or wheel speed sensor, and the car's lateral acceleration as measured by an on-board accelerometer. (If the turn is banked, or the car is sliding dramatically, these approximations become much poorer.)

If, in plan or top view, we construct at the c.g. a perpendicular to the car's instantaneous direction of travel, and define a point on that line a distance r from the car's center of mass in the direction of the turn, that point is the turn center.

In the car's frame of reference, the turn center can be anywhere from the rear axle line to well ahead of the front axle line. (It could even be behind the rear axle line, if the rear wheels have a negative slip angle. This could occur when negotiating a banked turn at low speed. Normally we can ignore this case when studying behavior at racing speeds.)

The simplest situation is a small-radius turn, taken so gently that tire slip angles are negligible. In this case, the turn center is on, or very nearly on, the rear axle line in plan view. For zero slip angle at both front wheels, the front wheel axes should ideally meet in plan view at the turn center. That implies that the front wheels will have substantial toe-out.

The larger the turn radius, and the larger the rear wheel slip angle, the further forward the turn center moves. At some point, the turn center will lie on the front axle line in plan view. In this situation, we have equal slip angles on the two front wheels when the toe-out is zero.

In high-speed, large-radius turns, the turn center is usually ahead of the front axle line. We now have a condition where the front wheels need to be toed in for the slip angles to be equal.

As a broad generalization, the front wheels are steered more in small-radius turns than in large ones, and the turn center is further rearward in the car's frame of reference. This argues for having at least some Ackermann in the geometry, but it is harder to come up with a general rule to calculate exactly how much.

In high-speed turns, steering inputs are generally very small, and consequently Ackermann effect has far less influence than static toe setting. Ackermann has greatest influence in autocross and hillclimb cars.

Does Ackermann aid turn-in? Basically, yes, and so does static toe-out, at least up to a point. Really excessive toe-out, whether from static setting or Ackermann, will hurt front grip and the effect will reverse, but within a sane range, turn-in will at least feel quicker with some toe-out. This is partly due to the early yaw moment from inside front tire drag as the handwheel is first turned.

Does using anti-Ackermann along with static toe-out make sense? It's certainly done successfully, on winged single-seaters on high-speed ovals. Logically, however, it makes more sense to use static toe-in with positive Ackermann. This is conventional practice in passenger cars.

What about completely unloading the inside front wheel using front roll resistance? Well, it does make Ackermann academic, at least during the time that the wheel actually is airborne, and it eliminates any concerns about tire drag due to the front tires fighting each other. Unfortunately, in many cases using that much front roll resistance will create excessive understeer.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January/February 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE EXOTIC VEHICLE LAYOUT IDEAS

Thanks for an interesting commentary on the diamond car recently. That certainly has had me thinking. Below are some topics I've been pondering.

Another unusual race car
Following up on the diamond shaped car, I came across another unusual car. This one is a special based on a Mazda Familia hatchback. The original car had four-wheel drive and was powered by a transverse mounted engine and gearbox assembly. The engine was a turbo four (of 1.8 litres I think) with a five speed manual box. The box drove all four wheels via a "central differential". The car was modified by moving the original engine/gearbox and front wheels rearwards. Another engine/gearbox (this time an automatic gearbox) assembly was mounted ahead of it. This powered an extra pair of road wheels up front. The car became a six wheeler with all wheels driven! The front four were steered. The car is reported to need more development work but has been successful on dirt and hill climbs, outperforming V-8 competitors on occasion! It is a good handler and fast with plenty of traction.

I'd be most interested in reading your analysis of this type of vehicle.

A second type of six-wheel car I've thought about is one with only the front four wheels driven. In this case the plan would be to use a pair of front wheel drive engine/gearbox/suspension/road wheel assemblies mounted at the front of the car. The two power trains would operate independently of each other. All four front wheels would be steered. It may be necessary to use a balance beam linkage in the steering to ensure that the front wheels on each side of the car developed similar lateral forces whilst cornering. The final pair of wheels at the rear of the car would be un-powered and non-steering, as is normal for front wheel drive vehicles.

I am aware of various attempts to use independent power-trains in four wheel vehicles. Famous examples include the Twinnie Minnie, which had an engine at each end. VW tried this for the Pikes

Peak hill climb, as did Monster Tajima in his Suzuki at the same venue. There are some interesting handling issues to be tamed with this type of system (where there are no links between the two power-trains). VW used traction control and an early form of throttle by wire. In the end Monster added a shaft linking his two engines via a computer controlled electro-magnetic clutch. The Twinnie and various replicas of it have encountered some handling problems, which appear to be solvable, or at least made controllable by careful engine tuning and modification to engine response and power delivery characteristics. Still there are questions about the ultimate handling behaviour on the limit with this type of set up.

In the case of the six wheeler with only the front four wheels driven it may be that these problems are not present since the driving wheels are close to each other and not at opposite ends of the car.

What is your analysis of these set-ups?

In most classes of racing nowadays, we don't have the option of more than four wheels, or more than one engine. For the most part, the rules prohibit such things. Nonetheless, there are scattered venues where one can at least run such vehicles for fun, and it is fascinating to consider the possibilities.

One fundamental argument against using a larger number of anything is that, as a rule, two little ones weigh more, cost more, and take up more room than one big one. This applies to engines, cylinders in the engine(s), and wheels. That doesn't mean that adding a component is always a bad idea, but there has to be a compelling functional advantage in adding anything to the car, to override this disadvantage.

Multiple engines are usually used for one of two reasons: either the engine in question is small, and using two of them is an economically appealing way to build an all-wheel-drive special, with much-needed additional power, or the vehicle is really large, as with land speed record cars, and high-output engines of sufficient size are either unavailable or unaffordable.

Until recently, controlling two unconnected engines driving the front and rear wheels was a real problem. Perhaps with the advent of by-wire controls, this can be overcome, and potentially even turned into an advantage. The question is, who is going to invest the money to do this, when the car is ineligible for any big-money racing class? Not only do we need to appropriately control the power and rpm of the two engines, we also have to make the two clutches take up in sync with a single pedal or other control, and make both transmissions shift in sync from a single lever or other control.

If we are talking about two identical engines, with unconnected two-wheel-drive transaxles, driving the front and rear wheel pairs, we have a 50/50 torque split front to rear, and no connection between the front and rear wheels. This means that unless the car is severely nose-heavy, the front wheels will tend to spin prematurely in hard forward acceleration. Either the driver will have to lift to control front wheelspin, or, if we have traction control, the power of the front engine will have to be suppressed, meaning we are not making full use of that engine. If we do make the car sufficiently nose-heavy so both engines can deliver full power, that will compromise cornering and braking. Linking the two transaxles with a driveshaft and some form of clutch allows some transfer of power from the front to the rear, but requires the front and rear wheels to be constrained to the same rpm at the traction limit, which compromises handling – unless we use different final drive ratios at the two ends, which would bring a whole set of new problems, chiefly in transmission control and rev limits. Adding such a mechanism to two existing two-wheel-drive transaxles also involves a lot of expense and special parts, which erodes the original economic appeal of using two cheaply available economy car powertrains.

All things considered, if we are designing from a blank sheet of paper (or screen), it is probably more appealing to have one large engine, with one transmission, and, if we want to drive all the wheels, to control torque distribution with the differentials and clutches or other locking devices in the drivetrain. Fundamentally, we have fewer parts, no synchronization issues, less metal, and an easier time getting the torque distribution we want.

It is desirable not only to have the power, or at least a preponderance of it, go to one end of the car, it is desirable that this be the rear end. One advantage of rear wheel drive in terms of car control is that it affords us some measure of independent steering control at both ends of the car. This allows us to create yaw accelerations, and control the car's attitude, with the rear wheels, even when the front wheels are at the limit of adhesion. With front wheel drive, we have two ways of controlling the front end, but nothing except perhaps the brakes to control the rear. If we wish to drive all the wheels, it is desirable, at least in a high-speed road or road racing car, to make the car throttle-steer somewhat like a rear-drive car.

If we are designing a front-engine car with all wheels driven, and we want to have a back seat and a reasonable overall length, we will have somewhere between 50 and 56 percent of the weight on the front wheels. To make the car throttle-steer controllably, we want at least 60% of the torque to go to the rear wheels. If we are designing a pure performance car, it is better to choose a rear-mid-engine layout, which will typically give us around 60% static rear weight. If nothing else, this layout gives us better braking and less yaw inertia than a typical front-engine layout. We then want somewhere in the range of 75 to 90 percent of the propulsion to come from the rear wheels.

There is a definite advantage in driving the front wheels, even if they receive a small percentage of the power. It may not be obvious, but we always have to drive the front wheels, particularly while cornering. When the front tires are running at an appreciable slip angle, they are generating significant drag, which must be overcome. This uses power, which must be supplied somehow. In a pure rear-drive car, the power to drive the front wheels is delivered from the rear wheels via the road surface. That means that this portion of the car's power passes through the contact patches, and uses up a portion of the tires' friction circle or performance envelope. If we power the front wheels, even moderately, by some other means, we have more grip available for cornering. This explains why cars with all wheel drive perform as well as they do, even in pure cornering, despite their inherent weight penalty.

The basic front/rear torque split is generally established by a planetary center differential. The planet carrier is driven in some manner from the transmission output shaft. The ring gear or annulus drives the rear diff. The sun gear drives the front diff. The torque ratio is the same as the pitch diameter ratio of the sun and annulus. If the planet gears are the same diameter as the sun, the diameter ratio is 3:1. The torque split is then 25/75. To get a 20/80 split, the planets need to be 1.5 times the diameter of the sun. Planets twice the size of the sun yield a 5:1 annulus to sun ratio, or a 16.7/83.3 split. To get a 10/90 split, the planets need to be 4.5 times the diameter of the sun. To get a 40/60 split, the planets need to be ¼ the size of the sun. It will be apparent that there will be practical limitations precluding extremely unequal torque splits, or ones between 40/60 and 50/50. (Actually, splits between 40/60 and 50/50 are possible, if we are willing to machine a conventional-style bevel gear differential with unequal-size side gears, cut with different bevel angles, and the pinion gears cocked at an angle. Opposing pinion gears have to be on independent shafts, rather than a common shaft running through the carrier.)

Usual practice is to design a drivetrain so that the front and rear pinion gears and propshafts run at equal speeds when there is no wheelspin and the car is running in a straight line. This minimizes wear on the center diff. However, it is possible to design the center diff to withstand having its gears in relative motion most of the time, and have the propshafts run at different speeds.

Running the propshafts at differing speeds has two effects: first, the split in propulsion forces changes, because we are using different ring and pinion ratios at the front and rear, and the torque multiplication at these gear sets is correspondingly unequal. The end with the faster-spinning prop shaft gets more power, for any given torque split at the center diff. Second, the car's response to center diff locking devices changes. It is possible to have the rear wheels slipping or overrunning the front wheels by a predetermined percentage when the center diff locks, rather than running at the same speed. This allows us to tune the car's throttle-steering characteristics. Ordinarily, we would arrange for the front propshaft to overrun the rear most of the time, and lock the center diff when the propshaft speeds equalize, or when the rear one overruns the front by some amount.

With electronic control of the lockup, we really don't need to juggle the front and rear final drive ratios. We can pretty much get whatever torque distribution and rear overrun we want, just with the center diff and the lockup strategy. However, there is some effect on car behavior when the center diff is completely or largely locked, even if the lockup is electronically controlled. With simpler, passive, mechanical lockup devices, juggling the front and rear final drive ratios offers some interesting possibilities.

For example, it would be possible to set up the ratios so that the front propshaft overruns the rear by 30%, and add a roller clutch that locks the front propshaft to either the planet carrier or the annulus whenever either of these tries to overrun the sun gear, but freewheels the rest of the time. That would allow the driver to spin the rear wheels enough to throttle-steer, but prevent runaway rear wheelspin. Front to rear propulsion force ratio short of lockup would be sun/annulus ratio, times the ratio of front final drive ratio to rear final drive ratio, times the ratio of rear tire effective radius to front tire effective radius.

Another possibility is to drive the front wheels hydrostatically or electrically, rather than mechanically. This could be a way of powering the front wheels very modestly, with far less hardware and more attractive packaging. One way would be to use a small gear pump, perhaps one stage from a dry sump pump, as a motor on each front wheel. These might be driven by a similar pump at the transmission, and the circuit might include a cooler for the fluid. It might also include a relief valve that would limit pressure to the motors.

It would only be possible to transmit small amounts of power this way, but we could eliminate or reduce the power otherwise transmitted through the contact patches to overcome the drag of the front tires when operating at a slip angle. If we can do that, we get the main cornering advantage of mechanical all-wheel drive, with a smaller weight penalty.

Hydrostatic motors and pumps are available in a wide range of sizes. However, power losses in hydrostatic drive systems are high, so transmitting large amounts of power hydrostatically is not an attractive proposition for a high-speed vehicle.

Electric drive, as used in locomotives, offers some of the same possibilities as hydrostatic drive, although it cannot be made a part of a transmission cooling circuit. Also, electric motors are more delicate than a gear motor, so mounting them outboard is probably not a good idea. Electric drive does lend itself nicely to computer control of torque distribution.

Mounting both motors and brakes inboard offers some interesting advantages. Inboard brakes do add weight and cost, and they occupy space. So does the ducting required to cool them. However, inboard brakes do offer advantages not commonly recognized. Everybody knows they save unsprung weight. But additionally, they make it easier to get balanced airflow to both faces of the rotor, reducing the rotor's tendency to dish as it heats up. Potentially, they allow a larger rotor diameter, at least in some cases. This is particularly true in low-built race cars, where the floor hangs below the wheel rims. Finally, they allow relatively unobstructed airflow through the wheel. In a full-bodied car, this greatly increases the effectiveness of the wheel openings as ducts for extracting under-car air.

Inboard front brakes do require beefy halfshafts. The torque transmitted can easily be as great as that transmitted by the rear wheels under power, even in a tail-heavy car, and the consequences of shaft or joint breakage are really nasty.

Either electric or hydrostatic drive to the front wheels offers the possibility of integrating the front drive system with an energy recovery system.

So much for the pros, cons, and nuances of driving more than two wheels. What about having more than four wheels on the vehicle?

In the case of tires, more rubber generally gives us more grip, at a penalty in speed on long straightaways, for reasons I have discussed in previous writings. We are limited, however, by our ability to keep a wide tire upright enough to use the full width of the tread. As the tire becomes really huge, suspension packaging can become problematic. So can steering geometry, in the case of a front wheel. Frontal area of the tire increases in direct proportion to width, or height, with a corresponding aerodynamic drag penalty. If there is standing or streaming water on the road surface, a wide tire is more prone to aquaplaning. If we are going to have to deal with snow, a wide tire is a disadvantage because of the greater force required to move it through the snow.

So there are reasons to consider having two small tires in line with each other, instead of one big one, or perhaps two big ones in line with each other to provide twice as much tire, or some compromise in between.

Two strategies immediately present themselves: double the number of wheels, and use eight, or make the car heavy at one end, and use four at the heavy end and two at the light end.

Those who have studied Formula 1 history will recall the Tyrrell P34 6-wheeler of 1976, which had the extra wheels at the light end. The idea there was to use much smaller front tires than usual, and reduce aerodynamic drag. This required special tires. Conventional, large tires were used at the rear. The reasoning was that the airflow was so disrupted by the time it reached the rear wheels that small tires there would not return as much benefit. Also, it is easier to add more non-driven wheels than to create the hardware to have more driven ones. The Tyrrell actually won The Swedish Grand Prix that year. In fact, the team got a 1-2 finish.

Because rear grip was essentially the same as a conventional car, there was no benefit in exploiting the potentially greater front grip from the four front tires. F1 cars of the time already cornered with the inside front very lightly loaded, so there was little room for adding more front roll resistance to use enhanced front grip to help the rear end stick. Consequently, the advantage of the design was mainly aerodynamic, and this came at a penalty in weight.

Shortly following the Tyrrell effort, March built a 6-wheeler the more obvious way: four rears, all driven. The transaxle was an adapted Hewland, with no center differential. The car was tested but never raced. March was in financial straits at the time, and development was lacking. In 1979, a retired March F1 car with the 6-wheel setup achieved considerable success in British hillclimbing.

There was also a similar car built by Williams in 1982, which likewise was never raced. After this, the FIA banned all four-wheel-drive systems from F1.

It would be possible to have four rear wheels and only drive two. The car wouldn't throttle-steer as one might hope, and propulsive traction would be poor, but it might be possible to get good lateral acceleration numbers that way, and perhaps better braking than with four wheels.

The idea of using four front wheels, and driving the frontmost pair from a separate engine, with a automatic transmission, is interesting. One drawback would be that the torque distribution would vary in a largely uncontrolled manner, because the torque multiplication at the automatic would not have a constant relationship to the torque multiplication in the manual transmission. The two transmissions would have differing ratios and wouldn't shift at the same time.

The questioner asks about steering four front wheels, and wonders about using some form of balance bar. I don't think a balance bar is appropriate for the steering. All four wheels need to have a fixed relationship to handwheel position. I think Tyrrell got it about right. As I recall, they steered the frontmost pair of wheels from a rack and pinion in the usual manner, and ran a drag link back from each of these to the wheel behind it. It would be normal practice to make the second wheel pair steer a bit less than the first. This could be considered a form of Ackermann effect. For conventional, or positive, Ackermann effect with four front wheels (i.e. to optimize for a turn center behind the front wheels), the second pair of wheels should steer less than the first because they track inside and behind the first pair. This is certainly the way it's done in heavy trucks with multiple front axles.

We have noted in previous newsletters that we want different dynamic toe-out for high-speed turns than for low-speed turns, and possibly even toe-in for the high-speed ones. I don't think we have this kind of reversal regarding how much the second wheel pair steer compared to the front pair. The second pair should steer less than the first pair whether the turn center is behind the front wheels or ahead. (At least, this is true as long as the front wheels are steered into the turn, as opposed to counter-steered to correct a rear wheel slide.) We might say that this is so because the second wheels always trail the first ones, whereas the outside wheel may either lead or trail the inside one.

Tyrrell did use balance bars on either side for springing and damping, and it seems to have worked okay. The first and second wheels on each side were linked by a balance bar, and a single pair of coilovers and a single anti-roll bar acted on these balance bars. This meant that, considered as a four-wheel group, the front wheels had a wheel rate of zero in pitch and warp, and were also undamped in these modes.

I'm not sure a pure balance beam arrangement like Tyrrell's is fully optimal, but the fact that it worked acceptably points up a potential advantage of having two close-coupled pairs of wheels rather than a single pair: we can make the system very soft in absorbing small bumps, while still having good control of sprung mass pitch, roll, and heave.

Does it make sense to have eight wheels, then? It might, if the application makes traction a priority over power-to-weight ratio and compactness. The optimum might be a rear-mid-engine configuration, with the rear four wheels driven mechanically and the front four driven hydrostatically or electrically, and inboard brakes on all eight. The most advantageous course for such a vehicle would have low grip and plenty of chatter bumps, and be uphill – but not be confining enough to make compactness too crucial. And of course, the rules would have to be set up to encourage innovation, rather than create a drivers' class with standardized or closely matched cars.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANOTHER EXOTIC CAR LAYOUT

A reader sent this link, with the following brief comment:

http://sitepalace.com/donmichael/reece/reece.html

Don't know what magazine this came from. If it's Racecar, you've probably seen it. Interesting chapter in circle track history!

The article is from *Open Wheel*, which was a companion publication to *Stock Car Racing*. What tipped me off was the mention of Doug Gore being the technical editor. I remember him being technical editor for both of those titles. He is now in a similar position at *Speedway Illustrated*. I asked Doug if he remembered the article and he said he did. *Open Wheel* ceased publication eight years ago.

The car shown is a supermodified with three right wheels and one left wheel. It was called a 3-to-1 car. It was built in 1979 by Kenny Reece. It was tested, but was outlawed before it ever got a chance to race.

The left wheel and the center right wheel were connected by a beam axle, and driven by the engine, which was an aluminum big-block Chevy. The front and rear right wheels steered. The rear one steered out when the front one steered in, but by a lesser amount.

This idea could be regarded as an oval-track variant of the Scalabroni diamond layout discussed in the March 2007 newsletter. The end wheels are just offset all the way to the right.

Tim Richmond drove the car in testing. The lateral forces generated were so great that his helmet strap broke. (I can't tell if this refers to the one under his chin, or an extra one to the side to keep his head upright, which some drivers used at the time.) The thin-walled fuel tank began to bulge laterally on the right side.

This is on a flat, half-mile, fully round skidpad, at Honda's Marysville, Ohio transportation research center. The car had no wing of any kind. It didn't even have a rear spoiler. It had a rounded tail with a tailfin, reminiscent of an Indianapolis car or a Detroit styling exercise of twenty years earlier. There can't have been any aerodynamic downforce at all. The tires were all twenty inches wide, and the total car weight was 1,315 pounds. According to the article, 75% of the weight was on the single left wheel, in static condition.

Even considering the engine being almost half of the mass, I'm a bit skeptical of the 75% left figure. Also, I'm sure all the weights are without driver. The driver is a bit to the right of center. I'd guess the car might have as much as 65% left, with driver. Supermodifieds have been built with 75% left, but they have wider tracks and the right wheels hang out considerably further from the main structure than in this car.

The big worry I'd have about such a layout is, what happens if you have to suddenly slow down and turn right at the same time? You don't have to do that to turn a fast lap on an oval, but there are times while racing when you have to turn right and slow suddenly to avoid a wreck. This is worrisome enough in a conventional supermodified with two left wheels.

I would also expect the car to veer rightward pretty heavily in hard braking. The article does not explicitly say this, but most supers run spools, so any braking on the drive axle slows both of those wheels. Then we have the other two wheels, acting entirely on the right side. In straight-line braking, the left wheel has to contend with about 65% of the inertia, but it only has something like 25% of the braking power. Even with two left wheels, it is difficult to keep a supermodified straight while braking hard, or make it turn left while trailbraking. The heavy side wants to keep going.

As with the diamond layout, trying to tune the handling properties by adjusting the wheel loads and spring rates presents some interesting quandaries. In the pictures, it looks like the three right wheels are not perfectly in line; the middle one appears to be set inboard slightly. But let's imagine that the right wheels are all in line, for simplicity.

In such a car, raising or lowering the left wheel suspension wouldn't change any of the wheel loads. Raising or lowering all three of the rights the same amount also wouldn't change any of the loads.

A conventional car has percentages for the ends, the sides, and the diagonals. The 3-to-1 car has side perccntages, a middle percentage, an end percentage (front and rear right combined), and separate front and rear right percentages.

If we raise just the right front, we increase the end percentage and decrease the middle percentage. All the load comes off the right middle wheel. No load comes off the left wheel, although its ride height might change a little bit, depending on exactly where we measure.

What does this do to the handling? Does it do anything? It probably changes it some, at least with the car in the article. It might not, if the car's center of gravity were right on the axle. However, the Reece car appears to have its c.g. well forward of the axle. The axle is somewhat aft of the wheelbase midpoint, and both the engine and the driver are forward of it. The only major mass aft of the axle is the fuel load.

If the car had very little end percentage, and a lot of middle percentage, we would expect it to push, or understeer. The centrifugal inertia would act forward of the tires with most of the load, and the front tire would have little ability to resist this.

Conversely, if there is less load on the axle, tire stagger has less influence. Usually, oval track cars with spools are over-staggered, compared to the stagger theoretically required for the turn radius of the racetrack. Consequently, anything that unloads the middle wheels tightens the car by reducing stagger-induced oversteer.

This could get tricky.

And then there's the steering on the rear wheel to play with.

Of course, this is all largely academic, since there is no place to race such a car. This case illustrates a point I have made in the past regarding highly visible or radical innovation: you have to think not only about whether your idea is legal under existing rules, but also how the rule-makers are likely to react when they see what you've come up with.

I do think there might be a role for really unusual cars like this as exhibition acts between races for legal cars, sort of the way jet cars have been used in drag racing. I'd probably turn out if I knew my local short track was going to let a car with three right wheels do some laps and see how fast it could go.

Now, if we wanted to build a no-holds-barred exhibition car that would break track records on any oval short track it showed up at, might we want to use more than two right tires? I think so. But I think we'd also want at least two lefts. Additionally, we'd want powered ground effects – some form of sucker car. If you're going to throw away the rule book, throw it all away!

Some readers will remember the Chaparral sucker car. It had two big exhaust fans in the tail, driven by their own two-stroke engine, sucking air from under the car and blowing it out the back. It had vertical Lexan skirts that dragged on the ground, to help exclude air from the underside. The skirts had slots where they attached to the body, so they could slide up and down. Their own weight was all that held them down. The skirts on the Chaparral were clear. I think if you had multiple right wheels, you'd want opaque skirts, so nobody could see how many wheels there were. Black would be good. In fact, it would be good theater to paint the whole car black, for that extra aura of mystery. And keep it under a cover whenever it wasn't actually running. People would go nuts trying to figure out what was underneath.

An exhibition car doesn't need a number, but you could put something on the side. An infinity symbol perhaps? Or maybe 666?

If the car only has to do a few quick laps, electric power starts to look really attractive. The batteries could be fairly light, and all to the left. All or most of the wheels could be driven by individual motors. The sucker fans could also be electric. The car would obliterate lap records, almost silently.

It might also be pretty good at drying and cleaning the track. That could be its job, once more advanced exhibition cars put it out to pasture.

My guess is that pretty quickly we'd reach a point where the driver's ability to withstand the accelerations would become the limiting factor. Also, if large numbers of people started building outlaw exhibition cars, the market for between-races exhibitions might become saturated. But at least in the US, there are a lot of short tracks that need a draw. I think a lot of people could have lots of fun before one factor or another called a halt to the party.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

LATE CORVAIR REAR SUSPENSION

I know it has been a while since you spent any effort analyzing the Corvair suspension (addressing the late model, fully independent only) but I wondered about your thoughts on the general layout and where improvements could be made. Most of the rulebooks now allow a lot more "deviations" from stock than they used to. I am building a late Yenko "Ringer" (looks like a Stinger) to do some open-track work. I can describe the suspension changes pretty easily: stiffer urethane bushings in the front end; a coil cut to lower it. The one variable I am still working on is the front mounting of the rear trailing arms. Stock had a rubber bushing with an auxiliary link part of the way back on the trailing arm, tied into an inboard location on the transmission mount. The vertical location of that inner pivot point determines the L/R movement of the trailing arm during roll. Chevy had at least two different mounts, attempting to add a bit of roll "Turn-in rear steer" probably to counteract the natural rear engine tail happiness of the Corvair. Several folks have just removed all the toe-control apparatus, replaced the rubber in the trailing arm with a Monoball, lowered the car, set the toe to some spec and gone racing. With fairly stiff anti-roll bars, the car still leans some in corners, but was pretty drivable on the road courses. Do you think the Monoball replacement, and the giving away of any additional toe change is a worthwhile tradeoff?

A picture of the Corvair suspension, scanned from the GM shop manual, accompanies this newsletter. This happens to be a car I know well, having owned one since 1976 (with a Crown V8 conversion since 1978).

For readers less familiar with the car, the Corvair is a rear-engined sedan (though with a fairly cramped rear seating area), produced by Chevrolet from 1960 through 1969. The engine is an air-cooled flat 6, hanging out behind the rear axle as in a Porsche 911. This puts around 62% of the weight on the rear wheels, varying somewhat with equipment and fuel load, and gives the car a lot of yaw inertia.

There are two basic designs of Corvair: early (1960-1964) and late (1965-1969). The early and late models have dramatically different styling. I call the early version the sardine-can body, for its somewhat brick-like shape with protruding lip all around the beltline. I call the later model the Coke-bottle body, for its curvacious fender line. There are also pickup truck and van models, called forward-control vehicles, similar in layout to VW transporter models, which were based on the early platform and were made from 1960 through 1965.

The early platform had swing axle rear suspension, and earned itself a bad reputation for not only having snap oversteer at the limit like a VW or Porsche, but also having an unusually compliant spring and tire combination from the factory that allowed the outside rear wheel rim to actually contact the road in extreme conditions and sometimes hook an irregularity and flip the car. The early models can be improved dramatically with lower ride height, stiffer springing, and good tires. Really, they like a lot of ride stiffness, especially at the rear. The factory attempted to add ride rate at the rear, and roll rate at the front, in 1964. At the rear, they added a transverse leaf spring acting only in ride, similar to the aftermarket EMPI Camber Compensator, and they added an anti-roll bar at the front. Porsche also used a swinging transverse leaf on the last versions of the 356, and Mercedes-Benz had a coil spring version of the same concept on some swing axle rear ends, as early as the 1930's. The "swing spring" on late-model Triumph Spitfires is another version of a swing axle system with a higher wheel rate in ride than in roll. "Zero roll" rear suspensions on swing-axle Formula Vees take the concept to its logical conclusion: no elastic roll resistance at all; springing that acts in ride only.

In any case, the early models have a handling problem. It can be fixed, at least somewhat and for certain purposes. Set the early cars up low and stiff, and they work much better than you'd expect, especially for low-speed events like autocross and hill climbing that favor a loose (oversteering) car. This is a case of any suspension working if you don't let it.

By 1963, when people started suing them, GM realized there was a problem, and set about remedying it for the next version of the car. The redesigned rear suspension was similar to the system introduced in 1963 on the new Corvette Sting Ray, and found on all C2 and C3 Vettes. In the transverse plane, it's a short-and-long-arm (SLA) system, using the halfshaft or driveshaft as the upper arm. In side view, it is a trailing arm system. The "torque control arm" is stiff in bending both vertically and horizontally. It does actually react torque in braking, although not under power. It also locates the wheel longitudinally and for toe.

In the Corvette version, springing is by a non-swinging transverse leaf spring, and the toe-controlling "front strut rod" shown in the Corvair suspension illustration is absent. The Corvette has a fairly serious deflection steer problem under power, due to compliance in the bushing at the front of the torque control arm. This bushing is much stiffer than in the Corvair, but it is rubber and it has to accommodate rotation in three axes. Fortunately, the wheels toe in under power rather than out, but the deflection is greater than one would wish, especially when the car is fitted with wide, sticky tires and large-offset wheels. Racers running C2 and C3 Corvettes eliminate the bushing deflection by substituting spherical joints (monoballs) for the bushings. The bump steer properties are not perfect either way, but they aren't horrible and deflection steer is largely eliminated. People racing late-model Corvairs often do the same thing, as the questioner notes.

The Yenko Stinger was a mildly modified Corvair, with the back seat removed and fiberglass C-pillar extensions to visibly differentiate it from a stock Corvair, and make it look more like a coupe and less like a sedan. The reason people road raced these and not regular Corvairs had to do with what SCCA was willing to recognize as a sports car. If a car had a back seat, that a small adult could squeeze into at all, it was considered a sedan. Sports cars didn't have back seats. If the owner removed a sedan's back seat, that didn't make it a sports car. However, if the car came from the manufacturer that way, then it was a sports car. Yenko was the manufacturer of record for a Stinger, rather than GM, and they took the back seat out, so that made the car a sports car. Shelby did the same thing with the Mustang when creating the original GT-350.

There weren't very many Stingers made, and they have become scarce with the passing years, so now SCCA is allowing people to make reproduction ones out of old Corvairs. That's what a "Ringer" is.

As to what we might do if the rules are liberalized, a great deal depends on exactly what those liberalized rules allow, and also what the budget allows.

Really, the biggest potential improvement isn't a suspension modification at all. The thing the Corvair needs the most, that I see most people missing, is bigger tires on the rear. The engineers evidently knew this, because they left extra room in the rear wheel wells. They couldn't actually put larger tires on the rear from the factory, because GM had a corporate policy that all cars had to have one spare tire that would work on all four corners. But there was no rule that they couldn't leave room for enterprising owners to install larger rear tires after purchase, so they did. Why Yenko didn't seize the opportunity has always mystified me. My recollection from road test reports back in the day was that street Stingers had the same size tires at both ends. All the Stingers and Ringers I've seen race have wider tires than stock, but the rears are the same size as the fronts. Perhaps this has to do with Yenkos being sold through Chevy dealers, and SCCA not allowing unequal size tires for racing if the street model didn't have them. At any rate, even when the car is set up so the inside front wheel lifts, it has visible limit oversteer. This works okay for autocross, but for full-scale road courses the car is in crying need of more rear rubber.

The Corvair came with really tiny wheels, I think 13 x 5½. Mine already had 14 x 7 chromies installed by the previous owner, with 185/70 tires. With zero offset (4" back space), these bolt right on at both ends of the car. I left those on the front, and tried 15 x 8 Corvette wheels at the rear, with 215/60 tires. This combination also bolts right on, at the rear only, clears the fenders in all conditions, and absolutely transforms the car. Camaro Rally wheels on the front and Corvette Rally wheels on the rear make a readily available, Chevy-looking package.

I still have antique aftermarket wheels on my car: heavy but nice-looking Shelby Vector cast aluminum, 14 x 7 on the front with 195/55 tires, 15 x 8½ on the rear with 225/50 tires. These rears only clear the fenders if I run about 3 degrees negative camber, and more ride height than you'd want for a track car. My wheels have 4" back space front and rear. There looks to be at least an inch of room available on the inboard side of the rim and tire, both front and rear. With more modern wheels, it should be possible to get at least 9" of rim width at the rear, and 8" at the front, without modifying the bodywork.

What diameter those wheels and tires should be depends on a number of factors. Some of these are straightforward packaging issues. Others have to do with the limitations of the suspension design.

To locate an upright using only tension/compression links, we must use five links. For a rear suspension, this normally means two longitudinal links and three transverse ones, or something approximating that. We can use fewer parts than that, but then some of them have to resist bending loads. The Corvair suspension has the three transverse links, with the halfshaft as one of them, but the designers simplified the longitudinal location by using a rigid beam, pivoted about a single point, anchored rigidly to the hub carrier, and loaded in bending.

This compromises the side-view geometry, at least when using outboard brakes. Because brake torque reacts through the arm but drive torque does not, the system's longitudinal anti's are dramatically different for lift under braking and squat under power. We would like to have moderate anti-squat under power, and moderate anti-lift under braking. With this design, and outboard brakes, that is not possible. The only way to keep the anti-lift halfway within reason is to have the pivot somewhat below the hub. However, that gives us pro-squat under power, so we can't lower the pivot too much. Making the arm longer improves the compromise, but the arm starts to get heavy, and packaging constraints limit us.

If we make the wheel taller, and lower the suspension to keep a given ride height, the anti-lift stays about the same, but the pro-squat increases, which is not good. We can live with a lot of anti-lift if our brake bias is such that we don't approach rear wheel lockup. We can live with squat under power if we have enough ride rate and available compression travel.

The transverse or front-view geometry presents us with a different picture. Here, the designers saved parts by using the driveshaft as a suspension link. This saves both cost and unsprung weight, but it entails geometric compromises. If the U-joint angles are to be kept moderate, the halfshaft needs to be fairly close to horizontal. If the halfshaft is horizontal, and we want 50% camber recovery in roll, the front view instant center for each rear wheel is at the center of the opposite wheel, and the roll center height is around half a loaded tire radius, or about 6 inches. That makes the force line angle of elevation roughly eleven degrees. That's enough to produce visible jacking if the tires are sticky.

The geometry gets better as we raise the hub or lower the car. Unfortunately, the U-joint angles rapidly increase. Up to a point, we can live with that, especially if power is modest, but there is a penalty in frictional losses even if we can keep the joints alive.

The C2 Corvette had front-view geometry much like the Corvair. With the skinny tires that car had, the jacking was not too noticeable, but when people started adding rubber, it became a problem. For the C3, Chevrolet lowered the inboard strut rod pivots. This increased the front view swing arm length, lowered the roll center, and reduced jacking. Crown offered a strut rod mounting bracket for the Corvair that does the same thing.

The problem then is that the suspension loses camber recovery in roll. This is usually addressed by having lots of roll stiffness. Camber change in ride is reduced, which is good in a straight line.

As long as we keep the halfshafts as the upper control arms, we face a three-way compromise among U-joint angularity, camber recovery, and geometric anti-roll. We can't optimize all three at once.

A kit for improving C2's and C3's is still sold by Guldstrand Motorsport (see at http://www.guldstrand.com/scripts/prodView.asp?idproduct=80). It uses the original Vette brake and hub carrier, and adds a bolt-on upright, two transverse lower links, and two parallel trailing links, all with Heim joints. The halfshafts are still used as suspension links, and the lower links are nearly horizontal, meaning the compromise is weighted in favor of roll center height, at the expense of camber recovery, as in the stock C3 layout. The trailing links run slightly uphill toward the front, giving moderate anti-lift and anti-squat.

It would be possible to make yourself a similar setup for the Corvair – perhaps even use a Guldstrand kit with minor modifications.

Better yet would be a suspension that doesn't use the halfshafts as suspension members. We would like a front-view instant center around 70 inches from the wheel, and five inches above the ground. This requires the upper link to slope down toward the middle of the car, more steeply than the halfshaft can slope and still transmit power. The halfshaft then has to use splines or tripod joints, so it can accommodate plunge.

Such modifications should be legal anywhere that a Guldstrand suspension is legal on a Vette.

While we're considering cost-no-object modifications, how about scrapping the Saginaw 4-speed and using a Hewland, Porsche, or ZF transaxle? Once you're redesigning the suspension and halfshafts, it's tempting to throw out everything from the flywheel to the crossmember and just start over.

Of course, that's expensive, and a cost-no-object Corvair-powered Corvair really is an exercise in fanaticism. It's probably no crazier than a high-dollar Camaro or Impala, though. People build those, and in my opinion the Vair has nicer lines. I don't have a very high opinion of the stock engine, but it does have a low c.g., and it makes a nice noise. For those of us who have more limited means, the stock hardware really is not so bad, and it is reasonable to ask what is the best low-buck approach.

I like keeping the toe-control link and the big, soft bushing at the front of the torque arm. It makes sense to take compliance out of the toe link. You can make one with rod ends, or you can use the stock one, with urethane or nylon grommets instead of rubber. The grommets are ordinary anti-roll bar drop link grommets. To use four of them per link, you may have to cut them down a bit, or you may not be able to start the nuts – or you may be able to start them but not get safe thread engagement. The stock rubber grommets squash down a lot when you tighten the nuts, and they stay squashed if they've been in for a while. If you use stock grommets on the inner side of the torque arm and outer side of the crossmember bracket, and hard ones on the outer ends of the link, everything goes together with no cutting of anything. You get a link that allows compliance toe-in more easily than compliance toe-out.

When you lower the inboard end of the link, the wheel toes in as the suspension compresses. This provides roll understeer. To some extent, the understeer this creates isn't real at the contact patches, or at least the effect is smaller than the driver thinks based on required steering wheel input. However, there is some actual effect even in terms of slip angles, at least as long as we're applying power, which we usually are when cornering hard in an oversteering car. When the rear wheels are aimed more toward the inside of the turn, the car's c.g. is further toward the outside of the turn with respect to the rear wheels' lines of thrust, and that actually does add understeer.

Having the wheels steer in any way with suspension motion is at best a mixed blessing. One thing I've noticed when running ample roll understeer is that you get heave oversteer when you go over a crest while cornering hard. The rear steps out as it comes up on the suspension, then comes back in as the car settles back down. This mostly just feels disconcerting, but it does require a steering correction by the driver, and if you're really on the limit, it could cause you to lose the car.

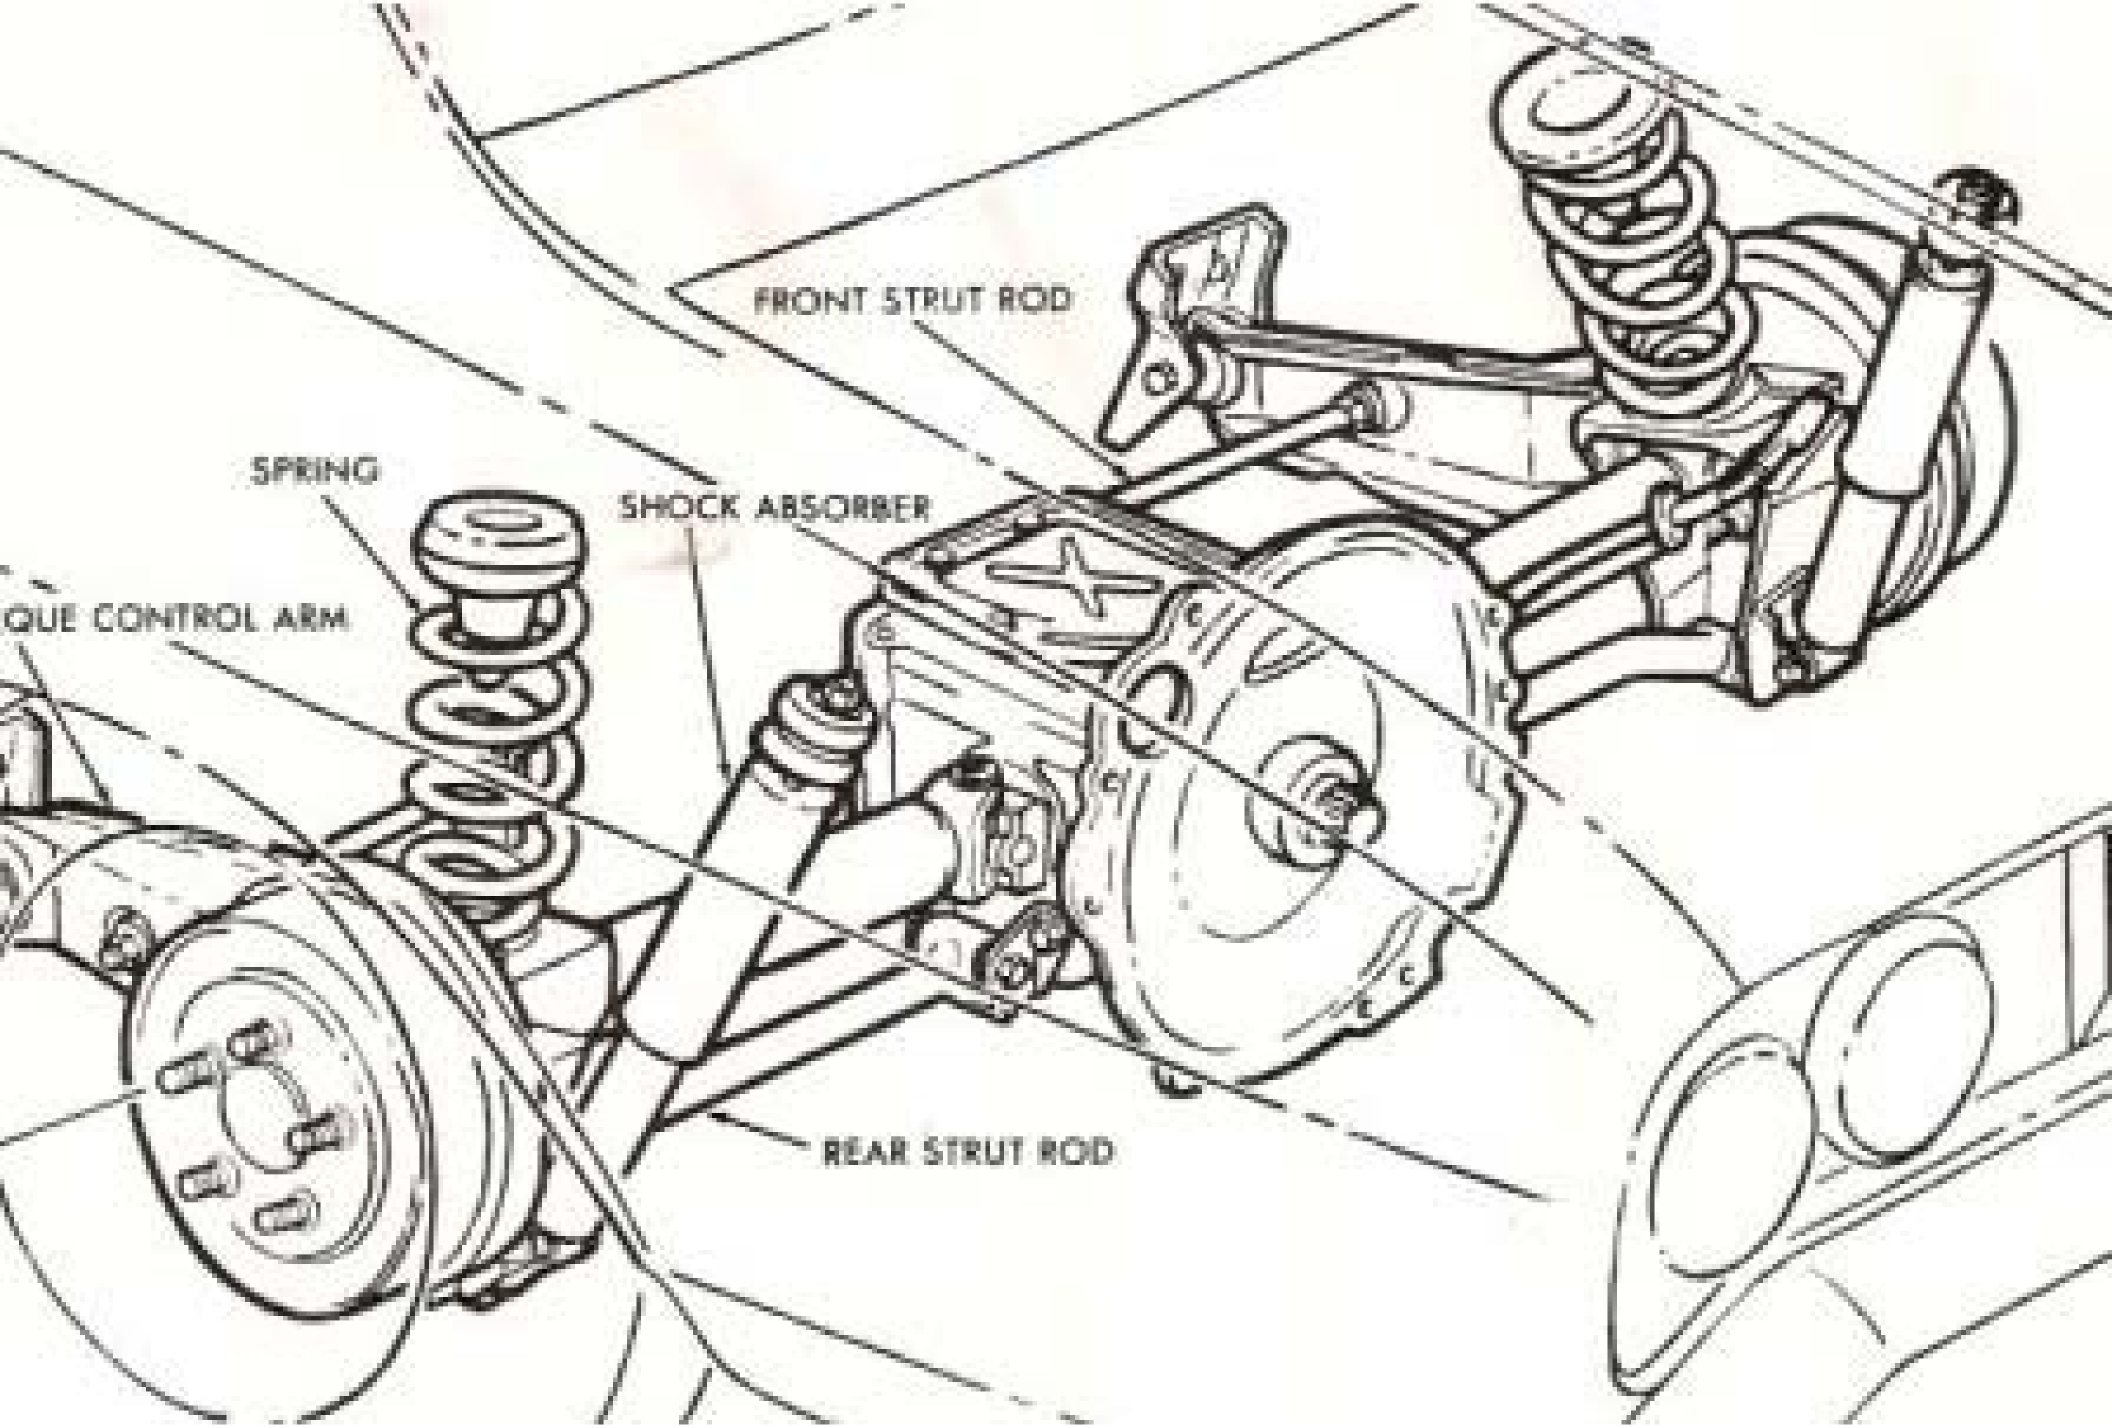
FRONT STRUT ROD
SPRING
SHOCK ABSORBER
QUE CONTROL ARM
REAR STRUT ROD

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANOTHER PONY CAR SUSPENSION FROM THE '60'S

Last issue we took up the independent suspension on the later-version Corvair and the C2/C3 Corvette.

I have received e-mails from a number of correspondents regarding the form of independent suspension shown on pages 61 and 62 of the October 2007 *Racecar Engineering*, in Richard Nisley's article entitled "the Program Makers". This is an independent suspension using a reverse wishbone as one of the links, to control toe. The article states that the system shown was used on the GT-40, and a number of other cars. This idea was invented by Ford engineer Klaus Arning and patented by Ford in 1958. Correspondents have inquired what I think of the system, as shown in the magazine. One of them sent me a photo of the system shown on p.61, in place under a Mustang. I am sending that with this newsletter, as a separate attachment.

This suspension is described as a 4-link system. As we noted last month, it really takes five links to locate an upright, if by a link we mean a simple rod or tube with joints at the ends, loaded purely in tension and compression. To reduce the number of locating members to four, it is necessary for one of them to be designed to take bending loads. The Arning system uses what has come to be known as a reverse wishbone to accomplish this. The lower control arm cosists of two converging tubes, with two joints at their outboard ends and a single joint at the inboard end.

I haven't read the actual patent, but I gather that what is claimed in it is the combination of a reverse wishbone and two longitudinal links, which can be either leading or trailing. The illustrations show a version with the lower link trailing and the upper link leading, so that in side view the system looks something like a Watt linkage for a beam axle. The photo on p.61 shows an outboard brake. The drawing on p.62 shows an inboard brake.

I was surprised to learn that the reverse wishbone idea was patented by Ford, because it was first used on a race car by Lotus – but not before 1958. Lotus was using it in 1960, on both the Mk. 18 F1 car and the Mk. 19 sports racer. In fact, all rear-engined Lotuses used it, for at least ten years, and by 1962 it was pretty much universal in F1 and for rear-engined sports-racing cars. The earliest use by Ford was on the V4-engined Mustang I prototype, which as I recall was first shown in 1962. By the time Ford showed the public a running car with reverse wishbones on it, racers everywhere were building cars that way, and a number of small companies were building them for sale.

I don't know of any lawsuits seeking royalties from any of these small builders. Some of the cars used Ford engines. Others didn't. It would be interesting to know exactly who communicated what to whom during that time. We can't necessarily assume that only one person came up with the idea. Sometimes when an idea appears in different places at about the same time, it really was one person's idea, and they communicated it. In other cases, the collective evolution of scientific thought leads totally independent individuals to similar conclusions nearly simultaneously.

However, all successful suspensions of this type have used two trailing links for longitudinal location, rather than the version with the leading upper link shown in the pictures. The system shown was used only on a prototype. It was never raced or mass-produced.

As a way of adapting independent rear suspension to a car originally designed for a live axle on leaf springs, the layout is appealing. It can anchor to the frame or unibody rails in approximately similar places. There are no upper trailing links intruding into the rear seat area.

However, for the layout to work as it should, the brakes have to be inboard. The reason why provides a good example of the difference between a case where torque acts through the suspension linkage and a case where the linkage only reacts thrust.

It is possible to arrange the geometry so the hub moves in a vertical straight line in side view. In that case, anti-squat under power is essentially zero, although it may go to a small positive or negative value with pitch. With inboard brakes, the anti-lift will likewise be zero, with small variations due to pitch.

If the links are parallel at static condition, the side view instant center (SVIC) is undefined in the x or longitudinal direction. However, it can be said to have a z or vertical location, at hub height. As the suspension compresses, the SVIC appears behind the wheel, at hub height. With further compression, it remains at hub height but rapidly migrates forward. A line drawn from the hub to the SVIC remains horizontal. But a line drawn from the contact patch center to the SVIC slopes downward below ground level toward the front of the car, and this slope increases rapidly as the suspension compresses.

The former line, from the hub, determines the anti-lift in braking when the brakes are inboard. The latter line, from the contact patch, determines the anti-lift when braking. So with inboard brakes, the anti-lift stays pretty much constant, but with outboard brakes it varies dramatically as the suspension moves. As the suspension compresses, it rapidly gains pro-lift. As it extends, the effect works the other way, and the suspension gains anti-lift. This means that when a bump compresses the suspension during braking, the force trying to jack it back up rapidly grows. As the suspension extends during braking, a force trying to jack it back down rapidly grows. The changes in jacking force oppose suspension compliance both ways.

I generally take issue with those who suggest that longitudinal anti effects lock up the suspension during braking, because if the jacking force remains fairly constant as the suspension moves, variations in load due to bumps still move the suspension. The range within which the motion occurs shifts, but the motion is not resisted more strongly. However, this suspension presents us with a case where the jacking forces vary in the same direction as the spring forces, as the suspension moves. Therefore, the total support force from spring and linkage combined varies more with displacement than it otherwise would, and the suspension truly does become less compliant.

This is also a reason not to use a Watt linkage on a beam axle for reacting brake torque. It is important to mount the caliper to the axle housing or use a separate brake floater instead. It shouldn't be on the birdcage.

Interestingly, if the upper link is trailing and the lower link is leading, the anti-lift with outboard brakes changes the opposite way. The jacking forces change oppositely to the spring forces, and tend to amplify suspension movements. Conceivably, that could be beneficial, at least in moderation.

The correspondent who sent me the photo says that the layout originally used a Jaguar XKE (Dana 44) rear end with the Jag inboard disc brakes, but the brakes were moved outboard during development.

REALLY QUICK STEERING

While driving a street-car I became interested in the question of why it is we have multiple turns lock to lock for the steering wheel function. This necessitates the driver having to remove his hands from the wheel for tighter turns. It can be an issue in racing as well, especially in rallying and some types of open wheel cars (returning to the pits or negotiating some chicanes). Karts and motorcycles have sub-one turn steering. Both are known to be controllable and in the case of a motorcycle (or any bicycle) it is impossible to conceive of a rideable system with more than a fraction of a single turn (I had fun imagining what a three turn lock to lock motorcycle steering system would be like!).

My question is, why not have single-turn steering for cars (that is, the steering wheel would turn approx 360 degrees lock to lock)? It would eliminate the problem of removing hands from the wheel in order to make some types of turn. It would also give the driver a far better chance of catching a car that had entered a slide. He would be able to command a faster response from the steering system.

The main concern relates to straight line driving at medium to high speed. In this situation the vehicle is very (too) sensitive to small steering inputs. Drivers perceived this as "twitchiness" or *excessive response on-center. In cornering the problem goes away, as lateral loads are introduced which make the "gain" of a fast ratio steering system work in favour of controllability. In that case the car is perceived as responsive but not excessively twitchy (the driver is working against transverse forces in the steering system which effectively "balance" the input from the driver and "damp out" the twitchiness or excessive response). So the trouble appears to exist only on-centre at speed.*

I have an interest in installing a single turn system in a road car but can't think of the best way to sort out the excessive on-centre response at speed. The power assist system does not appear to be much help as even disconnecting the pump has little effect on what occurs above 15 mph anyway (no assistance required in the medium speed on-centre cruising regime). It's only at low speed you need the power assistance and at low speeds "twitchiness" is not present! What would you recommend?

Whether power steering is needed at speed depends on the car. Some cars require a lot of effort if the power steering quits. A two-ton car with one-turn steering and dead power assist could really be hard to control.

I drive full-size cars most of the time. Occasionally I get to drive a go-kart. When I get on the kart, I find it hard to be smooth, even when cornering. Practice definitely helps. A driver can get used to really quick steering, given time to acclimate.

One-turn steering without power assist is common on really light racing cars. Formula SAE cars very often have less than a turn. People would probably use kart-style non-geared steering in FSAE, if the rules didn't require geared steering. Assisted very quick steering is used in larger pure race cars, such as F1 and LMP cars. Sprint cars and midgets have very quick steering. In road cars, it's less common, but it has been done, notably by Citroen on the SM.

The SM had power-assisted centering, which varied with speed. This was done with a heart-shaped cam, with a hydraulically-loaded roller bearing against it. The hydraulic pressure was varied to control the centering force. The car had the peculiar habit of returning its steering to center by itself, even when stationary, as long as the engine was running. A correspondent who drove one in an autocross describes it as having horrendous understeer when used for this purpose. He surmises the designers were thinking more of straight-line superhighway cruising. Of course, this is mainly a handling issue, as opposed to a steering problem.

The combination I have found hardest to deal with is quick steering combined with a lot of lash (or play, or slop). Move the steering wheel a little, and nothing happens. Move it some more, nothing happens. Move it some more, and the car suddenly darts, and you need to correct the other way, and of course the same thing happens again. Even in this case, driver acclimation helps. You get used to moving the wheel rapidly until you start to feel resistance, and being delicate and precise from there, but it isn't as easy to do well as it is to describe. If you want to use very quick steering and have it driver-friendly, it's important to keep the slop to an absolute minimum.

It is possible to make steering gears that speed up or slow down as you get toward full lock. Nowadays, it is most common to make the steering speed up toward the ends, but for heavy cars with non-assisted steering, it was once common to make the steering slow down away from center, to reduce effort when parking. Even a fixed-ratio rack-and-pinion system gets faster toward the ends, due to the shortening of the moment arm length at the steering arms as the wheels steer.

It is also common nowadays for passenger cars, at least upscale ones, to have power assist that diminishes as speed rises, without going away altogether.

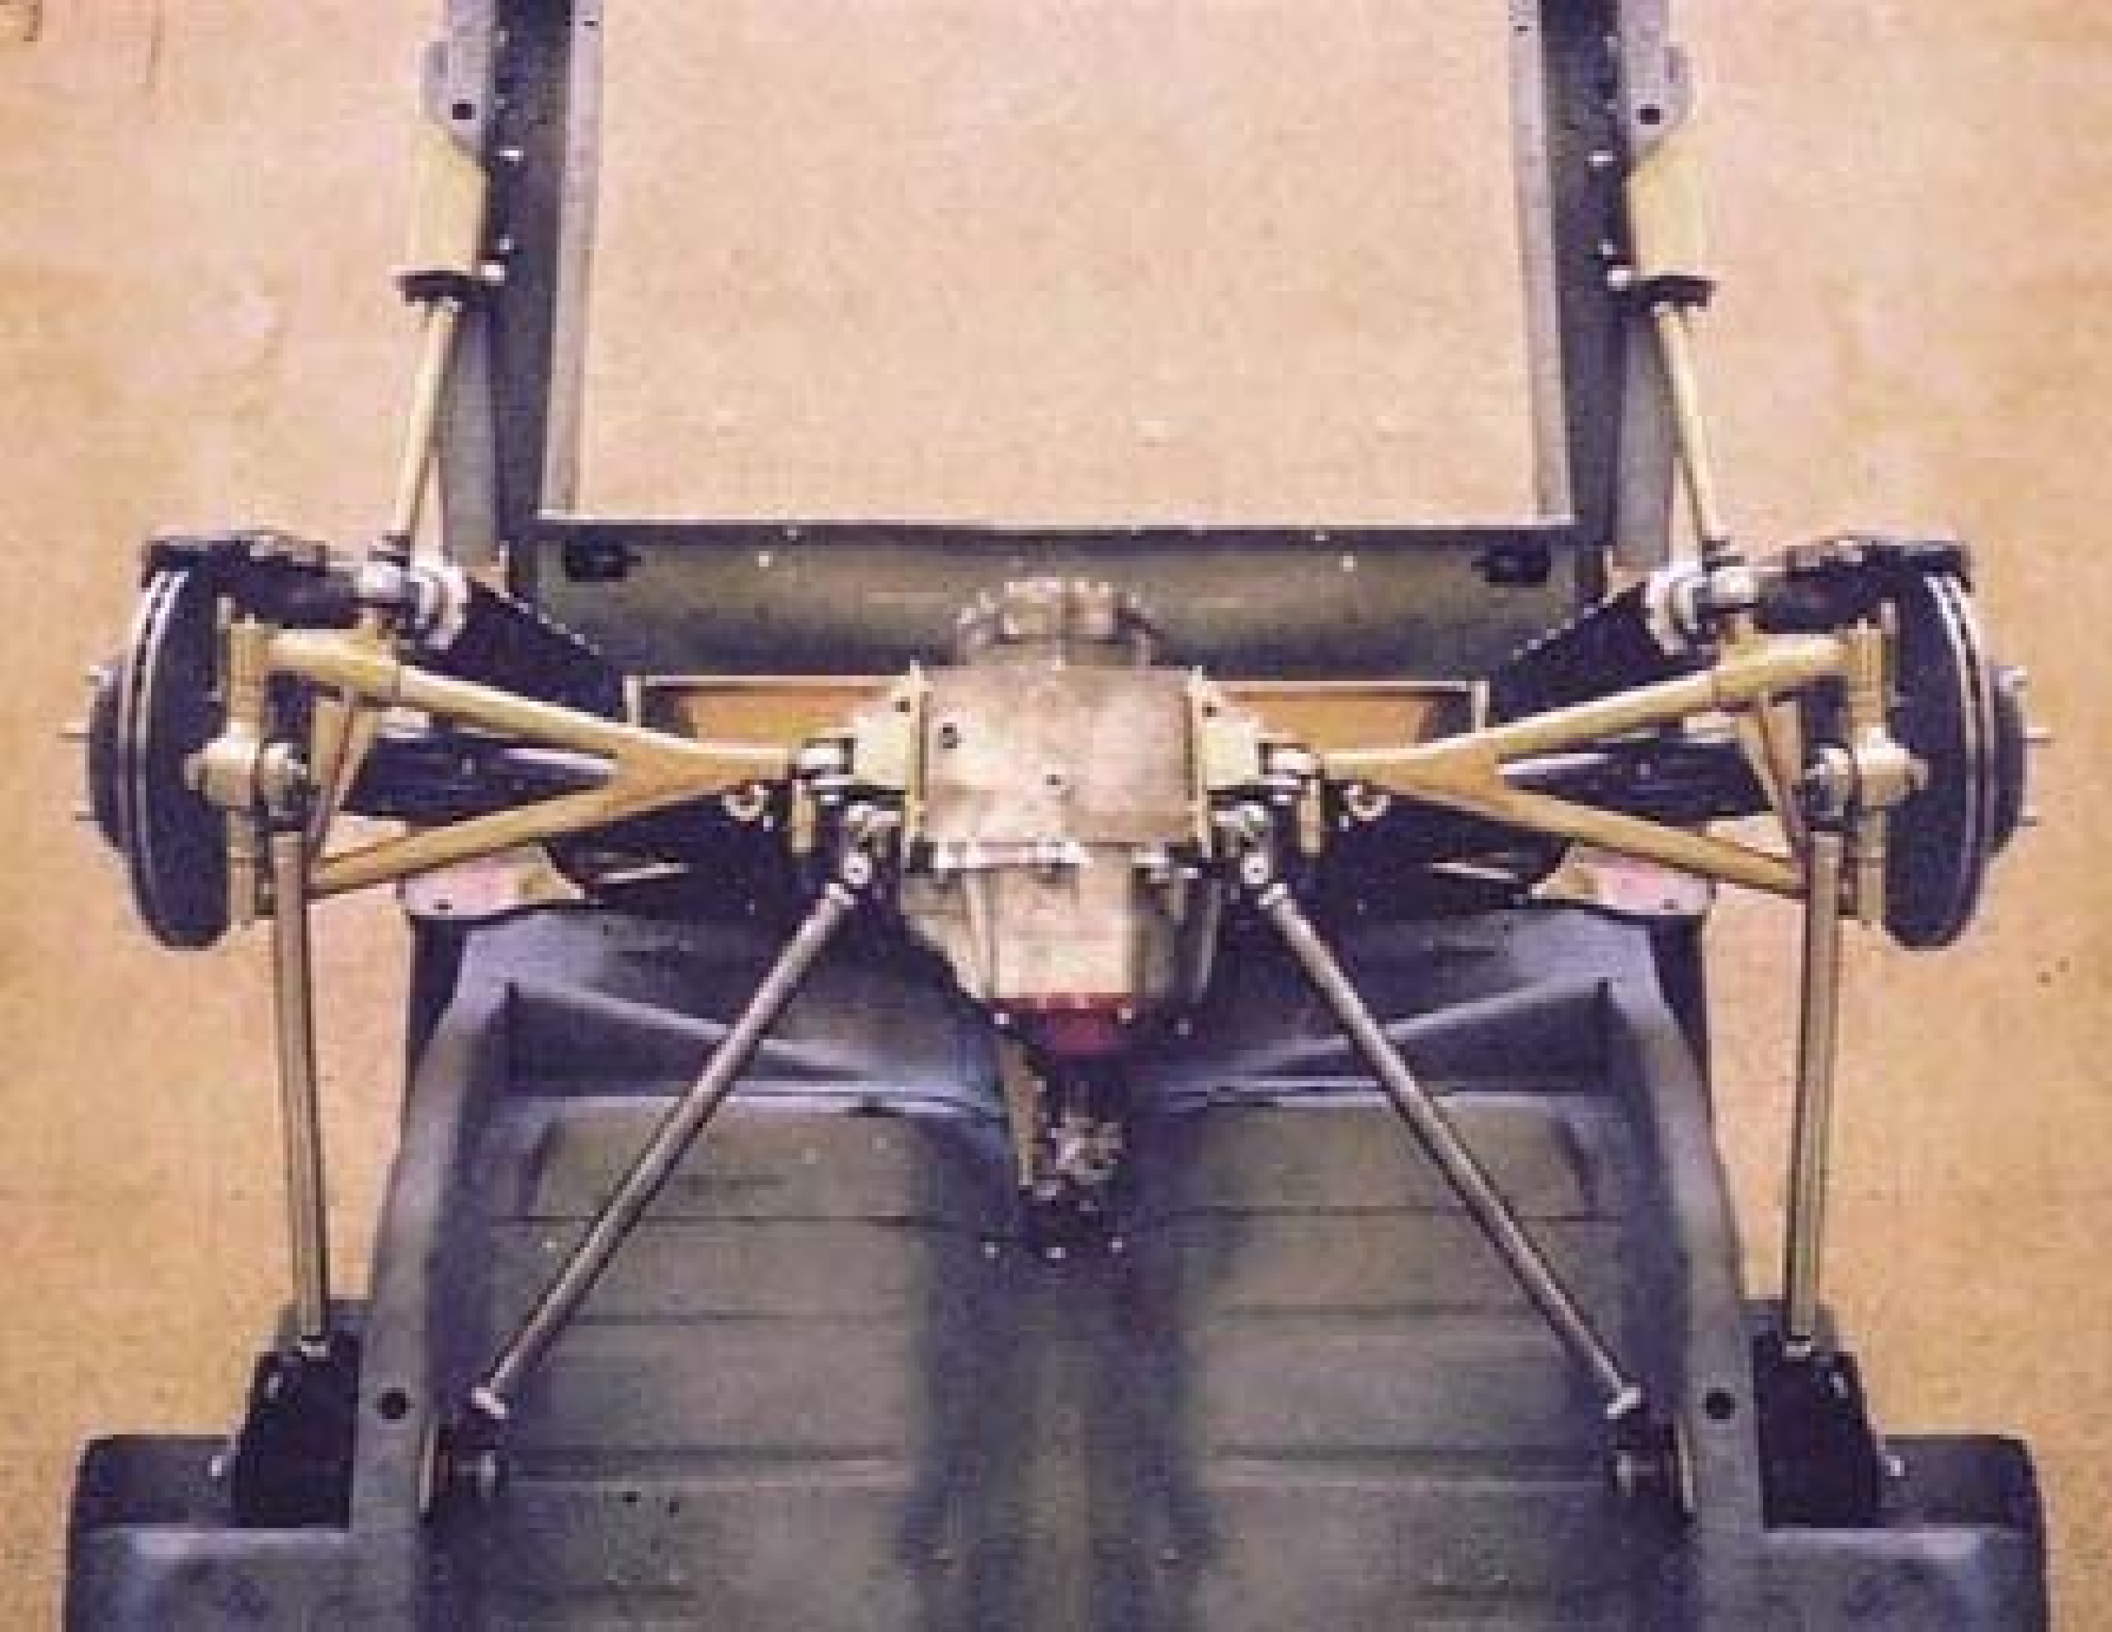

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHY NOT GO FOR 100% CAMBER RECOVERY?

I enjoyed this newsletter on a car dear to my heart [referring to the issue about Corvair rear suspension, April 2008]. *I had a question though. I am about to begin installation of a C5 Corvette rear end into my V8 car, using an adapter I designed to couple to a big car Saginaw 4 speed. I will need to center the engine in the car, and then design/build all new rear suspension. The question is how much camber gain to design in. You mentioned 50% as a typical gain. Why would you not design for more, even 100%?*

I am planning to use two trailing links, and SLA with the lower arm horizontal, and the upper arm angling down to the transaxle. The Corvette uses a beefy CV joint axle, which I will adapt to the corvair drive spindle.

Interestingly, if you remove the tailhousing from the Saginaw, the output shaft slides into and engages the pinion gear in the Corvette axle, with about 1" gap for the adapter plate. The splines match. Thank you GM! The ring gear is about 8.75" in diameter, and comes in ratios of 2.73, 3.15. 3,45. 3,73 and more. To line up the rear wheels with the axle will require the wheels to be moved back about 1.5", slightly less if the drive axles can angle forward a little.

Camber gain usually refers to the change of camber in degrees, per inch of wheel travel. Its units then are degrees per inch. It is commonly measured with the car stationary in the shop, the springs removed and the anti-roll bars disconnected. When you see a single number for this, it's most commonly the camber change in the first inch of compression from static. It increases as the suspension compresses and decreases as the suspension extends, if the upper arm is shorter than the lower. Conventionally, change toward negative in compression and toward positive in extension is taken as positive camber gain.

Camber recovery is the percentage of the roll angle that is recovered at the wheel when the car rolls. If the wheel leans the same amount as the body, that's zero camber recovery. If the wheel leans half as much as the body, that's 50% camber recovery. If it doesn't lean at all, that's 100% recovery.

For a car with a 57.3” track width (57.3 being approximately 180/π, the number of degrees in a radian), 1 deg/in of camber gain gives 50% camber recovery. To get 100% camber recovery, we have to have 2 deg/in of camber gain. The problem with that much camber gain is that the camber changes too much in situations other than roll: in longitudinal acceleration, over humps and dips, over bumps, and with fuel burn-off and other load variations.

With a passive independent suspension, it is impossible to have zero camber change in both ride and roll. The best we can do is to strike a compromise between the two, so that we don’t have any big camber changes in any situation. With a 57.3” track, one degree of roll is one inch of displacement difference at the wheels, or ½” per wheel. Two degrees of roll is twice that: 2” per wheel pair or 1” per wheel. So if the camber gain is 1 deg/in and camber recovery is 50%, we have one degree of camber change per inch of wheel movement in both ride and roll. The only way to get more camber recovery without more camber change in ride is to increase the track. This also reduces lateral load transfer and thereby increases cornering power, but unfortunately it makes the car bulkier.

The formula for instantaneous rate of camber gain is:

$$d\varphi/dz = \tan^{-1}(1/L_{vsa})$$

where:

φ = camber angle, degrees
z = suspension displacement at the contact patch center, inches
$d\varphi/dz$ = first derivative of camber with respect to suspension displacement (instantaneous camber gain), deg/in
L_{vsa} = length of virtual swing arm

This formula also works with units of length other than inches, provided you use the same units for z and L_{vsa}. Note, however, that 1deg/mm of camber gain is nowhere near the same as 1deg/in.

A simplified version is:

$$d\varphi/dz = 57.3/L_{vsa}$$

One question that sometimes arises is whether to measure L_{vsa} horizontally, or along the force line. Unless the instant center is at ground level, the true or vector length will be greater than the horizontal distance to the instant center. The radius on which the contact patch moves is actually the true or vector length, right?

Well, yes. But also, if the instant center is anywhere other than ground level, an inch of contact patch movement about the instant center is not a vertical inch; it's some lesser amount. And it turns out that the ratio of vector to vertical motion at the contact patch is the same as the ratio of horizontal to vector swing arm length. This ratio is also the cosine of the force line's angle of elevation.

So if you use vector distance for the swing arm length, and apply the formulas, you get the rate of camber change per vector inch of wheel movement, and if you use the horizontal distance, you get the rate of camber change per vertical inch. Convert a vertical inch to vector inches, and use the vector swing arm length to find the camber change rate per that unit of distance, and you get the same number as if you hadn't converted, and used the horizontal distance.

Formula for instantaneous rate of camber recovery is:

$R_\varphi = (1 - d\varphi/d\theta) * 100\% = (t/(2L_{vsa})) * 100\%$

where:

R_φ = camber recovery, in percent
φ = camber angle, degrees (measured with respect to the road, not the car)
θ = roll angle of sprung mass, degrees
$d\varphi/d\theta$ = first derivative of camber with respect to roll, i.e. the instantaneous rate at which the tire leans per degree of body roll
t = track width
L_{vsa} = virtual swing arm length, taken horizontally

GYROS FOR ROLL STABILITY

Has anyone ever used a vertical axis flywheel to damp pitch and roll through gyroscopic effects? Ignoring practical considerations and assuming a low drag (vacuum), high RPM, low weight, probably counterrotating disc layout, how beneficial would this be to the ride quality of an off road (desert racing) vehicle?

Would it be worth adding a system to slow either disc independently, via brakes and one way clutches, in order to gain active yaw control? Could such a system ever completely replace the normal steering system (ignoring reliability concerns)?

In your opinion, how would the gyro damping compare to interconnected suspension systems? Has anyone tried fully active suspension systems in a desert racing environment (thinking of the bose demonstration videos here)?

I will confess to some lack of expertise in desert racing, as I have had absolutely no clients or even acquaintances doing it, but to my knowledge nobody has used fully active suspension off-road. I would think it would eat a huge amount of power. Active roll control, and maybe also pitch control, employed to allow soft passive suspension to absorb bumps, might be more efficient. Power consumption was a big part of what killed Lotus's fully active system in the '70's. It ate about as much power as it took to drive the car in cruise mode on the highway, so it had nasty effects on fuel consumption. Fuel thirst might be more tolerable in a racing environment, but in off-road applications the energy consumption would be higher than in road use, due to the much greater movement of the suspension.

For road racing, active suspension looks more appealing, but it's illegal in all existing classes, for reasons of cost containment and car parity. Probably if it started to become a threat in off-road racing, the same thing would happen there, for the same reasons.

The gyro idea is interesting. I had to scratch my head awhile, and even play with a bicycle wheel, but I concluded that if you had two flywheels lying one on top of the other, spinning at equal and opposite velocities, you really would still have stabilization, and the precession effects in both perpendicular axes would cancel out.

I decided you could call this a gyro sandwich. (Sorry. I couldn't resist.)

This would make the sprung mass resist roll and pitch accelerations, and it might therefore allow use of softer suspension. If the suspension could be softened, it might be possible to have some benefits similar to interconnected suspension. The idea could also be used with interconnected suspension.

Have you ever seen a Segway? It has two wheels, one on each side. It uses an electric gyro to keep it from falling over in pitch. There was also a vehicle called the GyroCar about forty or fifty years ago that used a vertical flywheel. It was basically a full-bodied motorcycle, with a door to get in and out, and no holes to put your feet down. It had two small, retractable wheels like training wheels on a bicycle, which deployed for parking. With the flywheel running, it would hold itself upright while sitting still, without the parking wheels.

As for speeding up and slowing down the flywheels differently to create yaw moments and steer the vehicle, yes you could create yaw moments that way, but as soon as you did that the flywheel speeds would no longer be equal, and you'd start getting unintended precession effects. I think you'd probably want conventional steering.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

FRONT TO REAR TRACK RATIO

I've been building a wide-body Porsche 914 for a few years now, and had always assumed that the wider rear track would stabilize the rear of the car better (especially since I now have a somewhat heavy sprint car engine in the back). But lately I've been questioning whether that's correct. I found that F3 chassis currently in use run a wider front track. And on a trip to Laguna Seca Raceway I noticed a Ferrari F430 which appeared to have a wider front track as well. I even went to speak with some local FSAE students to get some answers. But the only thing they said is, "A narrower rear track gets around the cones faster." That was not the answer I was looking for.

Published lateral acceleration numbers from many of the current 3-wheeled chassis are impressive. But I cannot find any material explaining why.

I am in desperate need of someone to explain the chassis dynamics involved.

First, let's make sure we define what we mean by track width. This is the transverse distance between the center planes of the wheels. I sometimes encounter people who think it means the overall width at the outside of the tires instead. Engineers generally use the term to mean the center-to-center dimension.

Exactly what is the center can have some nuances, if we are trying to really be accurate. When the wheels have some camber, we may mean the distance between the contact patch centers, or more properly, load centroids. This is probably the most realistic and accurate, but those points are a bit hard to ascertain in the design phase. The most convenient approximation to this is the intersection of the wheel center plane and the ground, which is not quite the same point as the load centroid, but is easier to define in the design phase. Or, we can say it's the distance between the wheel midpoints, at hub height. This is roughly equivalent to taking the other two values at zero camber, and treating that as design track width. If we are just after general information, these subtleties aren't necessarily important.

There is no perfect ratio of front to rear track width, but there are practical considerations that may lead us one way or another. In the case of your V8-powered Porsche 914, you are probably keeping the suspension as nearly stock as practicable. The car will be tail-heavy, so it needs more tire in the rear. The only way you can add a lot of tire and rim width without major chassis fabrication and suspension redesign is to add it on the outboard side of the wheels, and flare the fenders. At the front, you don't need to add as much tire, and you want to keep the scrub radius from getting too big, so the front track won't end up growing as much.

The car will be okay this way, although I think you'll need to make a conscious effort to apex wide on tight turns to avoid clouting curbs with your inside rear. I have a street car that's wider at the back myself, and I have to constantly remind myself that the rear is wider than the part in my field of view. The car handles well, but this aspect requires some driver adaptation.

As to whether the car gets tighter or looser if you narrow or widen one end relative to the other, in most cases it will change, for various reasons. However, the factors involved vary from one case to the next, and we can't reduce the whole matter to a blanket generalization, even regarding the general direction of the behavior change.

If we are designing a race car from scratch, we aren't constrained by a need to preserve existing structure. For road course or oval track applications, we generally will be constrained by limits on tire and/or rim width, perhaps even a spec tire, and a limit on overall width. For wide, high-speed tracks, usually the widest car will be fastest, because there will be less lateral load transfer and we will be working the four tires more nearly evenly. In most rear-engined cars, we will have wider tires at the rear than at the front. If we have the same overall width limit at both ends, and the rear tires are wider, we need to have the track narrower at the rear than at the front, simply to take full advantage of the width limit at both ends.

In Formula SAE, and for SCCA solo events and other autocross events, the constraints are somewhat different. The turns are tight, the track is narrow, and some of the turns are defined only by the need to clear cones at the inside – slaloms are an obvious example. There is no maximum width limit in FSAE. The event itself prevents us from wanting a wide car. In tight turns, the rear tires track inside the front ones, even when there are significant slip angles. Not surprisingly, the car has a tendency to collect cones with the inside rear. It is consequently common to make the car about 4 inches narrower across the rear than across the front. The rules actually impose a limit on the track ratio: the narrower track can't be less than 75% of the wider. Most teams don't get anywhere near violating that rule.

There are certain cases in which we can readily predict whether adding rear track width will add understeer or oversteer. If we have a live axle, and we leave the springs alone and add track width by moving the tires out, the car should get tighter (more understeer). The car will have the same amount of roll, the same angular roll resistance at both ends, the same amount of load transfer at the front tires, and the same amount of load transfer at the rear springs, but less load transfer at the rear tires, because the rear roll-resisting moment will be taken out over a wider base. The total load transfer will be less, and the reduction will all come at the rear. That should make the rear stick better compared to the front. Of course, the car can always be re-balanced using other changes.

Another case we might consider would be to increase the track width at one end, while keeping the wheel rates in roll unchanged.

I probably should spell out exactly what I mean by the wheel rate in roll, as I have encountered some confusion on this from various quarters. Wheel rate in roll, as I use the expression, is the rate of elastic change in wheel load with respect to linear suspension displacement, when the two wheels of a front or rear pair each move the same amount in opposite directions – in English units, the pounds of load change per wheel when one side compresses one inch and the other side extends one inch. This relates to the angular roll resistance as follows:

$$K_\varphi = ½ * K_{roll} * t^2 * \pi/180$$
$$K_{roll} = (2 * K_\varphi) / (t^2 * \pi/180) = 360K_\varphi / \pi t^2$$

Or, approximating 180/π to three significant figures:

$$K_\varphi = ½ * K_{roll} * t^2 / 57.3$$
$$K_{roll} = 2 * 57.3 * K_\varphi / t^2$$

Where:

K_φ = angular roll resistance, lb-in/deg
K_{roll} = linear wheel rate in the roll mode, lb/in
t = track width, inches

In the case mentioned above, where we leave the springs alone on a beam axle and move the wheels out, we are increasing t, leaving K_φ unchanged and decreasing K_{roll}. Or at least K_φ remains unchanged if we are ignoring effects of tire compliance.

Now, let's take a simple case where all roll resistance is elastic, the c.g. is in the middle of the car, and K_φ, K_{roll}, and t are all equal at both ends. Then suppose we increase t at the rear by a factor of two, leaving the front end unchanged, and leaving K_{roll} unchanged at all wheels. This is a ridiculously exaggerated example of course, but the numbers are easy to work with in your head.

At the rear, K_φ increases by a factor of four. The total of K_φ at the front and K_φ at the rear increases by a factor of 5/2, or 2.5. Roll at a given lateral acceleration decreases by a factor of 2.5. Total load transfer at a given lateral acceleration is 2/3 as great, because the mean track width is 3/2 as great.

The rear suspension is now providing 80% of the roll resistance instead of 50%, but doing it with tires spaced twice as far apart. The load transfer at the front is (2/3)*(20/50) = 26.67% of what it was before we widened the rear track, and the load transfer at the rear is (2/3)*(80/50) = 106.67% of what it was.

As a check, (106.67% * 50%) + (26.67% * 50%) = 66.67%. That is, the sum of the new rear load transfer and the new front load transfer is 2/3 of the old total load transfer, as it should be with a 3/2 wider mean track.

Note that in this case we actually increased load transfer at the rear a bit when we widened the track, and dramatically decreased load transfer at the front. That should add oversteer rather than understeer.

This example is particularly relevant to go-karts, where the suspension mainly consists of the tires, and consequently wheel rates in all modes tend to remain constant as we change other things, provided we use the same tires and inflation pressures. Kart racers tend to find that the kart gets freer (understeers less/oversteers more) when the front track is narrower compared to the rear, although it depends to some extent on the turn in question and on how we widen or narrow the front track. If we widen the front track by spacing the wheels out on the spindles, we increase caster jacking with steer. This can mean that the vehicle is tighter in sweepers yet freer in tight turns.

So far, we have been analyzing effects of track width, and track width relationships, purely in terms of effects on load transfer. There are some other effects as well. If we have a locked rear, and we are road racing, meaning we can't use rear tire stagger, we generally have a greater tendency toward locked-axle push when we widen the rear track. With a limited-slip diff, there is a similar effect, only more subdued. However, if rear load transfer increases, that can make it easier to unload the inside rear tire in hard cornering, reducing locked-axle push instead.

It should now be apparent why I say the whole business can't be reduced to a blanket generalization. However, with sufficient thought it is possible to understand the various factors at play.

Regarding 3-wheeled vehicles, these can produce similar lateral acceleration numbers to 4-wheelers, if they don't tip over. They can be made very light, so with suitable tires they can generate plenty of lateral g's.

The key to giving a trike good overturning resistance is to put the preponderance of the weight at the 2-wheeled end. Do that, and keep the c.g. low and the track wide, and you can compete with cars on the skidpad. Put the heavy bits at the one-wheeled end, and you are creating job security in the legal profession.

In pure lateral acceleration, all lateral load transfer occurs at the 2-wheeled end. This will generally lead us to want larger tires at the 2-wheeled end, even though there are twice as many of them. The outside one will be doing most of the work.

The July-August 2005 issue of the Newsletter is entirely about trikes.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DIRTY TRICKS

I have a request...we are trying to get an asphalt Late Model working with some ideas normally reserved for dirt suspension. I see some of these techniques used more and more. I am thinking about:

- *LR spring in front of rear axle*
- *Double springs within single coilover*
- *Double link on top link for rearend.*

We play with these but don't necessarily understand all the concepts.

All of the features mentioned are used on dirt cars for one purpose: to add wedge or diagonal percentage (total of left rear and right front tire loads as a percentage of the total for all four wheels) when power is applied. Dual springs can be used for other purposes as well.

The reason we want to add diagonal percentage under power is to tighten the car (add understeer), so that the rear tires' grip is not entirely used up for cornering, and we have more available for propulsion. That lets us apply more power and get a faster exit. We want this understeering influence to come in as we add power, so the car won't understeer excessively the rest of the time.

We can have too much of this effect. More isn't necessarily better. If we overdo it, the car can have a power push. The idea is to tighten the car enough for good exit, but not tighten it excessively.

Having the left spring ahead of the axle and/or the right spring behind adds wedge under power if and only if the springs mount to the axle, not the birdcages, and if the torque reacting mechanism (pull bar/top link or lift bar/torque arm) is compliant and lets the axle rotate backward under power. The axle tries to do this because the pinion gear wants to climb the ring gear and lift the nose of the third member. Or to state it another way, in a live axle suspension, drive torque reacts through the suspension linkage.

With the left rear spring ahead of the axle and the right rear spring behind, the left rear corner of the car is lifted by axle rotation and the right rear corner is let down. That rolls the car to the right and adds diagonal percentage. Again, this only happens if the linkage controlling the axle lets it rotate.

Indeed, all strategies for adding wedge under power that involve the rear suspension work by rolling the car to the right. All strategies for adding wedge under power that involve the front suspension work by resisting roll to the right, or trying to roll the car to the left.

In a rear-drive car, we don't have any torque or any thrust at the front wheels under power, so we have less to work with. However, we do have the tendency of the front end to lift as load transfers to the rear. If the left front suspension resists lifting more strongly, or the right front resists lifting less strongly, that adds wedge under power. If the right front spring is softer, or the left front is stiffer, that encourages the front end to rise more on the right side under power than on the left side, and adds wedge under power. If that sounds backwards, remember that spring rate is the rate at which force changes with displacement, not the magnitude of the force. A stiff spring is more reluctant to extend, or creates more load change per inch of extension, than a soft one.

Dual springs are used on the left front of dirt Late Models to make the springing on that corner of the car get stiffer when the front end rises. Contrary to what you might suppose, the dual springs are used to give a falling rate rather than a rising rate. The rate increases when the suspension extends to a certain point, rather than increasing when the suspension compresses to a certain point.

This is done by stacking two springs on top of each other on the coilover, with a slider between them that can slide up and down on the coilover. The slider can slide upward as far as it wants, but there is an adjustable collar that stops it at some point when it slides down. When the slider can slide freely, the springs are both working, in series. Once the slider tops out against the collar, only the lower spring can extend further.

Two springs acting in series are softer than either of them acting alone. More precisely, the rate of the combination is equal to the product of the individual rates divided by their sum. For example, a 100 lb/in spring stacked on top of a 200 lb/in spring gives a combination rate of 20,000/300 = 67 lb/in. If either of the two springs is immobilized, we are left with the rate of the remaining one.

Applied to the left front suspension, the idea is to have the upper spring working until the driver gets on the power, and then top out.

The driver may or may not be able to actually feel the top-out point. If the driver isn't sure where the spring is topping out, or even whether it's topping out, there are ways you can still find out.

Using electronic data acquisition, you simply find the suspension displacement where the upper spring tops out, and then look at where the car is on the track when the suspension reaches that point, and what the driver is doing with the controls and what the driver says about the car's behavior.

If you are testing without data acquisition, you can cut a hole in the top of the fender and attach a rod or dowel to the spindle or some other part of the suspension, extending upward through the fender, as a travel indicator. Mark the travel indicator with bright-colored tape or paint below the point where the upper spring tops out. When the color disappears, the upper spring is topped out. The driver can watch this indicator, and so can trackside observers. You can also use video to record what it does.

The usual way of using compliant upper links on the rear axle is to have both of them near the middle of the car, and have one act in compression, for braking, and the other in tension, for power. If you use two rigid links, or even one rigid link and one compliant one, they fight each other and you have a bind.

The compliance should be damped somehow. One common strategy is to use a shock absorber (damper), with a bump rubber on its shaft, as the compression link. The damper acts as a compression link when it bottoms out on the rubber, and it damps axle rotation in both directions. On dirt cars, the shock is normally angled up at the front, and this creates some pro-lift when the shock sees compression. The pro-lift helps prevent wheel hop, which can otherwise be a problem as dirt cars often use a lot of rear brake.

Engine braking torque always acts through the shock. Torque from the brakes themselves may or may not. If the calipers are mounted to the axle tubes, brake torque acts through the shock. If the calipers are mounted to the birdcages, the torque from the brakes acts through the birdcage linkages instead.

You may not want pro-lift on a pavement car. It depends partly on how much rear brake your driver likes.

I like mounting at least the tension link offset to the left, behind the driver. Rules permitting, I also like the idea of using a two tension links, one to the left of the other, and choosing spring rates and preloads so that light drive torque acts mainly through the left link, and the right or center link takes up and absorbs more of the load as torque increases and the axle rotates more. This allows the rear suspension to add wedge most when the car needs it most – on a slick surface – while not adding excessive amounts when grip is good.

Pavement Late Model rules in the US are a real hodge-podge. NASCAR has three different Late Model classes, the differences being mainly in the engine rules. They prohibit most of the suspension tricks found in dirt cars. There are several regional sanctioning bodies, and I think all of them are more permissive than NASCAR. Then there are tracks that have their own unique rules. Consequently, the legality of the ideas described above will vary drastically depending on where you're running, and what class.

PORSCHE 993 REAR SUSPENSION

In your article on the Corvair rear suspension, you stated you have a V8 Crown Corvair. I'm building a V8 Corvair using a Porsche 915 5 speed and a complete suspension from a C2 993 911. They're complete sub-assemblies that unbolt from the body like the Corvair front end. I'm wondering what your opinion is of the Porsche 993 suspension and what sort of things I should be on the lookout for when I get to installing these sub-assemblies.

For an illustration of this suspension, see http://www.autozine.org/technical_school/suspension/tech_pic_sus_dw.jpg. This is a rear three-quarter view of a right rear assembly, removed from the car. Also see http://www.autometricsmotorsports.com/images/cars/rsr/933_strut_rear.jpg. This is a rear three-quarter view of a left rear assembly, in the car. The September 2008 *Road & Track* has a picture on p.77.

Another correspondent suggested to me that maybe Porsche vindicated GM by imitating its Corvair design for the 993. However, it will be seen from the pictures that the Porsche design bears more resemblance to current Corvette rear suspension. It is a true 5-link system, with the driveshaft not serving as a suspension link. It does have a bit in common with the Corvair and C2/C3 suspension, in that all of these are short-and-long-arm suspensions, and it does resemble the Corvair layout more than the 911 semi-trailing arm system did.

But if the 993 is like another car, that would be the front-engined Porsche 928. That was the first Porsche to have 5-link rear suspension. At the time, Porsche called this design the Weissach Axle. As used in the 928, it had compliance characteristics intended to toe the rear wheels in on trailing throttle, minimizing drop-throttle oversteer.

I have not worked with 993's, and I can't tell what the system's side-view geometry is like. It looks like Porsche's thinking on front-view geometry pretty much agrees with mine: camber recovery around 50%, roll center height somewhere in the 3 to 4 inch range. Road tests I've read don't contain comments that would point to any severe geometry problems.

So, absent better information, I would advise sticking close to the geometry Porsche uses, and mounting the structure to try to preserve that, while also trying to maintain solid structure and decent load paths.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MONOSHOCK SUSPENSIONS

I have seen photos of some of the alternate springing methods used on race cars these days including monoshocks and 3rd springs, but I don't understand how they work or what their advantages might be over coilovers in the usual methods of actuation. Could you explain how those work?

Z-BARS

I have a question on the fundamental principle of the anti-roll bar.

Talking about any standard U shaped ARB, when the vehicle rolls the outside wheel is loaded and the inside wheel wants to be unloaded but faces the resistance of the spring or torsion bar (depending on the suspension construction), and therefore the inner wheel is unloaded by an amount which is the difference of the weight transfer and the spring loading.

However, given the situation where you need a soft setup (ARB disconnected) would it be useful to use an S- shaped ARB instead? I have tried to find literature on this but found nothing. I would assume that on a vehicle with a relatively low CoG like a Formula Ford, trying to load the inner wheel to an extent could yield some additional grip? Of course a new jacking force would be created, but do you have any comments on this?

STOCK CAR FRONT ENDS WITH BUMP RUBBERS

Now that the NASCAR COT is allowing bump rubbers (and will move away from spring coilbind setups), there is lots of interest in the use of bump rubbers in a number of stock car series.

Do you have any good advice to give on this subject?

I have grouped these three questions together because there is a common thread through all of them. They all relate to ways we can get a wheel rate that rises as the suspension compresses.

To understand some important differences between ways of doing this, it is helpful to recognize that a four-wheel suspension system has four modes of movement – roll, pitch, heave, and warp – and the elastic resistance to wheel displacement, or rate of elastic loading change with respect to displacement, can be different in all four modes.

However, in most suspensions, the wheel rates at a given wheel are very nearly identical in pitch and heave, and in roll and warp. Also, in road racing cars, the right and left wheel rates are identical, or very nearly so. Consequently, for simple and conventional cases, we can break the suspension system down into front and rear systems, each having a wheel rate for ride, and another for roll.

In most suspensions, neither of these rates is truly a constant. When the system has to contend with large variations in load, it becomes desirable to have the suspension get stiffer in ride as it compresses. Before rising-rate suspension became popular in racing, it was widely used in road vehicles, particularly in the rear suspension on trucks and front-wheel-drive cars.

If we have a conventional suspension with two individual springs acting at the wheels, two dampers acting on the two wheels separately, and an anti-roll bar, and we arrange for the springs to have a rising rate, then the rate in both ride and roll rises as the suspension compresses in ride. This also applies if we provide packers or bump rubbers on the shocks or elsewhere.

One effect that this will have on car behavior is that when we set up the aero balance to add more downforce at one end than the other, to make the end with more downforce stick better at high air speeds, the change in roll resistance distribution will change the handling balance oppositely to the change in downforce. This makes the net effect somewhat more difficult to predict, and it may in some cases make the car unresponsive to the adjustment.

To get away from this, it is desirable to have a way to add wheel rate in ride that does not add wheel rate in roll. Third springs, Z-bars, and monoshock systems are various ways of doing this.

For those not familiar with all of these, a third spring is most commonly seen in pushrod or pullrod suspensions, with rockers and a T-bar for anti-roll. The T-bar is a torsion bar with the torsional part forming the leg of the T, and push-pull rods extending from the rockers to the ends of the cross of the T. The T-bar can rock backward and forward freely about its lower end, so when the rods move synchronously, there is no resistance. When the rods move oppositionally, the bar is forced to twist, and resistance results.

The "third spring" attaches to the middle of the top of the T, where the leg meets the cross. It is most often a coilover, but it can be only a spring on a slider or dummy shock, or a shock or slider with a rubber snubber, or a coilover with a snubber. Ordinarily, the third spring is designed to have a pronounced rising rate. The third spring acts only when its attachment point to the T-bar moves, which only happens in ride. The total arrangement thus provides a steeply rising rate in ride, without a rising rate in roll. Depending on the geometry used, it is theoretically possible to have rising or falling wheel rate in either mode. The nice thing is that we control the two independently of each other.

A monoshock suspension generally has a single rocker, with a single coilover, actuated by pushrods from both right and left wheels. The coilover acts only in ride, much like a third spring. Roll is usually accommodated by having a laterally sliding shuttle in the rocker. The pushrods anchor to the shuttle bar. When the pushrods move synchronously, they rock the rocker. When the pushrods move oppositionally, they slide the shuttle bar. Roll resistance is most commonly provided by stacks of Belleville washers on the shuttle bar.

This arrangement offers us independent control of roll and ride properties, much like a third spring, with reduced weight and cost. Generally, there is no damping in the roll mode, which is a disadvantage. Monoshock systems are generally used only at the front, in smaller-displacement formula and sports-racing cars.

A Z-bar is what the second questioner describes: a transverse torsion bar, similar to an anti-roll bar, but with one arm pointing back and one forward, so that it has the shape of a Z rather than a U. It affords a way to spring only the ride mode, in a suspension that doesn't have pushrods and rockers.

We should note that there is no such thing as a negative-rate spring. In certain unusual cases, it is possible to mount a spring so that the force generated by the assembly diminishes with increasing displacement, but there is no way to make a spring do that by itself. A Z-bar, at least as we normally encounter it, cannot, in itself, reduce roll resistance or promote roll. It just adds rate in ride without affecting roll. This means that we can use it to get reduced roll resistance with a given ride rate, if we use it in conjunction with soft individual wheel springs, or with no individual wheel springs at all.

We most commonly see this with swing axle rear suspensions. By having a high wheel rate in ride and little or no wheel rate in roll, we can minimize the limit oversteer and upward jacking that a swing axle rear end tends to generate. Even when the Z-bar is the only spring in the system, the suspension still resists roll and creates load transfer. It just does this entirely through its geometry, rather than elastically.

If we were to use Z-bars instead of anti-roll bars at both ends of a modern formula car with low roll centers, we would increase roll, but we would not increase the loading of both inside wheels. Increasing loading of the inside wheels – i.e. equalizing the loading of inside and outside wheels, reducing load transfer – is beneficial in all conditions, wet or dry, but there is no way we can do that for both wheel pairs by making changes to the suspension. The suspension only controls how much of the load transfer occurs at the front, compared to the back. The only way to reduce the total load transfer at both ends combined is to lower the c.g. or widen the track.

With stock car front ends, we have a somewhat similar situation, but with some differences. Generally, in oval track racing a given track's turns are all taken at similar speed, but the car has to pass a ride height check. With the stock car, we are trying to get the splitter or valance as close to the track as possible, and keep it there, without destroying it against the track surface, yet still have the car pass ground clearance check statically. So we run soft springs and a big anti-roll bar, and arrange for something to bottom when the car reaches the desired ride displacement. Until recently, at least in the upper NASCAR divisions, bump rubbers were prohibited, so people arranged for the springs to reach coil bind. Now bump rubbers are permitted, and teams are learning to work with them.

The bump rubbers are on the shock shafts. They unavoidably act in both ride and roll. They can be about as soft as we like, but we still face a tradeoff between controlling the splitter and having highly non-linear roll resistance at the front of the car. If we try to soften the rubbers to get some roll compliance, we give up some aerodynamic control.

On high-speed ovals, aero forces are so significant that reducing front roll resistance can actually cause a push, by costing us front downforce as more air gets under the nose. On short tracks, this is less of a concern, but still a consideration.

The problem with having a steeply increasing wheel rate in roll at just the front of the car is that it becomes difficult to get good handling balance through the entire turn, and as grip level varies. We can increase rear roll resistance easily enough, but giving it a non-linearity that harmonizes with the front is not possible.

Rules permitting, there would be a strong case for having a stiff Z-bar at the front of a stock car, in addition to the anti-roll bar, and having sliding drop links with bump rubbers for the Z-bar. Another possibility would be to have a third arm in the middle of the anti-roll bar, and a bump rubber acting on that. I'm sure NASCAR would not allow this, but if you are running with a different sanctioning body, or running to track rules, maybe you could get your officials to allow it.

I have heard of one instance of a Z-bar being allowed as an anti-roll bar. This was on the swing-axle rear end of a Triumph Spitfire, in SCCA. Aftermarket anti-roll bars were legal, and the officials ruled that the bar still qualified, even if one arm pointed the wrong way. To me, it's a bit of a stretch to call a device an anti-roll bar when it doesn't resist roll, but that was their ruling.

In many cases, the ruling would probably come down to whether the device is reasonably affordable, can be retrofitted to existing cars, and improves the show. Having a car that is more consistent is definitely desirable from the competitor's standpoint, but from the promoter's standpoint there is a strong case for deliberately making it hard to achieve consistency. When all the cars are inconsistent, some cars are better in some conditions and in some parts of a run, and others are better at other times. This makes for more lead changes and less predictable race outcomes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DIFFERENTIALS AND YAW CONTROL

I've been reading some statements on various Internet message boards (yeah, I know) regarding how a Quaife or Torsen type of differential adds some level of yaw control to a RWD vehicle above and beyond what a CLSD provides. Given that many CLSDs are adjustable for preload and lockup percentage, what additional (if any) yaw control will an automatic torque biasing differential provide?

Maybe a more direct question is: Can a passive differential of any type provide any yaw control beyond the inputs from a driver's right foot?

The quick answer is: no, a passive differential cannot provide yaw control, in the sense that we usually use the term nowadays. However, more loosely, the action of passive diffs does affect the car's yaw behavior, and the driver's ability to control the car's yaw behavior, so semantically maybe there's some sort of case for calling that yaw control. That's a stretch, though.

For readers less familiar with these devices, CLSD stands for clutch limited slip differential. CLSD designs commonly used in road racing have split differential carriers that can spread apart slightly when the carrier applies torque to the pinion shafts. This spreading is used to load packs of clutch plates on either side of the diff which create a locking effect. The amount of spreading force per unit of drive torque depends on the angle of flats on the pinion shaft, which bear against corresponding ramps on the carrier halves. There is one set of flats and ramps for forward torque, and another set for reverse torque. By selecting components with different ramp angles, the unit's properties can be varied for propulsion torque and deceleration torque independently. Additionally, the clutch packs can be preloaded to provide an initial locking torque which is always present, even when no torque is being applied by the engine.

A Torsen or Gleason-type differential uses worm gears rather than clutch discs to produce locking torque. Normally only one set of worm gears is available, so Torsens are less tunable than CLSD's. We can normally preload the gears to some extent, although the preload tends to be highly wear-sensitive. One interesting difference is that reverse torque in a Gleason will actually relieve the preload and reduce locking torque, at least up to some value, whereas with a CLSD, preload adds locking torque in all conditions. "Automatic torque-biasing differential" is another term for a worm-gear LSD.

Actually, all limited-slip diffs bias torque automatically. The torque to the output shafts is allowed to differ by an amount less than or equal to whatever locking torque the mechanism creates. Whenever the torque bias, or difference, is less than the locking torque, both shafts are forced to turn at the same speed. Any time the shafts are turning at different speeds, their torque differs by the locking torque.

By contrast, with an open differential, the torque to both shafts is always identical, and their speeds can differ freely. It is also possible to design open diffs that split torque unequally, in a fixed ratio. These designs are commonly used for center diffs in all-wheel-drive systems, sometimes in combination with an active or passive locking mechanism. Open diffs in a front or rear axle are always 50/50.

Yaw is rotation of the car about a vertical axis: change in the direction the car is pointing. Strictly, then, yaw control would be anything that controls such motion – the primary mechanism for this being the car's steering. However, in current usage, yaw control means active, computer-controlled creation of yaw moments, intended to augment the driver's control of the vehicle over and above that provided by the steering. Usually, this is done by selectively applying one or more of the brakes.

The propulsion or retardation thrusts from the tires create yaw moments. If we have two equal propulsion thrusts from the rear wheels, and the car's c.g. is exactly centered right to left, the two yaw moments exactly cancel each other, and there is no net yaw moment from propulsion. If the c.g. is not exactly centered, the rear tire thrust forces have unequal moment arms about the rearward inertia force acting at the c.g., and there is a yaw moment toward the heavy side. However, when the c.g. offset is small, this yaw moment is likewise small.

When there are unequal thrusts from the rear tires, and the c.g. is centered laterally or nearly so, we get a yaw moment away from the greater thrust. For example, if there is more thrust on the right, the car tends to turn left.

We noted earlier that a limited-slip diff allows the rear wheel torques to differ by the amount of the locking torque. So when the rear wheels have unequal grip, there is a yaw moment away from the side with more grip. The greater the locking torque, the greater this yaw moment can be.

If the road surface has dramatically and erratically varying grip levels – for example, randomly distributed patches of snow, ice, and bare pavement – and the diff has a lot of locking torque, under power the car will tend to erratically snake right and left. Absent any computerized yaw control, the driver will have to make constant corrections with the steering to keep the car pointed in the desired direction. With an open diff, the car will have better directional stability. Unfortunately, it will also have less available propulsive thrust.

When the rear tires have unequal amounts of grip, we face an inescapable tradeoff: thrust versus stability. No device or control strategy can give us both at once. Either we let the thrust of the drive wheels be unequal, to let the wheel with greater grip take advantage of its grip, and accept the resulting yaw moment, or we restrict that thrust to make it more equal to the lesser thrust, and gain yaw stability at a price in propulsion.

Even if we have a computer controlling the brakes and/or diff, the requirements for traction control and yaw control inescapably conflict. To achieve traction control, we need to brake the wheel with less grip. To achieve yaw control, we need to brake the wheel with greater grip. If we do both at once, we are merely turning fuel into heat and brake pad dust. The one thing we can do with computer control that we cannot do otherwise is to set some limits on the car's yaw behavior, below which the system functions to optimize thrust, and above which it changes character and optimizes stability instead.

The differential with the most driver-friendly properties, the one that makes it easiest for the driver to control the car's yaw behavior, is a completely open diff. Comparing limited-slips, the one that comes closest to this will be the one that has the smallest locking torque, under the conditions being examined. In general, worm gear diffs have less locking torque than clutch pack diffs, although this depends on the specifics of the units being compared. If the worm gear diff locks less, a car equipped with it will be more stable in yaw under power, and will therefore feel more like a car with computerized yaw control. However, the same car with more locking torque will be able to put more power to the road and will therefore be faster, provided the driver can keep it pointed straight.

So it is defensible to say that, in general, worm gear diffs provide a measure of yaw control, over what we have with CLSD's. However, this comes at a price in speed, and it is due to the worm gear diff acting more like an open diff than a CLSD does. If the worm gear unit does not provide less locking torque, it also does not provide greater yaw stability. And it is really not correct to suggest that either alternative is fully equivalent to computerized yaw control.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2008

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

CONTROLLING AWD FRONT/REAR TORQUE SPLIT ON A BUDGET

The last couple of months you have been discussing differentials and so I thought I would throw in my question(s) on the subject.

I am building an AWD car for road racing. Rules require production unit-body with modifications and the rules allow an alternate transmission but orientation must be as from the factory. The engine is transverse and in front of the front axle, will have around 475 hp, weigh about 3100 lb with driver and run the same size tire all around. So far, my research suggests that I should run a LSD in the front and rear positions and a diff (as opposed to a spool) in the middle. Research also brings up "torque split" to reduce on-throttle understeer which, in a passive system, is a final drive ratio difference front-to-rear that is usually achieved through a biased open center diff with a viscous clutch to control front-to-rear slip.

Sounds great, except I don't have access to that kind of center diff and viscous clutch arrangement. Can I achieve the above dynamics by running a different FD in the front and rear diffs, to achieve the "torque split", and run a Torsen or a Weismann diff in the center position? I considered running an open center diff in this case but it seems I would be converting the car to two wheel drive thereby losing the traction advantages of AWD. Or am I stuck with a 50/50 torque split?

The driveline is designed to take a Salisbury-type diff, a spool, a Weismann locker, and I presume an open diff. Any of these can be used in any of the positions. From this you can surmise that Weismann is building the driveline.

They claim their diff is "magic", without saying how it works, and have suggested running it at the front and the rear and a spool in the center. I suspect they may have run this configuration in off-road truck racing. I think this will exacerbate understeer and consequently benefit from a center diff. I have tried to research "Weismann Diff" and "Weismann Locker" but haven't found a good enough explanation of how it works to understand it. Some media say it is locker and some put it in *the same category as a Torsen, i.e. not an overrunning locker but a diff that is "torque-sensing" and biases the torque to the wheel with traction before any overrunning condition sets in.*

The overall problem I am grappling with is how to get a driveline that has the least throttle-on understeer while still maintaining the benefits of AWD traction. The torque split idea seems popular and makes some sense as the car would drive more like RWD but still have part of its power used at the front and in really slippery conditions it would drive like AWD.

What about the locker in the center position? It seems to me that the locker in the center position will act like a spool as a locker only differentiates when one of the outputs is <u>*caused*</u> *to overrun the speed of the diff and the other output (or in this case, one axle overrunning the other axle). The only situation that causes one axle to overrun the other is when turning as the front axle will scribe a larger arc than the rear. But in a race car, where the tires are operating at some degree of slip to generate a cornering force, this speed differential may be insignificant. An alternate scenario is to run different FD ratios front-to-rear in which case the locker will drive the taller gear, as the shorter gear functions as overrun, until the tires at the taller end slip enough to exceed the speed on the other side of the diff. The effect will be a driveline switching between 100% RWD and 100% FWD. AWD drive could be achieved if the slip at one end matches the output speed of the other side of the diff but then you are back to 50/50 torque split and the effect is probably transitory on the way between RWD and FWD. Am I overlooking something here?*

I tried a search for Weismann differential too, and likewise came up dry. Any readers who know more about this, including any working with the company, are invited to fill us in. I am familiar with the Weismann locker, which is a Detroit Locker style design, made to fit road racing dog-ring transaxles. It's been around since the late 1960's. There is no law of nature that says the company has to make only that design forever, but I can at least speak to the characteristics of this device.

For high-speed pavement use, I don’t like a spool in the middle. It’s not so much a matter of drag as a matter of being able to control the car with the throttle in a normal manner. Ideally, you’d like the car to be able to throttle-steer more like a rear-drive car, but have the front drive pick up as the rears begin to slip. A nose-heavy car with 50/50 torque split will sort of do that. When cornering hard, you can only get a little forward acceleration – hence little rearward load transfer – before the rears start to slip, so the rears should spin first. Trouble is, the car understeers steady-state unless you have it roll-stiff in the rear, making the inside rear wheel very lightly loaded in cornering – much like a similar layout with front-wheel drive. That means the rear diff must either be a locker or have enough preload to lock when one wheel is airborne or nearly so. Otherwise, when you try to throttle-steer, you just spin the inside rear and the outside rear stays stuck rather than sliding.

As I see it, if you can’t get a planetary center diff, you have three options: open center diff, with equal final drives or taller final drive at the front; worm gear or clutch pack center diff, with equal final drives; or locker in the middle with taller final drive at the rear.

If you try unequal final drives with a limited-slip in the middle, you'll have lots of drag, friction, heat, and wear, so you're pretty much limited to equal final drives in that case.

If you have an open center diff, and taller final drive in front, the gears in the center diff are working constantly, but at least you're not wearing out any locking mechanism. The mean front/mean rear torque split (at the respective ring gears or diff carriers) has the same proportionality as the respective final drives. For example, if we use a 4:1 at the rear and a 3:1 at the front, the rear wheels get 4/7 (57.1%) of the torque and the fronts get 3/7 (42.9%). The torque split is always in this ratio, and the speeds can be in any ratio.

With a locker in the middle, and taller final drive in the rear, the rear wheels do drive, and the fronts do overrun, most of the time, as the questioner notes. But when the rear wheels spin, and the locker locks, we don't get front-wheel drive. We now have all wheels driven. There is no controlled torque distribution. The torque goes where the resistance is. What is controlled is the ratio of front and rear mean wheel speeds. These will be inversely proportional to the final drive ratios.

Suppose, for example, that we have a 3.00:1 rear final drive and a 3.30:1 front final drive. When the locker is unlocked, we know all the torque is going to the rear. When the locker is locked, we don't know what the torque distribution is, but we know that the mean rear wheel rpm is 10% greater than the mean front wheel rpm. We know that we have to exceed 10% slip at the rear to get lockup. What torque it takes to get this slip will vary, but the 10% threshold is a constant, which we can adjust by our choice of final drive ratios.

This is probably a good set of properties. The one problem is that when lockup comes, it is abrupt. This is an inescapable characteristic of lockers. Either one output shaft is overrunning, or the whole thing is locked. When cornering hard, the lockup will come when the car is at the limit of adhesion and the power is on, and it may upset the car. I would expect you will need some seat time to get used to it.

The center locker does give us front-wheel drive in one situation: when we're backing up. In reverse, a locker drives the faster shaft, and it will not lock. This could be a drawback for off-road use. For racing on pavement, it shouldn't be any worry.

For the front diff, I'd try either a worm gear or clutch pack unit. Clutch pack with preload is best for minimizing inside front wheelspin, but adds understeer and "fight" at the steering wheel. Worm gear is smoother but prone to inside wheelspin. This is less of a problem than it would be in a front-drive car, because the rear wheels have to both spin for either of the fronts to spin.

My best guess for what to try first with your constraints would be: lockers middle and rear; worm gear in front; front final drive about 10% greater numerically than rear final drive; battery, oil tank, and any ballast at the rear.

IS THERE SUCH A THING AS TOO BIG (IN TIRES, THAT IS)?

My question this time is regarding slicks and the width thereof.

I have an MG Midget set up for road racing and at present run 8in slicks. The front is the usual double wishbone and the rear is live axle. All are converted to telescopic shocks.

Given the same suspension system, is there a point where tyres can be too wide? Another midget runs 7in slicks and the comment was made that 7in was as wide as was needed but the previous owner went to 8in as the tyres were easier to get at the time. Would there be any advantage in going to 10in wide? Would the setup need to be changed a lot to get any advantage? I was going to offset the rims so that the extra width added to the overall track.

If a road course were a skidpad, the answer would be fairly simple. In general, wider tires give greater lateral acceleration, within any practical limits, on practically any car, provided that setup and inflation pressure are optimized. However, most road courses have significant straights. In many cases, road races are won on the straightaways. In most of these cases, the straights are at least partially won in the turns: faster cornering lets us start each straightaway at a higher speed, and brake later at the end.

It may be that there are cases where a narrower tire actually will give more cornering power due to higher temperature. Some tread compounds – most, in fact – have a threshold temperature, below which they do not make good grip. However, hotter is not better without limit. Above a certain point, higher temperature hurts grip. It may also cause more severe heat cycling, which makes the tire go off more as a run progresses. If we do have a case where a narrower tire grips better due to temperature, we can say with great confidence that the effect will be highly weather-sensitive: the colder the day, the narrower the tire should be. That makes it difficult to state categorically what tire size we should put on a particular car.

In general, we can say that bigger is better for grip within any practical limits, especially if we have some freedom in choosing compounds. The main drawbacks to big tires are that they weigh more, have more rolling resistance, and make the car wider. Greater car width compromises the line we can take, and it adds aerodynamic drag. The line issue is more important for autocross or hillclimbs on narrow roads than on higher-speed road courses. In general, for road course work, the issue mainly comes down to grip versus drag.

I mentioned the possibility of winning the straightaway in the turn. Perhaps an example will help illustrate. Suppose you are on wide tires and are racing another car on narrow tires, and consequently you have more grip than your competitor, but he has less weight and drag, and therefore has better top speed and top-end acceleration. Suppose your grip advantage is good for a 1mph advantage in a particular turn. Suppose that this competitor is right on your tail approaching the turn.

You will be able to brake later, partly because your car can slow at a greater rate (not only because of greater grip; drag actually helps you here), and partly because you are able to take the turn faster. So you will start pulling out some distance, or at least some time, as soon as your competitor begins braking. You will continue to build your lead through the turn.

At exit, you are still 1mph faster, and still building your lead. Now your competitor can build speed faster than you can, at least if the turn was fast enough so that both cars are now power-limited rather than grip-limited. Is your competitor gaining time on you? Not yet. His drag and weight advantage has to get him up to your speed before he even starts cutting your lead. Once he starts cutting your lead, he has to catch you. Once he catches you, he has to pass you. He has to at least get up even with you before the next braking zone to be able to even have his nose alongside your tail entering the next turn. Remember, he has to brake earlier than you.

If the straight is long enough, he can get by. But the length required is greater than might be imagined.

It will be apparent that in at least some cases, the choice will depend on the course. Long straights and high speeds favor the narrow tires more. A track where most of the lap is spent cornering will favor wide tires more.

The amount of engine power influences the choice as well. Ample power favors wide tires. Meager power favors narrower ones.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

DAMPING RATIO

I would like to ask you an advice about the way to choose the right damping coefficients for a given race car.

I normally work using damping ratio as a index of how much damping I am using. Literature references about mechanical vibrations state that a common sense value for damping ratio should be 0.7, but I have seen a lot of very fast racecars running with damping ratios of more than 1 at least at low speed.

Very often, in these aero cars, there is then a limited difference in slope between low and high speeds of the shaft. Talking with guys who work on post rigs I heard that they choose a damping ratio of about 0.6 but they calculate it in a quite wide range of speeds, so this number is probably an average of what the dampers do in both low and high speed. But how to calculate the relationship between the two areas?

What is your thought about setting the correct damping on a damper and about building its curve from scratch? What is the influence of aerodynamic load on the car?

The really short answer is yes, 0.6 to 0.7 is optimal.

A slightly longer answer is that the matter is controversial, but 0.6 to 0.7 is about midway in the range that various people like.

A better answer than either of those would be that the concept of damping ratio is meaningful, but the behavior of actual cars is so complex that describing it in terms of a damping ratio is only a very crude approximation, albeit a useful one.

For readers who haven't heard of damping ratio, a bit of explanation will be helpful. The damping ratio is a measure of how stiff the damping is, for the setup and the mass of the car. The stiffness of a damper can be expressed as its damping coefficient (C). This is the amount of force the damper produces per unit of shaft velocity. It thus has units of force/velocity – in English units, pounds force per inch per second, or lbf-sec/in; in metric units, newtons per meter per second, or N-sec/m. It corresponds to the slope of the force versus velocity line when we test a damper on a shock dyno.

The damping ratio (ζ, zeta) is the ratio between the damping coefficient we have (C) and the damping coefficient that would give us what is called *critical damping* (C_{crit}). That is,

$$\zeta = C/C_{crit} \tag{1}$$

Critical damping is the least damping that will make the system *non-oscillatory* – that is, will make it return to its neutral or undeflected position after being disturbed, in the least time possible without any overshoot whatsoever, i.e. with no oscillation whatsoever.

For a given spring constant (k) and mass (m),

$$C_{crit} = 2\sqrt{(km)} \tag{2}$$

A professor I work with at UNC Charlotte told me he did some research on damped oscillatory systems and found that the system returns to a reasonable semblance of rest quickest when the damping ratio is such as to allow moderate overshoot: between 0.4 and 0.8. 0.6 would be right in the middle of that range. That is, the system may oscillate a little, but it returns to a condition of near-zero deflection, and near-zero velocity, sooner than it would if critically damped. If the system is more lightly damped than this, it oscillates perceptibly for a significant number of cycles, and stays perceptibly in motion longer.

All this theory has its limits, however, when we try to apply it to actual cars. Engineers like to think in terms of damping coefficients and ratios, because these are used in equations that model oscillatory behavior. If you don't have an input for damping coefficient, you can't use the equations.

Trouble is, real shocks don't have constant damping coefficients. Some come pretty close, but many are deliberately made non-linear. Not only are they made non-linear, often even relatively linear ones are deliberately made stiffer in extension than in compression, or sometimes vice versa. Some are deliberately made position-sensitive. That is, they have different force/velocity characteristics depending on where the piston is in the tube.

Is this complicated enough yet? We're just getting started. Not only is the force/velocity relationship often non-linear, but force is not exclusively a function of velocity at all – even a highly non-linear function. When we valve shocks stiffly, they start to act like springs some of the time: they sometimes make forces in the same direction as velocity rather than the opposite direction. At such times, the damping coefficient goes negative, and so does the damping ratio. Far as I can ascertain, this happens because oil is compressible, and the metal parts in shocks are not perfectly rigid. When the shock is making serious force and the piston reverses direction, the compressibility of the fluid and the expandability of the tube make the fluid on the high-pressure side of the piston spring back. Only when the fluid and tube are done springing back does fluid start to flow through the piston in the direction the designer intended.

Short of this extreme, non-rigidity of fluid and metal may not reverse direction of flow through the piston, but may reduce or increase the rate of flow. Reduction of flow occurs when the piston is gaining speed after a reversal of direction; increase of flow occurs when the piston is slowing down.

In general, it is undesirable to have the shock acting like a spring. Ways to reduce such effects include using a bit more bleed, using lower operating pressures, and making the parts stiffer. Unfortunately, more bleed may conflict with our need for low-speed damping. Lowering working pressure means either making the damper larger in diameter, or giving it a greater motion ratio and a longer stroke. For a given material, making anything stiffer usually means making it heavier. So there are inescapable tradeoffs between making the damper work well and making it small and light.

As if this weren't enough, the spring constant k, as a wheel rate, isn't necessarily constant either. It is often deliberately made to vary with suspension displacement. In some cases the damping coefficient C varies as k varies; in some cases not.

For example, we may use a bump rubber to get a rising k as the suspension enters some range of displacement. In that case, k increases and C does not, so ζ decreases. Or, we might use rocker geometry to make the motion ratio of a coilover increase as the suspension compresses. In this case, the damping and the springing get stiffer together.

It might be supposed that in the latter case, the damping ratio at least stays the same as the suspension stiffens. Nope.

Consider a case where the rocker geometry is such that the coilover-to-wheel motion ratio is 0.5 at static position, and increases by a factor of $\sqrt{2}$ at some amount of compression, making it about 0.7. What happens to k? It doubles. If the shock is linear, what happens to C? It also doubles. And what happens to C_{crit}? It goes up with $\sqrt{k}$, meaning it only goes up by a factor of $\sqrt{2}$. So what happens to the damping ratio, C/C_{crit}? It increases by a factor of $\sqrt{2}$. If it was 0.7 at static, it increases to 1.0.

Using some combination of third spring or ride-only snubber, shock shaft bump rubbers, and rising-ratio rocker geometry, it is possible to arrive at a combination that would give a nearly constant ζ through the full range of travel, if the shock is linear. However, most cars won't be this way, and there is no guarantee that a constant-ζ setup would give us a faster car.

As to the influence of downforce, again there's a short, simple answer: it doesn't do anything. However, as with the previous short, simple answers, the real story isn't nearly so simple.

If we imagine that we have a hypothetical car, with linear springing and damping, and we just add an aerodynamic downforce, what happens? The suspension compresses, but k, C, and m all stay the same, so C_{crit} and ζ don't change either. In particular, adding downforce is not the same as adding mass.

In the real world, however, when we decide to build a downforce car, everything changes. The days are long gone when people would take an existing car and just bolt on some wings. Downforce cars today are heavily dependent on ground effects, which in turn are heavily dependent on height and attitude of the car's underside with respect to the road surface. In order to keep the underside from moving much with respect to the road, we have to spring the car very stiffly. If we want to have at least a bit of compliance in the lower speed ranges, we will design in some amount of rising-rate effect.

So although downforce in itself does not affect damping ratio, the design implications of pursuing downforce, especially under current rules, affect damping ratio a lot.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TRAILING ARMS WITH TORSIONAL RUBBER SPRINGING

I wonder if you are familiar with torsion rubber suspension systems used on travel trailers and boat trailers. For instance see the Henschen dura torque axle system below:
http://www.inlandrv.com/tours/axle-plant/index.htm

There really is just a single trailing link with this system. Where is the roll center? Is it at the center of the axle tube? Is it essentially fixed?

Is the virtual swing arm length infinite?

Has anyone ever tried to use an axle like this on the rear of a fwd vehicle?

This system is a pure trailing arm system with no anti-roll bar. The front view swing arm length is infinite, or undefined. The side view instant center is at the pivot axis.

This creates a system with no camber change in ride, no camber recovery in roll, very little geometric anti-roll, and a lot of geometric anti-lift in braking.

The roll center is not at the pivot axis. It's at ground level in the unrolled condition, and slightly above ground level when cornering, i.e. with some roll and with more than 50% of the lateral force generated by the outside tire.

This particular design of trailer axle uses rubber in a combination of compression, shear, and torsion as the ride springing medium. The assembly includes a square-tube crossmember that mounts it to the frame. Each trailing arm has welded to it a piece of square bar that is inserted into the crossmember. The square bar is sized so that its cross-sectional width across points is slightly smaller than the crossmember tube's inside width across flats. In other words, the square bar is just small enough to fit inside the square tube and turn.

As assembled, the square bar sits at a 45 degree rotational angle inside the tube: its points are centered on the tube's flats. This leaves four roughly triangular cavities between the flats of the bar and the interior surface of the tube. These cavities are filled with rubber.

One might suppose this would be done by bonding custom extrusions to the metal, and also adding some metal retention system to keep the bar in the tube. But that's not what Henschen does. They buy round rubber bars or cords that, as purchased, would not go into the cavity. Then they squash them to the shape of the cavities in a hydraulic press, and freeze them with liquid nitrogen while they're still in the press. When they come out, they will hold their molded shape until they thaw. Then they will try to resume their cylindrical shape.

The rubber pieces and the trailing arm with the square bar welded on are assembled into the crossmember tube while the rubber is frozen. When the rubber thaws and tries to resume its former shape, it becomes a very tight interference fit in its cavity. Far as I can see from the website, this is deemed sufficient to retain the trailing arm to the crossmember with no additional hardware.

It would be possible to add a torsion bar connecting the square bars, to serve as an anti-roll bar, but this is not done. It would be possible to add hydraulic shocks, but that is not done either. The system relies on the internal hysteresis of the rubber for damping.

The result is a very simple suspension system, and a very compact one.

This exact design has not been used at the rear of a fwd vehicle, but certainly a trailing arm suspension with rubber springs has. That's what the original Mini had.

One problem with rubber as a spring material is that it cold flows under sustained load, causing the suspension to sag. The best application for rubber springing is a light trailer that sits unloaded between trips, or a light motorcycle or scooter that sits on its center stand with the suspension extended when not being ridden.

Probably the biggest car to use rubber-in-torsion springs was the Tucker. I'm told that owners of the few surviving Tuckers store their cars on blocks with the suspension at or near full droop to keep the cars from sagging. Sagging is less of a problem when the car is light compared to its cargo, the suspension is stiff, and longevity is not a high design priority. That's how the Mini got away with rubber springs.

With or without rubber springs, a suspension for a car needs proper dampers. Rubber does provide some internal damping, but a car suspension needs real shocks.

RIDE AND ROLL RATES WITH MONOSHOCKS

I was wondering if you might be keen to discuss and explain monoshock suspensions, with particular reference to Dallara F3 cars, and explain how one might go about calculating belleville stack rates, wheel rates, load transfer distributions (bearing in mind the oddity of one wheel going up and the other down at the front in roll) and so on. I am soon to race a 1998 Dallara in Monoposto, having been racing a conventional Reynard 883 for two years, and have been scratching my head as to how to apply the 'classical' equations for load transfer distributions and wheel rates etc. when using a single ride spring and belleville stacks.

The water is further muddied by the application of preload on ride springs and stacks, but I'm hoping to ignore them for now, as one ignores many dynamic aspects in such analysis. Is this wise?

Taking the last question first, no we cannot ignore preload, especially on the washer stacks. On the ride spring, it depends on whether the spring is preloaded at static ride height, or only at full droop.

If the ride spring is preloaded at static, the suspension is solid in ride (wheel rate approaching infinity, or undefined) until the preload is exceeded. If it is not preloaded at static, it provides a rate for the wheel pair that is equal to the spring rate times the square of the spring-to-wheel motion ratio. This is the number of pounds of load change for the wheel pair, per inch of ride travel. It corresponds to the sum of the two individual wheel rates in ride in a conventional suspension. So if you have an equation that calls for individual wheel rate in ride, you use half of that ride-spring-rate-times-square-of-motion-ratio quantity.

If the equation in question is for lateral load transfer, you use zero for the ride spring rate. The ride spring in a monoshock suspension acts only in ride and does not contribute to roll resistance or elastic lateral load transfer at all.

All the elastic roll resistance in a monoshock setup comes from the Belleville washer stacks. There is a stack on each side of the rocker where the shuttle passes through. Ordinarily, we use identical stacks on both sides of the rocker. The stacks act in parallel, so the rate of force change with respect to displacement at the shuttle is the sum of the rates of the individual stacks, if they are both active. They may, or may not, both be active. That's where preload comes in.

When both stacks are equally preloaded, there is a compressive load on each stack, and a reaction force from each stack trying to extend itself. But since these act on the shuttle in opposite directions, there is no net force trying to move the shuttle to either side.

Correspondingly, if there is an increase in extension force from one stack, and a decrease in extension force on the other, those force changes act additively, and the force trying to re-center the shuttle is the sum of the two.

That force acts on the wheels, through the pushrods and the rest of the suspension, at some motion ratio. The rate of force change with respect to displacement at the wheels is the rate at the shuttle, times the square of the motion ratio.

Like an anti-roll bar, the shuttle mechanism is an interconnective springing system. It generates force in response to a displacement difference between two wheels, i.e. an oppositional displacement of the pair. To define a motion ratio for such a system, we need to resolve the question of what we call an inch (or mm) of motion: is it an inch (or mm) at each wheel, meaning two inches (or millimeters) difference between the two, or is it half an inch (or mm) at each wheel, meaning one inch (or mm) difference between the two?

I prefer the former method, because it puts wheel rates for all modes, from all springing devices, in the same terms: force per unit of displacement per wheel. Using this method, the angular roll resistance is the wheel rate in roll, times the square of the track, times ½. That gives the angular rate in lb-in or N-mm per radian. To get lb-in or N-mm per degree, divide by 180/π or 57.3.

Most books use the method above to calculate the component of angular roll resistance due to the ride springs (again, this is zero for a monoshock suspension), and use a different formula for the component due to the anti-roll bar. The more common method for the bar component is as above, except the rate is taken as force per unit of displacement per wheel pair, and the angular rate from the bar is then as above, except with the multiplication by ½ omitted.

Both methods work fine, and give the same answer for total angular rate, provided you use the method that goes with your expression of rate.

The angular roll displacement is then the sprung mass, times the lateral acceleration, times the moment arm of the sprung mass c.g. about the roll axis, divided by the sum of the front and rear angular roll resistances. The front elastic lateral load transfer at that roll displacement is front angular roll resistance, times angular roll displacement, divided by front track.

If you are setting up a factory-built race car with a monoshock suspension, the factory will generally furnish a table stating the rates for various stacks. If you are building a car, or want to try stacks that aren't in the table, it helps to understand the basic rules governing the behavior of Belleville washer stacks.

For readers unfamiliar with Belleville washers, they are simply spring steel flat washers, dished a bit so they can serve as a short-travel compression spring. The dished washer has a concave, or cup, side, and a convex, or cone, side. If we stack a number of washers cup-to-cup and cone-to-cone, we have a compression spring with a useful amount of travel.

A Belleville washer is not a perfectly linear spring, but if we approximate its behavior as a constant rate, and call that k, and if we have a number of washers which we call N, we can state some rules about the rates of combinations of such washers.

If we stack N washers cup-to-cone, or nest them, they act in parallel and the rate is kN. If we stack them cup-to-cup and cone-to-cone, they act in series and the rate is k/N. If we have two stacks, each of an overall rate K, on each side of the rocker, and both are active, the rate of the assembly is 2K.

Both stacks are active as long as there is a load on both. Once the inside stack (the one toward the inside of the turn) unloads, it ceases to contribute to the rate, and the rate from that displacement on is K rather than 2K. If the stacks are not preloaded at all, only the outside one is ever active, and the rate is always K. If the stacks are loose on the shuttle at static, the rate is zero until the clearance on the outside stack is taken up, and then it's K.

It is possible to create stacks in which some of the washers are nested and some act singly. In that case we can have a rising-rate stack. If the portion with nested washers has a rate L and the portion with unnested washers has a rate M, then the whole stack has a rate of LM/(L+M). If we compress this stack far enough, the unnested washers bottom out (get squashed completely flat) and we are left with the nested ones still active. The rate then goes to L (possibly in a handbasket).

When we have variable-rate stacks, either rising-rate ones or ones where there is a chance of a preloaded stack unloading, we have to determine at what roll displacement we will get a rate change, and compare that to our predicted roll displacement using the rate we have at static. If this comparison shows that we will encounter a rate change, we have to work backwards and calculate at what lateral acceleration we encounter the rate change. We then calculate what load transfers we have at that lateral acceleration, using the angular roll resistance rates for the first part of the travel, and finally we calculate further load transfer for the additional increment of lateral acceleration, using the rate that applies for that increment.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
April 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

"MINDING YOUR ANTI" VIDEO NOW ON DVD

"Minding Your Anti: Understanding Factors in Load Transfer (Weight Transfer)" is a video of a one-hour lecture I delivered at UNC Charlotte in 2003. Until now, it has only been available on US-standard VHS video cassettes. I now have it on DVD. I also have a few VHS versions left in stock. Price for either format is US$50.00, which includes shipping and handling worldwide. To order, send check or money order to me at 155 Wankel Dr., Kannapolis, NC 28083-8200, USA.

RAISED LOWER CONTROL ARMS ON 2009 F1 CARS

I'm sure you've noted that recently there have been radical changes in the suspension geometry at the front of F-1 cars. The changes permit the front of the cars to be raised in order to increase the volume of air supplied to the aerodynamic features under the cars which produce down force.

Could you comment on the new geometry and how it affects the camber change curve, virtual swing arm length, instant center location, roll center height and any other significant ramifications.

First, a bit of explanation for readers who haven't been following Formula 1 lately: this year's cars have the lower control arms dramatically higher than previous cars. They are higher at both the inboard and outboard ends. Purpose of this is to get the arms up out of the airstream. This allows aerodynamic elements under the driver's legs to be more effective.

It's all related to efforts to keep the cars from losing so much front downforce when running in another car's wake, and thereby enable more passing.

The noses aren't really higher than before. Drivers' feet were already up to where they could barely see over their toes. What's different now is that the region below the footwell has been aerodynamically cleaned up.

Previously, there used to be either a single projecting pylon under the nose, or two of them, to pick up the inboard ends of the lower control arms.

If designers were just raising the inboard pivots of the lower arm, that would increase camber recovery in roll, increase camber change in ride, and raise the roll center. But that's not what they're doing. They're raising the outboard end along with the inboard points. This keeps the instant center position about the same as before, so the camber properties and geometric anti-roll are similar also. Looking at pictures of the cars, they all seem to have very little camber recovery, and roll center heights from 1 to 3 inches, which is about what they had in previous years.

In theory, it is possible to have any instant center with any control arm heights, packaging constraints permitting. The inescapable penalty if we raise the lower arms, or lower the upper ones, is that the loads on the arms in cornering and braking increase (assuming outboard brakes). So do the ball joint loads. That's why designers didn't raise the lower arms years ago. Now most of them have decided that aero considerations trump structural ones, under present rules.

Under present rules or previous ones, it is aerodynamically desirable to get the nose as high and as slender as possible. That means the master cylinders have to go in front of the pedals, rather than above or below, and the pedals have to pivot at the bottom. All of that means the steering rack has to lie above the master cylinders. Finally, for reasons of steering geometry and aerodynamics, the upper ball joints need to lie inside the wheel rim, whose diameter is dictated by the rules. These constraints require the tie rod and upper control arm to angle downward somewhat from the tub to the wheel.

With the upper arm angle constrained that way, the only way to get any serious camber recovery in roll would be to have a high roll center, which is undesirable because it causes jacking. So the designers accept having poor camber recovery in roll, and deal with this by just making the suspension so stiff that very little roll occurs, and running static negative camber.

PULLROD SUSPENSIONS IN F1

I am hearing all sorts of outrageous claims about the pullrod suspension used on the rear of the highly successful Red Bull F1 car. Does pullrod suspension actually offer any clear advantage over pushrod?

In terms of suspension dynamics, no. Both layouts can be made to have the same wheel rate, the same motion ratio/displacement curve, and so on. Both affect dynamic load transfer in the same manner.

The differences come down to packaging, with small effects on overall c.g. height, aerodynamics, and component accessibility.

Pullrod suspension generally places the shocks, springs, and rockers low in the car. Depending on whether something else has to be moved to make room, this generally lowers the overall c.g. of the car, which is good. Since these components are a fairly small part of the total mass, the benefit is correspondingly small, but it is real.

Accessibility of the shocks, springs, and rockers generally suffers. It is easier to get at these parts when they are on top of the transaxle or footbox than when they are down under other components, a few inches off the ground. However, when millions of dollars are on the line, it can make sense to put up with some extra hassle to gain a small performance advantage.

The shocks, springs, and rockers take up space, and create a wider package near the ground. If they get in the way of under-car airflow, it is often a better choice to use pushrods.

This last consideration was crucial in making pullrods a rational choice for the rear of the car under th new 2009 rules. As part of the effort to reduce and narrow the upwash behind the car, diffusers were required to start further back than before, and the flat portion of the underside had to extend further aft. That creates the opportunity to put components low down alongside the transaxle, without aerodynamic penalty. The Red Bull team simply saw this opportunity, and used it.

Some have suggested that there is a reduction of overall weight with pullrods. I am skeptical. It is true that a pullrod can be made more slender than a pushrod, and perhaps lighter, because it does not have to withstand large compression loads, and therefore does not need as much buckling resistance. However, compressive loadings on the upper control arms increase, and the weight can easily come back there. Any net weight reduction probably depends on the particulars of the individual design.

The Red Bull car's overall success should not be attributed to this one feature. Rather, it stems from a larger willingness to consider new possibilities, combined with the engineering understanding to properly evaluate, select, and apply such possibilities, and to integrate and optimize the total package.

RATE AND LENGTH NUMBERS ON TORSION BARS

What do the numbers stamped on the end of torsion bars mean (race car ones, like for sprint cars and midgets)? I understand one is the length and the other is the rate. The one for the length makes some sense: it appears to just be the overall length in inches. But I can't see what the one for rate relates to. Can you explain?

The number people sometimes call "rate" is actually the effective diameter for the active part of the bar, in thousandths of an inch. The active portion is the turned-down portion in the middle, where most of the twisting occurs.

The effective diameter is the actual diameter, if the bar is solid. If it's hollow, it's the diameter of an equivalent-rate solid bar.

From these numbers, you can calculate a rate at the end of a known-length lever arm that equates to the rate of a coil spring acting at a similar point. In most cases, you don't need to do that, because torsion bar manufacturers offer charts in their catalogs and on their websites. But here's the equation:

$$S = (\pi G d^4) / (32R^2L) = (.098Gd^4) / (R^2L) \qquad (1)$$

where:

S = rate at lever arm end, lb/in
G = shear modulus modulus of material
11,500,000 lb/in^2 for most steels, 11,800,000 lb/in^2 for spring steel
d = diameter of the active portion of the bar, inches
R = effective length (moment arm length) of lever arm, inches
L = length of active portion of the bar, inches

For d, you use the "rate" number stamped on the bar, with a decimal point three places from the right. For L, you use either your own measured length for the active portion, or, as a rule of thumb, the stamped or cataloged length minus four inches. R is the distance from bar axis to lever arm end, not along the arm if it's curved or angled, but perpendicular to the bar axis.

If you have a bar with no stamped numbers, or with the numbers worn off, and it's hollow, you can calculate the rate using measured inside and outside diameters of the bar, with the following formula:

$$S = \left(\pi G(d_o^4 - d_i^4)\right) / (32R^2L) = \left(.098G(d_o^4 - d_i^4)\right) / (R^2L) \qquad (2)$$

where:

d_o = outside diameter, inches
d_i = inside diameter, inches
All other variables are as in Equation (1).

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
May/June 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: mortiz49@earthlink.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

THE ANTI CONTROVERSY

A number of people have asked me to write a response to an article in the March 2009 *Racecar Engineering* by Danny Nowlan, in which the author makes some assertions that challenge conventional thinking (if there is such a thing) about anti-squat and other geometric pitch and roll resistance phenomena, and contradict some of what I say in my "Minding Your Anti" video. I will try to do that this month.

Let me begin by emphasizing that I intend no disparagement to Mr. Nowlan's integrity, intelligence, or intentions, nor to those of any of my colleagues in the field of vehicle dynamics, even if they have reached different conclusions than I have about how things work. One might suppose that vehicle dynamics, being merely Newtonian physics, would not be controversial or hard to understand for educated people in the twenty-first century. But any one who has delved into the field will have noticed that there is a lot of disagreement, and sorting through all the contradictions is anything but easy. In this situation, surely anybody who makes an honest effort to understand the laws of nature, and to deliver honest goods and services to the public, deserves the benefit of the doubt regarding their intentions.

I should also mention that Mr. Nowlan explicitly invited responses in his article, from any who might disagree with what he was saying.

Some readers will not have seen Mr. Nowlan's article. I will summarize his assertions, as faithfully as I can, in my own words:

1. *We must not be slaves to orthodoxy. We must be willing to think things through from first principles, and entertain any conclusions that may result, even if they contradict what accepted authorities say.*

2. *The car has only one center of mass. It is not a string of separate masses, and should not be analyzed as such.*

3. *All dynamic effects should be analyzed as forces and moments about that single center of mass.*

4. *The laws of physics act similarly in all directions. Therefore, we should be able to use fundamentally similar methods to analyze front-view and side-view suspension geometry.*

5. *In front-view geometry, the point where the force line intercepts the c.g. plane is the roll center.*

6. *Similarly, in side view, the point where the side-view force line intercepts the c.g. plane may be taken as a pitch center.*

7. *All forces exerted by the tire upon the road, and by the road upon the tire, act at the ground plane, not at the wheel center.*

8. *Any time a tire is exerting a longitudinal force, it is necessarily transmitting a torque.*

9. *Therefore, there is no dynamic difference between systems with inboard and outboard brakes, or between systems with sprung and unsprung differentials.*

10. *Accepted theory asserts that with inboard brakes or sprung diffs, forces are applied to the car at wheel center height.*

11. *Accepted theory thus predicts ample longitudinal anti for such systems, but data traces from actual cars contradict this. Suspension displacements are very nearly what we would expect from pitch moments about the ground, reacted entirely by the car's springs, implying little or no anti, or a low pitch center.*

Taking these more or less in order, I agree with #1. Willingness to re-examine and question accepted theory is the difference between science and dogma, and is an important foundational principle of human liberty in general.

#2 is correct, almost. That is, the car is a collection of masses that are rigidly attached to each other, and can therefore be considered a single body – except that when the car has a suspension system, some of the attachments within it are not rigid, and some components can move with respect to others, albeit in a (hopefully) limited and carefully-controlled manner. For some purposes, rigorous analysis requires us to treat these masses separately. Moreover, the masses are separate to different degrees in different modes. The nuances of this will make a good future article or newsletter.

However, we can examine the principles of anti effects using an assumption that all of the car's mass is sprung, and the masses and inertias of unsprung components are zero. Real cars aren't like that of course, but imagining so greatly simplifies the discussion and the associated equations.

Since we will be presenting equations, it is time for some definitions.

First, the vehicle axis system:

Per SAE convention, x is longitudinal, y is transverse, and z is normal or vertical. Unless otherwise stated, sign conventions are: x positive forward; y positive rightward; z positive downward. Unless otherwise stated, vehicle origin is per SAE aerodynamic axis system: on the ground plane, midway along the wheelbase, and centered with respect to the front and rear tracks. However, for particular purposes, we may use local origins and sign conventions may deviate, so it is necessary to pay attention to context. Even then, x will always be longitudinal, y transverse, and z vertical, in some sense that is appropriate to the context.

SAE vehicle axis convention takes as origin a point on the roll axis, directly below the sprung mass center of mass or c.g. I don't care for that, because the points in question move and are subject to uncertainties. Actually, there is no origin that is completely immune to some uncertainties and relative movements, but for our purposes here, we will use the SAE aerodynamic convention.

x, y, and z may denote coordinates in this axis system, or linear displacements in this axis system.

Next, angular quantities about the axes:

φ (phi) is roll, or angular movement about an x axis, conventionally positive rightward, or clockwise as seen looking forward.
M_x or M_φ is a moment in roll, or about an x axis.

θ (theta) is pitch, or angular movement about a y axis, conventionally positive rearward, or clockwise when looking from left to right.
M_y or M_θ is a moment in pitch, or about a y axis.

ψ (psi) is yaw, or angular movement about a z axis, conventionally positive rightward, or clockwise when looking down.
M_z or M_ψ is a moment in yaw, or about a z axis.

L is center spacing between two wheels of a pair under consideration.
L_x is wheelbase.
L_y is track.

r_t is tire loaded radius. For simplicity, we will ignore distinctions between loaded and effective tire radius.

F is force. F with an x, y, or z subscript is force in an x, y, or z direction.
m is mass.
a is acceleration.
F = ma. When m is expressed in pounds, a should be in g's.

H is height.
H_{cg} is height of the center of mass or center of gravity (c.g.).
H_{rc} is roll center height.
H_{pc} is pitch center height.

dx/dz is the instantaneous rate of change of a point's x coordinate with respect to change in its z coordinate: the first derivative of x displacement with respect to z displacement.
dy/dz, similarly, is the first derivative of y displacement with respect to z displacement.

dx/dz and dy/dz are also the instantaneous slopes or inclinations, from vertical, of the contact patch center's path of motion in side and front view respectively.

A force line is a notional line of action for the vector sum of an x or y ground-plane force at a tire contact patch and the induced z-direction support force within the suspension system that results from angularities in the suspension linkage. In front-view geometry, the force line is the line from the contact patch center to the front-view instant center. The front-view force line is an instantaneous perpendicular to the contact patch center's path of motion in front view, and has a slope or dz/dy, relative to ground plane, equal to the contact patch center's dy/dz. The side view force line is an analogous construction in side view, whose slope is the contact patch center's dx/dz. There are some subtleties to assigning that dx/dz, which we will address.

The front-view and side-view force lines define a plane, which we may call a force plane.

Returning to Mr. Nowlan's assertions, #3 isn't necessarily correct. In particular, it is incorrect to suppose that forces at the contact patches act along the force lines or force planes of the suspension geometry, and therefore exert moments about the center of mass according to their vertical or perpendicular distance from the center of mass. This misconception stems from a misunderstanding of what the force lines and planes represent.

The car is not a body floating in space, acted upon by angular forces in the force planes. Nor is it an airplane with four wings or control surfaces each exerting a lift and a drag force. It is a body with a center of mass above the ground plane, supported through compliant suspension at four points in the ground plane, accelerated in the x and y directions by forces in the ground plane. These forces do not act on the car as a whole along the force planes; they act on the car as a whole along the ground plane. The ground-plane forces induce support forces <u>within</u> the car, in the suspension systems, according to the slopes of the force planes. These induced forces create moments within the vehicle that may oppose or exaggerate the suspension displacements in roll and pitch, and can alter the distribution of tire normal forces.

This may seem merely a rhetorical distinction, but it's not. In particular, the anti-roll and anti-pitch moments do not depend on the proximity of the force plane to the center of mass. Rather, they depend only on the spacing of the wheels, the magnitude of the ground-plane forces, and the slopes of the force planes. The anti-roll and anti-pitch moments do not depend in any direct way on where the center of mass is.

#4 is correct. The laws of physics do act similarly in all directions, and we should be able to apply fundamentally similar methodology for both side-view (pitch-related) and front-view (roll-related) phenomena. However, #5 is incorrect, and consequently so is #6. The roll center is not where a single wheel's force line intercepts the c.g. plane. Correspondingly, applying similar thinking to the side view is likewise incorrect.

Indeed, we cannot meaningfully speak of a roll center or pitch center for a single wheel, even in simple two-dimensional modeling. Even in 2D, we need two points of support to resist an overturning moment, and it is the combined effects at these two support systems that create any moment. Likewise, we cannot meaningfully speak of right-side or left-side roll, or front-end or rear-end pitch.

What is the roll center, then? It is a notional coupling point for transmission of y forces between the notional front or rear axle and the sprung mass. Its height is the height at which the centrifugal inertia force for that end of the car would generate a roll moment exactly equal to the anti-roll moment induced by the linkage geometry, so that no roll would result.

Importantly, the roll center is not any sort of coupling point for z forces. It is analogous to a horizontal link, or a roller in a vertical slot, not a pin in a hole. When the car is in motion, a sustained z force does not induce a geometric roll moment. Engineers were confused about this for years, partly because when a car is tested sitting still on a shop floor, with no slip plates under the tires, asymmetries in its anti-roll actually will produce roll moments in response to z forces. But when the wheels are free to track in and out as they roll along, that effect goes away. This is why we roll a car before we scale it.

The equation for the roll moment M_x or M_φ produced by a total wheel pair y force F_y acting at a height H is:

$$M_x = F_y H \tag{1}$$

The equation for the anti-roll moment induced in the suspension by right and left wheel ground-plane y forces F_{yR} and F_{yL} is:

$$M_x = \left(F_{yR}(dy/dz)_R - F_{yL}(dy/dz)_L\right)L_y/2 \qquad (2)$$

Or, in English, the geometric anti-roll moment is equal to the difference of the two induced support forces, times half the track. (The sum of the two support forces is the net jacking force for the wheel pair.) Stated a bit differently, the moment in roll for the wheel pair induced by the linkage geometry has to equal the difference in the support forces induced, reacted on the half-track. If our modeling gives us some other value for this moment, there has to be an error in our modeling.

Note also that this equation includes no term representing c.g. location.

Combining equations (1) and (2), for a condition where H is the height of the roll center, H_{rc}, and the roll moment and the geometric anti-roll moment are equal in magnitude:

$$F_yH_{rc} = \left(F_{yR}(dy/dz)_R - F_L(dy/dz)_L\right)L_y/2 \qquad (3a)$$

Solving for H_{rc}:

$$H_{rc} = \left((F_{yR}/F_y)(dy/dz)_R - (F_{yL}/F_y)(dy/dz)_L\right)L_y/2 \qquad (3b)$$

This can also be written:

$$H_{rc} = \left(L_y(F_{yR}/F_y)(dy/dz)_R - L_y(F_{yL}/F_y)(dy/dz)_L\right)/2 \qquad (3c)$$

Or, in English, the portion of the total y force at the right wheel, times the track, times the right wheel's force line slope, minus the portion of total y force at the left wheel, times the track, times the left wheel's force line slope, all divided by two, is the roll center height.

We should note that when both wheels have anti-roll, i.e. when both force lines slope upward toward the center of the car, those two force line slopes are mathematically of opposite sign, so the difference of the two products in the large parentheses ends up with a larger absolute value than either product by itself: the two products either add, because we're subtracting a negative, or add negatively, because we're subtracting a positive from a negative.

Doing the same thing graphically, we construct what I call a resolution line. This is a vertical line in our front-view drawing, whose y position relates to the distribution of y force between the right and left tires. If, for example, the right tire is estimated to make 70% of the cornering force for the pair, the resolution line is 70% of the track from that wheel (30% of the track from the left wheel). We then find the two points where the front-view force lines intercept this resolution line, average those heights (add them and divide by two), and that's the roll center height.

My video includes illustrations.

How does this compare with Bill Mitchell's force application point method? He draws a vertical line at the c.g., finds the force line intercepts of that, and does a weighted average of those, with the weighting based on the F_{yR}/F_{yL} distribution.

Turns out that the FAP method gives the same answer as the resolution line method, when the c.g. is midway between the wheels. It also gives the same answer when the force lines have equal and opposite slope, i.e. when they cross in the middle of the car, even if the c.g. is offset. When the c.g. is laterally offset, and the suspension is also markedly asymmetrical, it creates an error. If the c.g. offset is small, the error is too small to measure experimentally on a kinematics and compliance rig. It would become a concern in a car that has a large lateral c.g. offset, and markedly asymmetrical independent suspension – for example, perhaps a Supermodified with independent suspension, in a rolled condition. (Such cars once existed, but Supermodifieds are now required to have beam axles at both ends.)

The oldest method of assigning a roll center is to take the front-view force line intersection as the roll center. The force line intersection method can also give a correct answer, but the only time it necessarily does is when the force line slopes are equal and opposite. The more asymmetry we introduce in the geometry, the more wildly anomalous this model tends to become. The most obvious anomalies occur when one wheel has slight anti-roll and the other has slight pro-roll, so that the force line intersection lies a long way outside the vehicle. The intersection can then be far above or below the ground, despite the fact that the suspension clearly generates little net pro-roll or anti-roll moment. And if we further suppose that z forces generate moments about this point, the modeling anomalies can get really outrageous. Finally, in some cases the force lines are parallel, and they have no intersection.

Since the roll center does not transmit z forces, it doesn't matter what its lateral location or y coordinate is. I call it undefined. If you like to think of it as being under the c.g., that's fine. If you like to think of it as midway between the wheels, that's fine. Just don't think of it as something that can transmit force vertically.

Regarding #7 through #11, it is true that the car and the road can only apply force to each other where they touch, i.e. at the contact patches. However, this does not mean that conventional thinking on anti-squat is incorrect, although it may be that some accepted presentations could be clearer.

It is also true that if there is an x force at the contact patch, there is a torque at the wheel, with the exception of the case where the x force is induced by the tire running at a slip angle when cornering.

However, the torque may or may not act through the suspension linkage, and accordingly may or may not contribute to the M_θ in the suspension.

Let's imagine a free-body diagram for the wheel and stub axle in an independent rear suspension, in side view, when the wheel is propelling the car. The tire exerts a rearward force upon the pavement, and the pavement exerts an equal and opposite forward force F_x upon the tire, which propels the car. The wheel and axle assembly exerts an equal forward force on the rest of the car (remember we are imagining that the wheel assembly doesn't have any mass of its own). The remainder of the car, acting through the linkage and upright and bearings, exerts an equal and opposite rearward inertia force upon the axle. Thus the sum of the wheel and axle assembly's forces is zero, and there is a couple, equal to the wheel's torque M_y, which must be reacted somehow, for the sum of the moments to also be zero.

$$M_y = F_x r_t \tag{4}$$

In most independent suspensions, the torque is applied through the drive shaft. The drive shaft is jointed so that it can only apply rotational force. It applies this force directly to the axle. It does not apply rotational force to the upright. The shaft sees a rearward torque at its outboard end and a forward torque at its inboard end. The stub axle at the diff, driving the axle, sees a rearward torque from the axle. This in turn reacts through the diff, the pinion bearings, and the transaxle mounts.

This torque is thus applied to the car as a whole, but the load path does not include the suspension linkage.

Imagining the free-body diagram for the upright, then, we have a forward force (thrust force) exerted at the wheel bearings, but no torque applied at the wheel axis. Both longitudinal links, or side-view projected control arms (SVCA's), exert rearward force upon the upright and forward force upon the sprung structure. These forces are both smaller than the total F_x at the wheel center, and both in the same direction. Their magnitudes are inversely proportional to the distances from the hub or wheel center, and there are induced z forces within the suspension depending on the magnitude of the x forces and the SVCA angularities. The overall z force F_z induced is F_x times the dx/dz value at the wheel center:

$$F_z = F_x(dx/dz) \tag{5}$$

Why not the dx/dz at the contact patch? Actually, the dx/dz at the contact patch is the same! If we constrain the wheel rotationally (lock it) at the inboard end of the drive shaft, and then move the suspension, the wheel doesn't rotate with the upright, and the point at the static contact patch center stays directly below the wheel center through the entire travel. Thus, its dx/dz is the same as that of the wheel center. This differs from a situation where torque is applied to the upright and does act through the linkage, as with an outboard brake. In that case, the dx/dz that matters is that of the contact patch center, assuming the wheel rotates with the upright. The force line then runs from the contact patch center to the instant center, and has a slope equal to the contact patch dx/dz for a wheel that rotates with the upright. Also, the forces at the SVCA's are opposite in direction, with the force on the lower opposite to and greater than the F_x at the contact patch, and the one on the upper in the same direction as the contact patch F_x and smaller than the force on the lower: clearly a different picture, and one that will usually produce a different induced z force.

Kinematics and compliance testing consistently confirms that it matters mightily whether torque reacts through the linkage or not. The only case where it does not matter is where the upright does not rotate in side view as the suspension moves (SVCA's parallel to each other).

With independent suspension or inboard brakes, the force line has the same slope as a line from the wheel center to the side-view instant center (SVIC), not the slope of a line from the contact patch center to the SVIC. This is what some authors mean when they say that the force acts at the wheel center. They don't mean that the rearward pitch moment M_θ is equal to $F_x(H_{cg} - r_t)$. It's still F_xH_{cg}. So conventional theory does not predict huge anti-squat for independent suspension. On the contrary, anti-squat by conventional theory is generally very modest for independent suspension.

So, in #11, Mr. Nowlan has misunderstood the theory he is attempting to refute, which is in fact confirmed by the data trace he presents.

Taking Terry Satchell's chapter in Milliken and Milliken's *Race Car Vehicle Dynamics* as an example of a conventional presentation, when F_x is considered as acting at wheel center height, a different rule for determining percent anti is used, and the result is a correct calculation when the full method is taken together.

If we insist on considering the force line as passing through the instant center, that method is correct. I prefer to say that the force line passes through the contact patch center but has the same slope as a line from wheel center to instant center. I then use the same rule for percent anti as for a system where the torque does act through the linkage. I find this a slightly clearer way of thinking about the matter, but both methods give the same answer.

At this point, we should probably review what is meant by percent anti-squat, since Mr. Nowlan seems a bit confused about that. Percent anti-squat is the percentage by which the anti effect reduces rear suspension compression, compared to what would occur with zero anti. It is not the same as percent anti-pitch.

Here is the equation for percent anti-squat P_{as}, per my method:

$$P_{as} = \left(L_x(dx/dz)/H_{cg}\right)*100\% \qquad (6a)$$

Or, by Mr. Satchell's method:

$$P_{as} = \Big(\big((r_t + L_x(dx/dz)) - r_t\big)/H_{cg}\Big)*100\% \qquad (6b)$$

It will be apparent that (6a) and (6b) are algebraically equivalent.

Graphically, what Mr. Satchell and others do is shift the force line up to the wheel center, but then determine anti-squat based on where the force line intercepts the front axle plane, not relative to the ground and the c.g. height, but relative to a point a tire radius above the ground and a point a tire radius above the c.g. In other words, they draw the same picture I do, except everything is moved up a tire radius, and the final answer is the same.

Mr. Nowlan poses an example where $L_x(dx/dz) = .20H_{cg}$ (force line intercepts front axle plane at 20% of c.g. height), and the c.g. is in the middle of the wheelbase, so that the height of the force line at the c.g. plane is 10% of c.g. height. He says that the c.g. plane/force line intersection should be considered the pitch center, and also the anti-squat is 10%.

Actually, the anti-squat is 20%. But this doesn't mean the car pitches 20% less than it would with no anti. It means the rear compresses 20% less, and the front extends the same as it would with no anti-squat. Remember, there is no such thing as rear pitch. Pitch is the relative displacement of the front and rear, or alternatively, the angular displacement resulting from that. For the simplest case to analyze, where the front and rear have equal wheel rates in pitch as well as equal static weight, if the front and rear displacement with no anti is identical in magnitude, and we call that absolute displacement z, then with no anti, the relative displacement front to rear is z + z, or 2z. With the anti, the relative displacement is z + .80z, or 1.8z. That's 10% less than 2z. So, for this simplified case, a 20% squat reduction produces a 10% pitch reduction.

Note, however, that we could move the c.g. forward or back, and if we don't change the springs, we still get the same deflections in response to a given z force, and the same anti-squat. If we make just the front springs softer, the pitch increases, because the front end rises more, but the squat and anti-squat stay the same.

Therefore, we cannot say that the c.g. plane/force line intercept can be called a pitch center, nor that it necessarily tells us anything definite about what the car will do.

Mr. Nowlan does not elaborate on what he recommends when both front and rear wheels are making x forces, but definitely neither will the average of the c.g. plane/force line intercepts serve as a pitch center, nor will a weighted average (except when the c.g. is midway along the wheelbase), nor will the side-view intersection of the force lines.

Well then – is the idea of a pitch center analogous to a roll center useful or relevant at all? I think it is. And I think we can assign it using a method closely resembling our front-view approach. Of course, this will only work if the front-view approach is correct, and if we adapt it correctly.

One issue we need to address is that the term "pitch center" is already in use in another area of vehicle dynamics, with a meaning not analogous to a roll center. Ride engineers speak of pitch centers and bounce centers. The pitch center is a center of rotation about which the car oscillates in response to sequential perturbations at the front and rear, as when going over a short speed bump. It is determined by the relationship of the front and rear static deflections or natural frequencies, and does not relate to pitch moments created by x accelerations, or opposing or exaggerating pitch moments induced within the suspension as a result of x ground-plane forces.

What would a pitch center analogous to a roll center be? It would be a notional coupling point between sprung mass and wheel pair suspension, for the total x force of a right or left wheel pair, with front/rear x force distribution close to actual for the situation being modeled. It would have a height H_{pc} such that the total F_x for that side of the car, applied at a height H_{pc}, would produce no net pitch, because the geometry would produce an equal and opposite anti-pitch moment. The equations for this would be of the same form as equations (3a), (3b), and (3c):

$$F_x H_{pc} = \big(F_{xf}(dx/dz)_f - F_{xr}(dx/dz)_r\big)L_x/2 \tag{7a}$$

$$H_{pc} = \big((F_{xf}/F_x)(dx/dz)_f - (F_{xr}/F_y)(dx/dz)_r\big)L_x/2 \tag{7b}$$

$$H_{pc} = \big(L_x(F_{xf}/F_x)(dx/dz)_f - L_x(F_{xr}/F_x)(dx/dz)_r\big)/2 \tag{7c}$$

Or, in English, the portion of the total x force at the front wheel, times the wheelbase, times the front wheel's force line slope, minus the portion of total x force at the rear wheel, times the wheelbase, times the rear wheel's force line slope, all divided by two, is the pitch center height.

For anti-squat in a rear-drive car, $F_{xr}/F_x = 1$, and $F_{xf}/F_x = 0$. With all-wheel drive, F_{xf}/F_{xr} depends on the front/rear thrust ratio. In braking, F_{xf}/F_{xr} depends on the overall front/rear retardation force ratio.

With sprung diffs or inboard brakes, dx/dz is the value for a wheel that does not rotate with the upright. For outboard brakes, live axles, or odd cases such as Humvees and VW Transporters where the upright contains drop gears, dx/dz is the value for a wheel that does rotate with the upright or axle.

Doing the work graphically, we construct a vertical resolution line, positioned according to the F_{xf}/F_{xr} relationship. For propulsion with rear drive, the resolution line is at the front axle. For braking, it is aft of the front axle by a distance equal to the front overall retardation percentage times the wheelbase.

Force lines are as described above: from the contact patch to the SVIC where torque does react through the linkage, and where torque does not react through the linkage, through the contact patch center and parallel to a line through the wheel center and SVIC.

We then find the intercepts of the front and rear force lines with the resolution line, average these, and that's our pitch center height.

The left and right pitch centers define a pitch axis, analogous to a roll axis. The left and right F_x values, times the respective H_{pc} values, divided by the respective wheelbases, are the left and right geometric load transfers. The total F_x for both sides, times the moment arm of the c.g. about the pitch axis, divided by the total angular elastic pitch resistance rate, is the pitch angle. The right and left angular elastic anti-pitch moments, divided by the respective wheelbases, are the right and left elastic load transfers.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net . Note that this is a new e-mail address. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

THE UNSPRUNG COMPONENT IN LOAD TRANSFER

What is the correct way to treat the unsprung components when analyzing or modeling cornering? It is customary to treat the unsprung masses as not acting through the springs or the linkages. But is that really true when the suspension is independent and the wheels lean with the car? What happens when there is some camber recovery in roll? What is the influence of roll center location on this? What is the correct way to treat the unsprung components for braking and acceleration?

This is a subject that I first received some correspondence about several years ago, and have been puzzling about ever since. I finally feel ready to present some conclusions.

First, a brief introduction to the subject under discussion. Dynamic load transfer (weight transfer) in response to ground-plane forces (forward, rearward, lateral) at the tire contact patches breaks down into four basic categories:

- **Elastic load transfer:** load transfer through the springs, anti-roll bars and any other springing devices
- **Geometric load transfer:** load transfer through support forces within the suspension, induced by linkage geometry
- **Frictional load transfer:** load transfer resulting from damper forces and friction in suspension pivots
- **Unsprung load transfer:** load transfer resulting from the overturning moments created by the inertia of the unsprung components, which does not act through the suspension.

In addition to these, there is also load transfer with steer, which results from caster jacking, and load transfer due to z or road-vertical forces, including vehicle accelerations from crests, dips, and turn banking, and aerodynamic lift/downforce. Load transfer with steer depends on steering geometry and front and rear angular warp stiffness and track widths. Load transfer with z force depends on relative wheel rates in roll, pitch, and heave, and on mass locations and center of pressure location relative to the contact patches. Neither of these depends in any direct way on suspension geometry.

Finally, with live axles, we have load transfer due to driveshaft torque.

To simplify discussion of the principles at issue, we will be ignoring effects of tire deflection, even though in some cases these can be significant.

What concerns us here is unsprung load transfer, and its relation to elastic load transfer.

To understand what unsprung load transfer is, imagine an axle with two wheels rolling down the road on its own, with no car sitting on top of it. The assembly has mass, centered at about hub or wheel center height. For it to go around a corner, the tires must exert a lateral or centripetal force, which they will do at ground level. The mass of the assembly will react with an equal and opposite inertial or centrifugal force, acting at hub height. The resulting couple will try to turn the assembly over, toward the outside of the turn. The tires will keep it from overturning, by exerting an increased force against the ground on the outside tire, and a correspondingly decreased force on the inside tire. This increase or equal decrease is the unsprung load transfer.

The axle does this just the same when it does have a car sitting on top of it.

It has been accepted practice to treat the unsprung masses in an independent suspension as if they were equivalent to an axle. However, many are now realizing that this isn't really correct, and that different parts of the unsprung masses can act as sprung masses in roll, and therefore a good analysis or model should take this into account.

There are also some interesting subleties to deciding whether some parts of the suspension are sprung or unsprung at all for a given type of analysis, which we will explore. We will also consider how this aspect relates to common methods of measuring unsprung weight in the garage or on a test rig.

Since we will be needing some equations, it is time to define some variables. Readers will recognize many of these from last issue. I am repeating those here, partly for the reader's convenience, and partly for the benefit of new subscribers.

There are also some additions, appropriate to our present discussion.

First, the vehicle axis system:

Per SAE convention, x is longitudinal, y is transverse, and z is normal or vertical. Unless otherwise stated, sign conventions are: x positive forward; y positive rightward; z positive downward. Unless otherwise stated, vehicle origin is per SAE aerodynamic axis system: on the ground plane, midway along the wheelbase, and centered with respect to the front and rear tracks. However, for particular purposes, we may use local origins and sign conventions may deviate, so it is necessary to pay attention to context. Even then, x will always be longitudinal, y transverse, and z vertical, in some sense that is appropriate to the context.

SAE vehicle axis convention takes as origin a point on the roll axis, directly below the sprung mass center of mass or c.g. I don't care for that, because the points in question move and are subject to uncertainties. Actually, there is no origin that is completely immune to some uncertainties and relative movements, but for our purposes here, we will use the SAE aerodynamic convention.

x, y, and z may denote coordinates in this axis system, or linear displacements in this axis system.

Next, angular quantities about the axes:

φ (phi) is roll, or angular movement about an x axis, conventionally positive rightward, or clockwise as seen looking forward.
M_x or M_φ is a moment in roll, or about an x axis.
K_φ is the elastic angular roll resistance rate. Depending on context, this may be for only a front or rear wheel pair, or for the entire vehicle.

It is also customary to use φ to denote camber of a wheel. This can create some confusion. Where we need to use this convention, φ_t is camber, and φ with no subscript is roll, unless otherwise indicated.

γ (gamma) is tire inclination. This is the same as camber except for sign convention. Camber is positive when the top of the tire tilts outboard with respect to the car. Inclination is positive rightward, or clockwise looking forward, for either wheel. By this convention, for a right wheel $\gamma = \varphi_t$, and for a left wheel $\gamma = -\varphi_t$. In less formal, more conversational usage, inclination may also be understood to be positive into the turn. Note also that both camber and inclination refer to the wheel's angle with respect to the road or ground plane or a perpendicular to it, not the body or sprung mass.

θ (theta) is pitch, or angular movement about a y axis, conventionally positive rearward, or clockwise when looking from left to right.
M_y or M_θ is a moment in pitch, or about a y axis.
K_θ is the elastic angular pitch resistance rate. It is analogous to K_φ, only in the pitch plane. Accordingly, it may refer to the rate for the whole vehicle, or for a right or left wheel pair, according to context.

θ_u is angular position or displacement of an upright or axle housing, e.g. caster or caster change. This is somewhat analogous to φ_t, only in side view. As with φ_t or γ, this is an angle or displacement with respect to the ground plane or a perpendicular to it, not the body or sprung mass.

ψ (psi) is yaw, or angular movement about a z axis, conventionally positive rightward, or clockwise when looking down.
M_z or M_ψ is a moment in yaw, or about a z axis.

L is a length. In our discussion here, it particularly refers to center spacing between two wheels of a pair under consideration.
L_x is wheelbase.
L_y is track.

r_t is tire loaded radius. For simplicity, we will ignore distinctions between loaded and effective tire radius.

F is force. F with an x, y, or z subscript is force in an x, y, or z direction.
m is mass.
a is acceleration.
F = ma. When m is expressed in pounds, a should be in g's.
Or, to be a bit more scientifically correct, weight W (pounds) is the force mg exerted on a mass m by a gravitational field in which a free-falling body accelerates at g (ft/sec^2). An acceleration in g's is a in ft/sec^2 divided by the number of ft/sec^2 per g. When we substitute mg, or W, for m, then F = ma becomes F = (mg)(a/g) = Wa/g if a is in ft/sec^2, or F = Wa if a is in g's.

Users of the metric system may be snickering at the use of pounds weight for mass, but weighing things in pounds, and getting accelerations from data acquisition systems in g's is the norm in English-unit vehicle engineering. But is weighing things in kilograms any more rational? Using kilograms for mass is correct, and a can then be in M/sec^2. However, one could similarly quibble that objects should really be weighed in newtons, since weight is force, not mass. a should then be in g's, or the weight properly taken in newtons should be divided by gravitational acceleration to obtain mass.

Is this confusing enough yet? Is it of only academic interest? I guarantee it will become hugely relevant, just as soon as we start racing on other planets.

Until then, it can at least be simple to apply F = ma for practical engineering, if we use these rules:
If m is in pounds, a has to be in g's. F will be in pounds.
If m is in slugs (pounds divided by 32.2), a has to be in ft/sec^2 (g's times 32.2). F will be in pounds.
If m is in kilograms, a has to be in M/sec^2 (g's times 9.8) F will be in newtons.
If m is in newtons (kg times 9.8), a has to be in g's. F will be in newtons.
Don't use pounds directly with ft/sec^2, or kilograms directly with g's – unless you want F in kgf.

m_S is sprung mass.
m_U is an unsprung mass.

H is height.
H_{cg} is height of the center of mass or center of gravity (c.g.). Unless otherwise indicated, this means the overall c.g. for the whole car.

H_{cgS} is the height of the sprung mass c.g.
H_{cgU} is height of an unsprung mass c.g. This is typically approximated as being equal to r_t, absent better information. Since we have more than one unsprung mass, further subscripts are applied per below.
H_{rc} is roll center height.
H_{pc} is pitch center height.

Subscripts applied to various quantities are as follows:
$_L$ is left.
$_R$ is right.
$_r$ is rear.
$_f$ is front.

For example:
m_{Sr} is the portion of the sprung mass statically supported by the rear wheels.
m_{SR} is the portion of the sprung mass statically supported by the right wheels.

dx/dz is the instantaneous rate of change of a point's x coordinate with respect to change in its z coordinate: the first derivative of x displacement with respect to z displacement.
dy/dz, similarly, is the first derivative of y displacement with respect to z displacement.

dx/dz and dy/dz are also the instantaneous slopes or inclinations, from vertical, of the contact patch center's path of motion in side and front view respectively.

A force line is a notional line of action for the vector sum of an x or y ground-plane force at a tire contact patch and the induced z-direction support force within the suspension system that results from angularities in the suspension linkage. In front-view geometry, the force line is the line from the contact patch center to the front-view instant center. The front-view force line is an instantaneous perpendicular to the contact patch center's path of motion in front view, and has a slope or dz/dy, relative to ground plane, equal to the contact patch center's dy/dz. The side view force line is an analogous construction in side view, whose slope is the contact patch center's dx/dz.

$d\gamma/d\varphi$ is the rate of tire inclination change with respect to angular roll. If $d\gamma/d\varphi = 1$, the wheels lean the same amount as the sprung mass, as with parallel control arms or pure trailing arms. If $d\gamma/d\varphi = 0$, the wheels don't lean at all with roll, as with a beam axle. We can say that $(1 - d\gamma/d\varphi)*100\%$ is our percent camber recovery, or that $d\gamma/d\varphi$ is the camber non-recovery, expressed as a decimal. (I suppose if we really want to be fastidious, we should speak of inclination recovery – only that isn't common usage.)

Analogously, $d\theta_u/d\theta$ is the rate of caster or side-view inclination change of an upright or axle, with respect to angular pitch. If $d\theta_u/d\theta = 0$, that means that θ_U does not change in pitch. This requires a side-view swing arm length x_{svsa} equal to half the wheelbase. It also means that $d\theta_u/dz = -d\theta/dz$. That is, the rate of caster change with respect to linear suspension displacement in the absence of pitch is equal in magnitude and opposite in direction to the rate of pitch displacement with respect to magnitude of linear displacement per wheel. If $d\theta_u/d\theta = 1$, that means that θ_u changes in pitch identically with sprung mass angular pitch. This requires parallel side view projected control arms, an x_{svsa} that is infinite or undefined, and $d\theta_u/dz = 0$ in the absence of pitch. $d\theta_u/d\theta$ can be considered to be the rate of caster non-recovery, expressed as a decimal.

r_{cgSx} is the moment arm of the sprung mass about the roll axis. That is, it is the side-view perpendicular distance from the sprung mass c.g. to the roll axis.

r_{cgSy}, correspondingly, is the moment arm of the sprung mass c.g. about the pitch axis. That is, it is the front-view perpendicular distance from the sprung mass c.g. to the pitch axis.

Mathematically, we can compute r_{cgSx} and r_{cgSy} as follows:

$$r_{cgSx} = H_{cgS} - \left(H_{rcf} + (H_{rcr} - H_{rcf})(m_{Sr}/m_S)\right) \left(\sqrt{\left(L_x^{2} + (H_{rcr} - H_{rcf})^2\right)}\right) / L_x \qquad (1a)$$

$$r_{cgSy} = H_{cgS} - \left(H_{pcR} + (H_{pcL} - H_{pcR})(m_{SL}/m_S)\right) \left(\sqrt{\left(L_y^{2} + (H_{pcL} - H_{pcR})^2\right)}\right) / L_y \qquad (1b)$$

Equation (1b) applies for a case where front and rear track are equal. Where we have unequal track front and rear L_{yf} and L_{yr}, for best accuracy we need to substitute for L_y a weighted average of L_{yf} and L_{yr}, which will be $\left(L_{yf} - (L_{yf} - L_{yr})(m_{Sr}/m_S)\right)$. A similar correction could even be required for unequal right and left wheelbases, but ordinarily the wheelbases are identical, or differ by so little that a simple average will suffice.

These equations can be simplified in certain cases. If the front and rear roll centers are similar height, there is little difference between the perpendicular distance from sprung mass c.g. to roll axis and the vertical distance. Similar reasoning applies regarding the moment arm in pitch, if the pitch centers are similar heights at the right and left. For a simple case where front and rear track are equal, the sprung mass c.g. is laterally centered, and the roll and pitch axes have little slope, the equations simplify to:

$$r_{cgSx} = H_{cgS} - \big(H_{rcf} + (H_{rcr} - H_{rcf})(m_{Sr}/m_S)\big) \tag{1c}$$

$$r_{cgSy} = H_{cgS} - (H_{pcL} + H_{pcR})/2 \tag{1d}$$

We use r_{cgSx} to calculate the overall elastic roll moment $M_{\varphi E}$ – the moment resisted by the springs, anti-roll bars, and any other springing devices that resist roll, for a given lateral acceleration a_y.

$$M_{\varphi E} = m_S\, a_y\, r_{cgSx} \tag{2}$$

Knowing the elastic angular roll resistance rates for front and rear wheel pairs $K_{\varphi f}$ and $K_{\varphi r}$, we then calculate the angular roll displacement φ.

$$\varphi = M_{\varphi E} / (K_{\varphi f} + K_{\varphi r}) \tag{3}$$

We now know the angular roll displacement, and we know the angular rate at which the front and rear suspensions resist roll, so we can calculate how much elastic roll resisting moment we have at each end of the car.

$$M_{\varphi Ef} = \varphi\, K_{\varphi f} \tag{4a}$$

$$M_{\varphi Er} = \varphi\, K_{\varphi r} \tag{4b}$$

Finally, knowing the front and rear track widths L_{yf} and L_{yr}, we can calculate the front and rear elastic load transfers ΔF_{zEf} and ΔF_{zEr}. (Note that ΔF_z is the load change at one wheel. The resulting change in load difference between the two wheels is $2\Delta F_z$.)

$$\Delta F_{zEf} = M_{\varphi Ef} / L_{yf} \tag{5a}$$

$$\Delta F_{zEr} = M_{\varphi Er} / L_{yr} \tag{5b}$$

Next, we calculate the geometric load transfer at each end, ΔF_{zGf} and ΔF_{zGr}.

$$\Delta F_{zGf} = (m_{Sf}\, a_y\, H_{rcf}) / L_{yf} \tag{6a}$$

$$\Delta F_{zGr} = (m_{Sr}\, a_y\, H_{rcr}) / L_{yr} \tag{6b}$$

Frictional load transfer is difficult to calculate with good accuracy. Ordinarily, suspension bearings are treated as frictionless – not because we really think they are, but because we cannot reliably calculate their frictional loads at a given instant. For steady-state analysis, we assume that the suspension systems are at zero velocity: the car has assumed a steady attitude, in response to steady ground-plane forces. When the suspension does have some velocity, the dampers are making forces. These are customarily modeled as being in a consistent relationship to shaft velocity, often a linear relationship, although behavior of actual dampers is considerably more complex. At any rate, for steady-state cornering, or combined cornering and longitudinal acceleration with unchanging a_x and a_y, frictional load transfer is generally considered to be zero.

We now have all the components of load transfer due to ground-plane forces, except the unsprung load transfer.

So far, I have merely been explaining conventional theory. Now we are going to break some new ground. Traditionally, the total unsprung mass at the front and rear has been taken as equivalent to an axle assembly, even where the actual suspension is independent, and the front and rear unsprung load transfer ΔF_{zUf} and ΔF_{zUr} have been calculated as follows:

$$\Delta F_{zUf} = (m_{Uf}\ a_y\ H_{cgUf}) / L_{yf} \qquad (7a)$$

$$\Delta F_{zUr} = (m_{Ur}\ a_y\ H_{cgUr}) / L_{yr} \qquad (7b)$$

In my opinion, this is correct for any beam axle, and also for any independent suspension with 100% camber recovery in roll. However, for most independent suspensions, it is incorrect. The equations really should be:

$$\Delta F_{zUf} = \Big((m_{URf}\ a_y\ H_{cgURf})\big(1 - (d\gamma/d\varphi)_{Rf}\big) + (m_{ULf}\ a_y\ H_{cgULf})\big(1 - (d\gamma/d\varphi)_{Lf}\big)\Big) / L_{yf} \qquad (7c)$$

$$\Delta F_{zUr} = \Big((m_{URr}\ a_y\ H_{cgURr})\big(1 - (d\gamma/d\varphi)_{Rr}\big) + (m_{ULr}\ a_y\ H_{cgULr})\big(1 - (d\gamma/d\varphi)_{Lr}\big)\Big) / L_{yr} \qquad (7d)$$

It will be apparent that when $d\gamma/d\varphi = 0$ for both wheels, and the unsprung mass c.g. height is the same right and left, equations (7c) and (7d) simplify into equations (7a) and (7b).

It will also be apparent that when $d\gamma/d\varphi > 0$, there is now a component of load transfer that is unaccounted for in the equations for the unsprung load transfer. No matter what the value of $d\gamma/d\varphi$ is, the components still have the same mass, and the same c.g. height, so at a given a_y the lateral load transfer from those masses must be a constant regardless of $d\gamma/d\varphi$. So what happens to the remaining component?

It tries to roll the car on its suspension, and thus becomes a part of the elastic load transfer! When camber recovery is less than 100%, it really is possible to create roll in the suspension by exerting a y force near the hub of a wheel, resisted by friction at that wheel's contact patch, without applying any other force to the sprung structure.

We therefore have to revise equation (2), $M_{\varphi E} = m_S a_y r_{cgSx}$, to include the elastically reacted component from the unsprung masses, as follows:

$$M_{\varphi E} = (m_S a_y r_{cgSx}) + \big(m_{URf} a_y H_{cgURf} (d\gamma/d\varphi)_{Rf}\big) + \big(m_{ULf} a_y H_{cgULf} (d\gamma/d\varphi)_{Lf}\big) + \big(m_{URr} a_y H_{cgURr} (d\gamma/d\varphi)_{Rr}\big) + \big(m_{ULr} a_y H_{cgULr} (d\gamma/d\varphi)_{Lr}\big) \tag{8a}$$

If we wish, we may also write that this way:

$$M_{\varphi E} = a_y \Big((m_S r_{cgSx}) + \big(m_{URf} H_{cgURf} (d\gamma/d\varphi)_{Rf}\big) + \big(m_{ULf} H_{cgULf} (d\gamma/d\varphi)_{Lf}\big) + \big(m_{URr} H_{cgURr} (d\gamma/d\varphi)_{Rr}\big) + \big(m_{ULr} H_{cgULr} (d\gamma/d\varphi)_{Lr}\big)\Big) \tag{8b}$$

Or:

$$M_{\varphi E} = a_y \Big((m_S r_{cgSx}) + \sum\big(m_U H_{cgU} (d\gamma/d\varphi)\big)\Big) \tag{8c}$$

Note that we do not simply take an effective mass $\big(m_U (d\gamma/d\varphi)\big)$ at height H_{cgU} and add it to the sprung mass m_S. Nor do we add such a mass to m_S when calculating the geometric load transfer ΔF_{zG}.

One of the correspondents who helped set me to thinking about this whole business suggested that what mattered was the lateral movement of an unsprung mass's c.g. with respect to the origin, or to the sprung mass, as the car rolls. That would mean that both the x and z coordinates of the front view instant center would matter. That is, both camber change and track change with suspension motion would matter. After considerable thought, I have concluded that what really matters, for an unsprung mass or a sprung mass, is the rate of lateral motion of the mass's c.g. with respect to the relevant contact patch or patches, per unit of roll.

For an individual wheel, the relevant contact patch is the wheel's own. For the sprung mass, it's a weighted average of the contact patches – weighted according to the right/left distribution of cornering force contribution. This is what we are getting to when we assign roll centers and roll axes the correct way, which I have described in my video and in other documents.

That is why we take a moment about the roll axis for the sprung mass, but take individual moments about a wheel's own contact patch, with a camber non-recovery multiplier, for the unsprung masses, and then sum these get a correct $M_{\varphi E}$.

The whole thing can be thought of as an application of force/motion relationships. The more a particular mass's c.g. moves laterally in response to its own centrifugal inertia, with respect to the relevant contact patch(es) where the centripetal acceleration force is applied, to roll the suspension a degree, the more roll a pound of force applied at that point will produce. This is directly analogous to a lever working against an elastic resistance. A longer lever produces more angular deflection for a given force at the lever end, in direct proportion to how far that lever end moves per degree of deflection. (For large angular displacements, the relationships start to get non-linear, but for the small angles we normally deal with in analyzing roll and pitch of automobiles, the non-linearities are very small.)

What components are we talking about when we refer to unsprung components? One might think this would be a simple question, but when we get down to the fine points, it isn't quite so simple.

For the parts with the most mass – the tire, wheel, upright, brake rotor, and caliper – it's easy: they can be treated as entirely unsprung, with a single center of mass as an assembly. The whole assembly has the same value for $d\varphi_t/d\varphi$, and the above discussion and equations apply.

But what about the spring? The unsprung portion of the damper? The control arms or links? The pushrod or pullrod? The rocker? An anti-roll bar or torsion bar arm? If we take out the spring, support the sprung structure on stands or on a test rig, and note the weight of the unsprung components at a wheel scale or the wheel support pad on a rig, how accurate is that? Do we add half the weight of the spring?

In my opinion, using weight as measured above, without adding half the spring, will in most cases yield a very good figure for use in the calculations we are considering here. It will be less accurate if we are seeking a value for use in ride analysis, or modeling of unsprung mass response to road perturbations. For that, we need to know the effective inertia mass for the components, in wheel z acceleration.

For example, suppose we have a coil spring of total mass in pounds m_{spring}, in a typical big-spring stock car front end, acting at a motion ratio of 0.50. Half of m_{spring} acting at a wheel scale through that 0.50 motion ratio would increase the scale reading by $m_{spring}/4$. However, when the wheel hits a bump, and sees an acceleration in g's a_{zt}, half of the spring sees an acceleration of $a_{zt}/2$. The inertia force at the spring is then $a_{zt}\ m_{spring}/4$. That inertia force at the wheel is $a_{zt}\ m_{spring}/4$, times the 0.50 motion ratio, or $a_{zt}\ m_{spring}/8$. The scale weight contribution of half the spring's mass is that mass times the motion ratio, but the inertia contribution is that $m_{spring}/2$ mass times the square of the motion ratio. In this example, an ounce on the spring only has 1/8 as much effect on the wheel's ability to ride bumps as an ounce on the wheel itself.

In roll, how does the spring act? It pretty much acts as part of the sprung mass, and can be treated as such. That is, we don't need to add any part of it to our unsprung mass for purposes of roll analysis.

What if we have a formula car with pushrod suspension, and a coilover lying horizontal, operated through a rocker? Here again, parts of this mechanism undergo accelerations when the wheel sees an a_{zt}, but they pretty much move with the sprung mass in roll, not with the tire and upright. No part of a horizontal coilover adds gravitational force at the contact patch, but the parts that move when the suspension moves definitely contribute to unsprung mass inertia when the wheel hits a bump.

Or, consider the case of an upper control arm that serves as a rocker operating an upright inboard coilover, as was popular for open-wheel cars in the '60's and '70's. Here, the moving bits inboard of the fulcrum appear to have negative weight when measuring on a wheel scale. If we hang enough lead on the inboard end of the arm, we can get a reading of zero at the wheel scale. That doesn't mean we are reducing unsprung mass or unsprung inertia, or improving the wheel's ability to ride bumps.

It is sometimes said that one measurement is worth a thousand expert opinions. There is some truth to this, but only if the measurement in question is an appropriate one for the phenomena under consideration. For most dynamic analysis, we are concerned with inertia, not weight, and a given object exerts different amounts of inertia force in different modes of movement.

To get a valid experimental measurement of unsprung mass inertia in response to an a_{zt}, we need to subject the wheel to an actual vertical acceleration of known magnitude, with a dummy shock to eliminate damper forces, and measure the resulting inertia force. To do that, we need a seven-post rig, and the acceleration needs to be measured in terms of wheel or displacement pot motion, not ram or contact patch motion, despite the fact that we will necessarily be measuring the resulting force at the ram or contact patch. Chances are that we can get at least equal accuracy by weighing the parts individually on scales, and applying some expert opinion in the form of sound mathematics.

Depending on the design of the suspension, the various inboard components may not exactly move as part of the sprung mass in roll. Fortunately, these components are generally small in mass compared to the rest of the system, so any error resulting from simply treating them as part of the sprung mass will be correspondingly small. Therefore, I suggest treating them that way as a general-purpose approximation.

Now, how does all of this apply for longitudinal accelerations (a_x)? Rather similarly, except that the wheel is round in side view, and only generates a ground-plane force when a torque is applied to it. Also, the wheel is spinning in side view when the car is moving, and it undergoes rotational accelerations a_{yt} about its axis of rotation, which create inertial moments M_{yt} in addition to the linear inertia of the wheel. These are relatively small, and it is common to ignore them, but they are real.

We will consider the linear inertia effects first. We know that the unsprung components will react to any a_x with an inertia force $F_{xU} = m_U a_x$.

If we have a car on a kinematics and compliance (k&c) rig, and we exert, say, a forward force directly upon a wheel at about hub height, what does the car do? Does it pitch? Unless we lock the wheel somehow, the car doesn't pitch; it just moves forward. To resist the force, we have to apply the brakes or otherwise lock the wheel. The way the forces react once the wheel is locked depends on the load path for the braking or locking torque, as it does when determining longitudinal anti and pitch center height. However, as with the unsprung component in roll, we are concerned with the unsprung mass's motions relative to is own contact patch center rather than a weighted average of two contact patches.

When the torque reacts through the linkage, the situation is directly analogous to roll. The portion of the unsprung mass that creates a pitch moment in the suspension is proportional to the caster non-recovery $d\theta_u/d\theta$. When $d\theta_u/d\theta = 0$, or when the torque does not react through the linkage, as with inboard brakes, the unsprung mass acts entirely as unsprung, and creates no suspension pitch. The equations that result are of the same form as those for roll, as follows:

When $d\theta_u/d\theta = 0$, or when torque does not react through the linkage,

$$\Delta F_{zUR} = (m_{UR}\, a_x\, H_{cgUR}) / L_{xR} \tag{9a}$$

$$\Delta F_{zUL} = (m_{UL}\, a_x\, H_{cgUL}) / L_{xL} \tag{9b}$$

Otherwise,

$$\Delta F_{zUR} = \Big((m_{URf}\, a_x\, H_{cgURf})\big(1 - (d\theta_u/d\theta)_{Rf}\big) + (m_{URr}\, a_x\, H_{cgURr})\big(1 - (d\theta_u/d\theta)_{Rr}\big)\Big) / L_{xR} \tag{9c}$$

$$\Delta F_{zUL} = \Big((m_{ULf}\, a_x\, H_{cgULf})\big(1 - (d\theta_u/d\theta)_{Lf}\big) + (m_{ULr}\, a_x\, H_{cgULr})\big(1 - (d\theta_u/d\theta)_{Lr}\big)\Big) / L_{xL} \tag{9d}$$

And when a component of the unspung mass acts to create a suspension pitch displacement,

$$M_{\theta E} = (m_S\, a_x\, r_{cgSy}) + \big(m_{URf}\, a_x\, H_{cgURf}\, (d\theta_u/d\theta)_{Rf}\big) + \big(m_{ULf}\, a_x\, H_{cgULf}\, (d\theta_u/d\theta)_{Lf}\big) + \big(m_{URr}\, a_x\, H_{cgURr}\, (d\theta_u/d\theta)_{Rr}\big) + \big(m_{ULr}\, a_x\, H_{cgULr}\, (d\theta_u/d\theta)_{Lr}\big) \tag{10a}$$

Which we may also write:

$$M_{\theta E} = a_x \Big((m_S\, r_{cgSy}) + \big(m_{URf}\, H_{cgURf}\, (d\theta_u/d\theta)_{Rf}\big) + \big(m_{ULf}\, H_{cgULf}\, (d\theta_u/d\theta)_{Lf}\big) + \big(m_{URr}\, H_{cgURr}\, (d\theta_u/d\theta)_{Rr}\big) + \big(m_{ULr}\, H_{cgULr}\, (d\theta_u/d\theta)_{Lr}\big)\Big) \quad (10b)$$

Or:

$$M_{\theta E} = a_x \Big((m_S\, r_{cgSy}) + \sum\big(m_U\, H_{cgU}\, (d\theta_u/d\theta)\big)\Big) \quad (10c)$$

It is possible to imagine cases where $d\gamma/d\varphi < 0$ or where $d\theta_u/d\theta < 0$ – that is, where camber or caster recovery is greater than 100%. In such cases, the equations still apply, and the unsprung components actually create an anti-roll or anti-pitch moment in the suspension, reduce elastic load transfer, and create unsprung load transfer as if they had more than their own actual mass. Such cases are rare, however.

Finally, the rotating components have rotational inertia. As mentioned earlier, these forces are smaller than the linear unsprung inertias, and it is common to ignore them completely. They do not add to the unsprung load transfer. However, they can produce some forces within the suspension linkage, and they present a rather interesting analytical challenge.

The rotating portions of the unsprung assembly – typically the tire, wheel, hub, brake rotor, and driveshaft – can be thought of as a single mass m_r, acting at a radius of gyration k_y about the wheel axis. At a given vehicle longitudinal acceleration a_x, the inertial moment M_{yui} about the wheel axis is equal to m_r times the square of the ratio of radius of gyration to tire radius:

$$M_{yui} = m_r\, a_x\, (k_y/r_t)^2 \quad (11)$$

One way to think about this is to compare the effect of an ounce at the tire tread surface to an ounce halfway in toward the hub. The ounce at the tread surface accelerates circumferentially at a_x, and acts at r_t, so its inertial contribution rotationally is equal to its linear contribution. An ounce saved there is worth two off the frame, as it were. The ounce halfway in accelerates circumferentially at only $a_x/2$, and its circumferential inertia force acts at half the tire radius. Thus its inertial contribution rotationally is only ¼ of its linear contribution, and an ounce saved there is only worth 1.25 off the frame.

Probably the best way to find k_y is to computer-model the various parts. Failing that, a better-than-nothing approximation would be to assume $k_y/r_t = 0.5$ where the assembly includes a brake rotor, or $k_y/r_t = 0.7$ where the brake rotor is not included.

When looking at anti-pitch geometry last issue, we noted that the big forces acting through the linkage can include both a thrust and a torque (outboard brake) or only a thrust (inboard brake). With the much smaller forces from the rotating unsprung components, we have some different possibilities.

In an ordinary car with four-wheel outboard brakes, each brake has to exert an extra increment of torque to overcome rotational inertia at its own wheel. This torque does not show up as a longitudinal force at the contact patch, but it does react through the linkage: a torque without a thrust.

If the additional torque is generated through an inboard brake or transaxle, and transmitted through a jointed shaft to the wheel, the additional torque does not act through the linkage: neither a torque nor a thrust through the linkage.

Finally, we have the unique case of a wheel being neither driven nor braked, such as a front wheel of a rear-drive car under power. In this case, the torque to overcome rotational inertia comes to the front wheel through its contact patch. The creates a small rearward inertial thrust, and this does act through the linkage. Moreover, that thrust ultimately comes from the rear wheel, as an equal forward thrust at the rear contact patch. That reacts through the rear suspension linkage, and the additional increment of torque may too, if the rear suspension is such that drive torque does that.

In the first case, the torque to overcome rotational inertia creates a pitch moment $M_{\theta Eui}$ which must be reacted by the suspension's elastic components that resist pitch, whose magnitude depends on the caster non-recovery $d\theta_u/d\theta$, and these on the front and rear at each side create an anti-pitch moment as follows:

$$M_{\theta Eui} = M_{yuif}\,(d\theta_{uf}/d\theta) + M_{yuir}\,(d\theta_{ur}/d\theta) \qquad (12)$$

The only way this small pitch moment can affect wheel loads is if it is different on the right and left sides of the car. However, it may slightly affect pitch displacement even if it is the same on both sides.

In the second case, the rotational inertias do not affect wheel loads or suspension displacements, but they do affect brake requirements, or power requirements in the case of all-wheel drive.

In the third case, the thrust F_{xuif} at the undriven wheel creates an F_z that depends on the dx/dz at the wheel center, or at the contact patch of a free-rolling wheel, which is an identical dx/dz.

$$F_{zuif} = F_{xuif}\,(dx/dz)_f = M_{yuif}\,(dx/dz)_f\,/\,r_t \qquad (13)$$

The rear wheel F_x is equal in magnitude to the front one but opposite in direction and hence in sign.

$$F_{xuir} = -\,F_{xuif} \qquad (14)$$

There is then a F_{zuir} induced in the rear suspension which depends on the dx/dz that applies for the type of rear suspension the car has. There is simply a slightly greater propulsion force at the rear contact patches than would be otherwise be needed to produce a_x, and that produces forces in the rear suspension in accordance with its anti-squat properties.

If the rear suspension has anti-squat, and the front wheel has thrust anti-dive (that is, if its hub moves forward when the suspension compresses), then there are small upward forces induced in both the front and rear suspensions, and the resulting jacking force for the pair is the sum of the two, F_{zuir} + F_{zuif}. The difference of the two, times half the wheelbase, is the pitch moment resulting from the rotational inertias' reactions through the linkages.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Please note that this is a new e-mail address. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

CHASING A PUSH

This month we are departing a bit from the usual format of a single answer to a single question, and presenting something more closely resembling a series of messages between me and a client. Ordinarily, such exchanges occur in confidence, but this client consented in advance to publication. I have edited and added to the material slightly for publication. Client's words are in italics.

I work with a road-racing RX-8, and recently we have been experiencing high front tire wear. It was never previously an issue, but in preparation for a recent race, a softer front anti-roll bar was fitted to the car and a bit of rake was added. Our front tire wear was measurably poorer at that event. It is my suspicion that the softer front bar resulted in less diagonal weight transfer to the rear and consequently, more front tire wear. Do you have any thoughts on the subject?

Anti-roll bars do not transfer wheel load diagonally in the sense of transferring it from one wheel of a diagonal pair to the other, and they do not transfer it front to rear. We might say they transfer wheel load diagonally in the sense that they transfer it from one diagonal pair to the other diagonal pair. Putting it another way, they change the dynamic diagonal percentage, but not the diagonal front, rear, right, or left percentage.

The RX-8 has double wishbone suspension in front, unlike all RX-7's, correct?

A couple of questions back to you: did you increase the rake by raising the rear, or by lowering the front? How do your static camber and toe settings compare, before and after the change? Where do the front tires wear, i.e. what part of the tread?

Ordinarily, a softer front bar, or a stiffer rear, helps front grip and front tire wear. Does the car in fact have less understeer?

The chassis that we are running indeed has a double wishbone front suspension. When we increased the rake, we added height in the rear; we try to run the front as low as possible at all times. We were sure to set the camber and toe to the same setting as before our changes. We are usually able to achieve very high setup accuracy. The tire wear was strikingly even across its full width. As far as the car's balance, we have been experiencing mid-corner understeer for the entire life of the car, and that did not go away with these changes. It didn't get any worse either though.

The fact that the understeer wasn't affected by the a/r bar change, together with the fact that it is most evident mid-turn, suggests that something in your front end is running out of travel at full roll. The coilover or shock could be bottoming, or it could be something else. Other things that can happen in lowered cars, particularly when you've found a bit more bump travel at the shock, include ball joints (usually upper) running out of travel, and control arms (usually lower) hitting something, like the frame. Sometimes tie rods or a/r bars hit things.

Do you have travel indicators on your shocks? If so, and if they don't indicate that the shock is bottoming, take out the springs, disconnect one end of the a/r bar, and move each front suspension through its indicated travel with a jack, then a bit beyond that. See if anything binds or hits. See if the steering still works freely.

If everything checks out at full compression, try the same thing, for full droop. In lowered cars, it is less common for the inside wheel to top out than for something to bottom on the outside one, but this can happen in cases where a shorter shock has been used to get more bump travel.

There is one other possibility, particularly if you are satisfied that nothing is binding or bottoming.

In some cases, mid-turn push can be caused by the brakes failing to release as they should, and dragging for a time following trailbraking. This can even be driver-induced, when a left-foot braker unwittingly fails to release the pedal completely. If you can rule out the driver-induced situation, it can be a bit hard to tell if the brakes are releasing slowly. One quick test is to try hard cornering, not immediately preceded by braking. That is, either test on a skidpad, or drive into a turn fast enough to put the car at the limit going through the turn, but approach the turn at that speed so you don't have to brake. If the car still pushes, you can probably rule out the brakes. If the push does not appear when you don't brake immediately before, suspect dragging brakes.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2009

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

LEADING AND TRAILING WHEELS FOR OVAL TRACK

My question relates to the wheelbase on a 2700 lb. Outlaw bodied Latemodel. We have no rules on the wheelbase difference from side to side. My question is, any benefit to altering the wheelbase from side to side? I was thinking of shortening the left side by toeing out the LR and toeing in the RR, then squaring the rear ties to the chassis to create lead in the lr tire. In the front I was thinking of moving the LF suspension back. This would create a shorter wheelbase on the left but have the tires still pointing straight ahead. Is there any benefit to this approach and how much could you go?

It comes down to whether this helps you work whatever rules you do have.

What is the wording of your engine setback rule?

What rule(s) do you have regarding wheelbase?

What rule(s) do you have regarding rear axle housings?

Engine setback is 4in. measured from #1 plug to center of upper balljoint.
There are no rules stated for wheelbase.
There are no rules stated for rear axle housings.

This is pretty much an open competition class. We run a 9 in. treaded tire so corner speed and momentum is the key to going fast and passing. I am just trying to think outside the box a little.

No wheelbase rule at all? Really?

Is there a minimum right side weight or a maximum left percentage?

How much ballast is in the car, and where is it?

What are your current static left and rear percentages?

Are you mostly running bowl-shaped momentum tracks, with large turns, short straights, and narrow speed range, or paperclip-shaped tracks, with tight turns, pronounced straights, and wide speed range? How important is braking? Do you use the brakes hard, or hardly use them?

No wheelbase rule at all.

No leftside weight rule - my car is about 62 % left and 49.5 rear.

Car weighs about 2100 lbs no ballast no driver. Rules 2700 lbs. We use about 400 lbs. of ballast mostly on left side.

The track in question is a 3/8 high bank bowl shaped momentum track with narrow speed range most definitely. We use very little brake if any on this track.

If the track is one where steady-state cornering is the main requirement, you're probably about where you should be on rear percentage. On a more paperclip-shaped track, I'd suggest more rear, to aid propulsion and braking.

It sounds like if you want to, you can move some ballast rearward and get more rear percentage, without needing to move anything else. Is that correct?

62% left is a lot, but if you can get even more, I bet you'd go faster. You want to start being careful doing sudden darts to the right, but turning left I bet you'd gain speed all the way up to 70% left, if you could get to that.

Have you tried more rear? More left?

Can you get the ballast any lower than it is?

Returning to the original question, I don't think leading the right front and left rear wheel, or leading both right wheels either, makes the car turn left more readily, if that was what you were wondering about. The car is very sensitive to the aim of the rear wheels, but if you change their location slightly without changing the direction they point, that shouldn't make much difference.

Relocating wheels does potentially give you a way to get a little more rear percentage, while meeting the engine setback rule, depending on how they measure. If they lay a square along the side of the head and let its leg go out to the ball joint, perhaps extending that line using a string or another straightedge, moving the right front wheel forward doesn't change the measurement, yet it gets you some rear percentage. Ditto moving either rear wheel forward, or both of them. Moving both rights forward makes sense when you have a wheelbase rule, and a wheelbase difference rule, and you're after more rear percentage. In your case, you don't have those requirements, so I doubt that you'd see an advantage. I doubt that it would hurt, either, nor would moving the right rear back and the left rear forward. It would mostly just be extra work for little or no effect once you dial the rest of the setup back in to suit the change.

One other question: if there's no axle housing rule, does that mean you can run cambered rear axles? If so, maybe a bit more rear percentage might even help you mid-turn.

There are reliability and efficiency issues as you try to angle the axle snouts with respect to the tubes. Beyond about 1 degree, you need to run the axles with the barrel-shaped splines if you want decent life and reliability. Even then, you are turning some horsepower into heat as the wheels turn and the splines rub.

Depending on the properties of your tires, you should see an improvement in cornering force up to at least a degree and a half of inclination into the turn (positive camber at the left rear, negative on the right rear). My expectation would be that by the time you get the camber optimized, the splines will be rubbing quite a bit as the wheels turn, and you probably don't want to make that worse by adding top-view angle to the snouts.

In other words, if you're going to angle the snouts with respect to the tubes, go for camber rather than toe, unless you are thoroughly satisfied that you've already optimized camber.

You may find that as you increase tire inclination, your optimum pressure will be less. That is, cambered tires will usually want a bit less air. You want to find the camber that works best at its own optimum pressure, rather than at the pressure that worked best without camber.

It sometimes happens that the camber and pressure that are fastest will blister tires. This definitely happens in the upper NASCAR divisions. If you encounter that problem, you may have to be more conservative on your settings to make sure the tires survive.

If you run a spec, treaded tire, you may find it desirable to shave them. This is hardly a secret, and you may be doing this already. If you do shave the tires, and you can run cambered axles, I suggest shaving the tires roughly straight across, rather than camber cutting them. Camber cutting makes sense when you can't run cambered axles, but when you can adjust camber, it's better to get the tread as shallow as possible all the way across.

Stagger will affect camber somewhat, so when optimizing camber you want to be working with an optimal amount of stagger. There are ways to calculate that from a measurement of turn radius and banking, but you can also find it experimentally. Drive into the turn on your racing line, at a bit below racing speed, disengage the clutch, and coast through the turn. If you run a spool, see how much speed you lose. If you have a locker, turn off the engine and see if you hear it click. The idea is to find the amount of stagger that gets you through the turn under power with the least drag.

This will be worth some cornering speed, even if the car is not power-limited when cornering. You can get the car to behave decently with some other amount of stagger, but then the tires are using up some of their grip budget overcoming the incorrect stagger: they're fighting each other to some extent. The grip they use doing that is grip they can't use to make total lateral force.

You can also use a similar technique to optimize toe, both front and rear. Again, you want the setting that drops your speed the least when coasting through the turn at slightly below racing speed. You may find that at the best setting, the car feels a bit nervous or wanders a bit, although generally that is more of a problem on straights than in turns, with least-drag toe. You do want to be comfortable with the car at the limit, but the idea is to run the least draggy setup that satisfies that requirement.

Note that this recommendation applies only to tire drag, not necessarily aerodynamic drag. You don't actually want aerodynamic drag, especially if you do little braking, but for almost any short-track application, you need maximum downforce from the body, even if it costs you some drag.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

USE RIDE HEIGHT TO TUNE CORNERING BALANCE, OR ANTI-ROLL BARS?

I have a FF1600 (i.e. no ground effects) and the person who runs the car and I differ on how it should be set up:

1. *He first gets the handling roughly right by raising the rear ride height and then fine tunes the handling with the rear ARB. This assumes the ARB is already of the required size,*
2. *I think the car should be kept flat so the CofG is as low as possible and only use the ARB*
3. *He says that having the car flat will require an harder rear ARB (which is true) and this in a turn lifts the rear inside wheel more*
4. *My response is that if the handling is the same the weight transfer must be the same. This means the rear roll stiffness is the same so the inside wheel will be lifted by the same amount.*

So who's correct? Point 4 isn't quite true because a lower CofG will affect the handling but the general principle still applies. I could try this out in the workshop but this would be a lot of unnecessary effort if the answer is already well known.

Basically, the questioner is correct here. The tires don't know if the load transfer comes from geometry or springs or anti-roll bars.

Raising the rear ride height affects handling mainly by raising the rear roll center. This adds load transfer, just like stiffening the rear bar does.

In general, the car will go fastest with both ends as low as possible. How low we can go will generally be determined by the need to avoid bottoming, with perhaps some additional influence from aerodynamic considerations. Both of these factors may call for slightly more ground clearance at the rear than at the front.

It is common for race cars, especially rear-engined ones, to have lower natural frequencies in ride in the rear than in the front. That will mean the rear suspension deflects more at the bottom of a dip than the front suspension.

To get good air flow under the car, it generally helps to have the under-car space open up a bit toward the rear. This will generally reduce both lift and drag.

But the idea is still to run both ends of the car as low as these constraints will allow. It is not a good idea to balance the handling by running either end higher than necessary.

ROLL CENTER AND ANTI-LIFT IN TRIANGULATED 4-LINK BEAM AXLE

I am in the process of converting the rear of a front wheel drive rallycross car from its present format to a solid beam axle, primarily to save weight. Tthe current set up actually works very well, but by removing both the rear chassis rails and floor and replacing these with a light weight semi space frame setup clad in either aluminium or a carbon /Kevlar sheet we can achieve a considerable weight saving, and by adopting the triangulated 4 link set up direct the loads favorably into the car.

My question is twofold: if we adopt a set up similar to the Satchell link with the angled links at the bottom of the axle tube and the tubes parallel to the car's centreline at the top, where will the roll centre be? Also, if from necessity from a structural standpoint the angled links slope upwards in respect to the ground will this affect anything except braking squat or anti squat on this un-driven axle?

Secondly, how do we ensure that with a fully rose jointed set up there is no bind over the full suspension travel? My mockups seem to indicate that some compliance may be necessary.

To find the roll center, find the point where the centerlines of the two diagonal lower links intersect. In a symmetrical geometry, there will be an intersection somewhere aft of the axle. If the linkage is not symmetrical, we approximate by finding where the two lines cross in top view, and averaging their heights (z coordinates) at that (x,y) location. We then construct a line containing that point, parallel to car centerline in top view and parallel to the parallel upper links in side view. Where that line intercepts the axle plane (vertical transverse or y-z plane containing the axle line), that's the roll center.

If the upper links are horizontal, and the lower links slope up toward the front, the car will have some anti-lift in braking. If the lowers slope up toward the axle, the system will have pro-lift in braking. With front wheel drive, there will be no anti-squat or pro-squat at the rear, because there are no ground-plane forces at the rear tires.

Regarding binding, to minimize that, the points where the lower links attach to the axle should be as close to each other as possible. Depending on packaging and structural constraints, it may also work to have the lower links converge at their front ends instead of the axle, or to use either three trailing links or four parallel ones, and use a diagonal Panhard bar.

MEASURING GRIP

I have a non-aero RWD racecar equipped with a good data acquisition system and data on plenty of tracks. I would like to assess the level of grip of each track. The idea is to be able to quantify what the driver or we usually say about a track, i.e. this one is a low-grip track, this one is a very high-grip one, etc. I would ideally classify each track in categories or something like that…

It's easy to explain but the related question is now 'What's grip?'. To stay with simple things, for me, grip can be considered as a coefficient of friction, i.e. the ratio between vertical force and horizontal force applied on the tyres. But how can I measure 'grip' with usual sensors. The relevant sensors I have now on the car measure lateral and inline aceleration, 4 wheelspeeds and 4 damper displacements. I would be happy to add a few sensors (if not too expensive) to my system to achieve our goal. However I don't think that adding 4 tyre slip sensors for example would be an option.

The data I have is a series of fast laps on various tracks. Using this data would be the first preference, but if not possible asking to the driver to drive a dedicated lap could be a second option.

The tyre is the standard championship slick tyre, so it's always the same. To be fair the aero package on our car is almost nothing. We have a very small rear wing but you can consider that the global level of downforce is negligible.

Our championship is made of 12 rounds on various tracks and they are pretty much the same in the calendar every year. We go often testing on the same tracks but during the season we are not allowed to test on the championship tracks. The maximum number of tests per year on a same circuit is 4 times but more often we go testing once or twice a year on a same test track.

When you have little lift or downforce, use the same type of tire on two or more different tracks, and other aspects of your package don't change, you can reasonably compare track surfaces by comparing accelerometer readings. The accuracy of this method depends considerably on exactly how consistent the tires, car, driver, and everything else are.

You can also get a good idea whether you really do have negligible downforce or lift by comparing suspension displacements and accelerometer readings for high and low speed turns on a given track.

If you want to prevent turn banking from leading you to false conclusions, an accelerometer for the z axis might be a worthwhile addition. However, when there is little aero downforce, you can get pretty much the same thing from a math channel that averages the four damper displacements.

Absent downforce or lift, and absent z accelerations, limit lateral (y axis) acceleration is a good measure of the tires' mean lateral coefficient of friction, and limit longitudinal (x axis) acceleration in braking is a good measure of the tires' mean longitudinal coefficient of friction, <u>if</u> we have the handling and the brakes well enough balanced so that we are using all four tires reasonably fully.

A track doesn't offer the same grip level at all points, but if we're after a general characterization of the surface, we can take an average of all the peak regions of lateral or rearward acceleration. We don't want to take the average acceleration for a whole lap, because the car isn't necessarily grip-limited all the way around the track. Sometimes the car is power-limited instead. On a high-speed track, more of the lap will be done in a power-limited state, and the accelerometer readings will be reduced accordingly. We want to try to pick segments or points where we can be confident that the car is grip-limited, and average those. It is best to pick sampling zones where the car is as purely as possible in lateral or longitudinal acceleration, not both at once. Generally, the longitudinal accelerations will be greater than the lateral ones.

A given track's grip level will also vary with age, and with weather, and even from race to race on the same weekend. Therefore, it is best to average readings from a number of outings, if the aim is to characterize the track.

Comments above assume that we're dealing with dry conditions, and that we are throwing out "outlier" data, such as data we might get after a car has been dropping oil, or after a car has thrown dirt onto the racing surface, or when the car is balked by slower traffic, or when the driver is dicing with another car.

One fact that helps us is that it is impossible to use the tires more than fully, so if we look at the best values from a run, we can be confident that those were situations where the coefficient of friction was the limiting factor, and not something else. (Or at least that's true if there was no contact with anything but the track and the air!) Looking at laps with quick times is a good start.

Using this method in the wet is probably less advisable, simply because we encounter so many degrees of wetness that it's hard to get an "apples to apples" comparison between one track and another. Our accelerometer readings in the wet will also vary according to the driver's success in finding productive ways to use off-line unworn pavement, and the magnitude of this effect will depend on how fresh or polished the track is on the racing line, how smooth the pavement is in the unpolished state, and how much traction-reducing dirt, marbles, and debris are on the off-line areas.

It will be apparent that there are substantial uncertainties and variabilities. I'm sure you'll find that the coefficient of friction is not precisely predictable. (That in itself is worth knowing.) However, with enough of the right kind of data, differences among the tracks should become visible.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
February 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

EFFECT OF CHANGING LOCATION OF JUST THE WHEEL CENTER

We run a 410 winged super sprint (World of Outlaw type car). Some guys claim that they have seen gains by using deeper backspaced wheels (particularly on the RR), and then using the axle spacers to achieve the same tire spacing that you would normally run. (measured from the center of the contact patch to a specific point on the chassis) My question is does the chassis, axle, or tire really know where the wheel center is, as long as the distance from the center of the contact patch to the measured point on the chassis is the same?

I see no reason to suppose there would be any effect from moving just the wheel center, or at least any effect that one could measure or detect when driving the car. The wheel center does have a little bit of mass, and the location of that small mass would change. There could be small changes in the patterns of deflection within the wheel and axle, but it would take sensitive strain gauging to measure those. Compared to the deflections elsewhere, particularly the very soft tires and suspension, they should not be significant.

It might also be possible for a person to think they had the same tire location, but measure or calculate this incorrectly, and reach an erroneous conclusion that way.

Overall, though, changing only the location of the wheel center shouldn't affect the car. If a driver thinks there's a difference, this could be attributable to "placebo effect", or it could be that the car really did go faster after the change, but it really was due to a better tire on the new wheel, or due to the track coming in while the change was being made to the car. It's always difficult to draw conclusions about a car change from lap time or car feel when racing on dirt, because the surface is changing so rapidly that it's impossible to get controlled back-to-back runs on identical racing surfaces. It is possible to get a better idea if we try the same change many times, but in most cases we'll be trying other things as well over a series of races, which again throws some additional ingredients into the pot. One testing technique that sometimes helps is to go back to the original setup and see if the lap times or car behavior go back to original. However, on dirt even this can be tricky, because usually the track first comes in and gets faster as it gets packed and dries from sloppy to tacky, then gets slower again as it dries from tacky to dry slick.

EFFECTS OF DROOP LIMITING

I have some thoughts/questions about droop limiting. This subject comes up from time to time on internet message forums, and I've never read a (or successfully derived and written my own) conclusive explanation on whether or not it "works". Even the way in which it "works" is never clearly specified.

I always approached droop travel in terms of the basic equations in Milliken - it should be proportional to bump, which is itself constrained by ride height/ground clearance and spring rate, but it is not a tuning tool all by itself. Hence when I see a race car with very limited droop travel demonstrated, say, during a pit stop, I would ASSume it is due to the car being fairly stiffly sprung in bump.

Conversely, internet lore seems to be that if you artificially limit droop, with tethers or modified dampers, it will decrease the peak roll angle of the sprung mass. Here's a direct quote from one such discussion: "If the inside can't extend any further, than the roll has to come from the outside compressing." I see. While I find the idea of a non-linear roll gradient interesting (especially in terms of camber compensation), I'm not sure the math supports this conclusion. Maybe what I should say is that I can do math that will support it, but I'm not sure if I'm basing it on solid assumptions.

To confuse me further, in your September 2007 newsletter, in response to question on rear shock droop travel on a Mustang, you wrote: "I can say this with certainty: when the inside rear suspension tops out, the rear roll resistance increases dramatically, and that makes the car looser (produces oversteer)."

My questions:

1. Specific to the September 2007 scenario, what is the relationship between inside rear suspension topping out and roll resistance? In other words, are there solid axle effects involved?

1b. Have the Smithees folks in Australia [referring, I believe, to persons basing their thinking on the writings of the late Carroll Smith] *been right all along when saying that when a damper reaches full droop, all of the remaining load on that side is instantly transferred to the other side? My gut feeling is that if running out of travel caused two wheeled motoring, there would be a travel component in the basic lateral load transfer equation... and there isn't such a term.*

2. Regarding limiters that prevent the suspension from extending fully... the scenario I am picturing is a car in steady state cornering, below the point where it has transferred all of its load from the *inside to the outside, which results in a sprung mass roll angle of X degrees. If at 80% of that roll angle, the inside wheels run out of travel, the lateral load transfer is not high enough to cause/support two wheeled motoring, so the inside wheels will remain on the ground. So now what happens when the lateral acceleration increases up to the delimited example? If the outside continues to compress with increased lateral load, and the inside remains on the ground, this seems to be lowering the dynamic CG.*

When the suspension at any corner of the car runs out of travel in either direction, the suspension at that corner becomes much less compliant: any attempt or tendency to move the suspension beyond that point meets with much stiffer resistance. This is true if the outside suspension bottoms out, as well as if the inside suspension tops out. In banked turns, it is even possible to have bottoming on both sides of the car at once. Cresting a rise, it is possible to top out both sides at once. It is possible to top out or bottom out anywhere from one to all four corners of the car.

It is possible to have a condition of travel limiter contact, but no travel limiter load – that is, a condition where the travel is just barely used up, but the suspension would not move further, even without the travel limiter. We might call this the point of takeup, or impending takeup, or impending bottoming or topping out.

At the point of impending takeup, the travel limiter has no effect on the car at all. The travel limiter starts to affect the car as we start to load the travel limiter. The more we load it, the more it affects the car.

So it is not correct to say that all the load is on the outside wheel as soon as the inside one tops out. Rather, all the load is on the outside wheel when the inside wheel reaches the point of impending lift. This often requires an additional increment of acceleration, although the two thresholds may be reached so close together in time as to appear simultaneous.

When we top out the inside suspension at just one end of the car, any lateral acceleration beyond the value required for impending takeup meets with added elastic roll resistance, just at that end. That does increase load transfer at that end and decrease it at the opposite end.

All types of suspension are subject to these principles, but there is a bit of a difference in the case of a beam axle. Usually, the travel limiter is some distance inboard of the wheel. This means that when we hit the limiter on only one side of the car, further roll does involve some further displacement of the near wheel, although not as much as the far wheel.

For example, if we have a beam axle with the shocks midway between the wheel plane and the car centerplane, and we top out the inside shock in cornering, roll beyond that point involves compression at the outside wheel that is three times as great as the extension at the inside wheel. In effect, the axle has to pivot about the topped-out shock.

If the car has independent suspension, a topped-out shock, or the equivalent, stops displacement at that wheel, and further roll can only occur by compression on the opposite wheel.

Either way, the effect on elastic roll resistance for the wheel pair is a dramatic increase. When we top out the inside suspension, the car does lose ride height and c.g. height as it rolls further. Conversely, if we bottom out the outside suspension, the car gains ride height when it rolls further.

When we bottom out or top out one or both wheels on a side of the car, the effect on elastic pitch resistance for a right or left wheel pair is also a dramatic increase, and this is very important. Topping out the inside suspension on a road racing car has an effect similar to left-stiff spring splits on an oval track car: braking while cornering de-wedges the car and loosens it (adds oversteer); applying power while cornering wedges the car and tightens it (adds understeer).

By "tying down" the suspension in droop, we have a way to make the car "inside stiff" in both right and left turns.

Furthermore, we don't have to tie down the front and rear at the same point, or with the same degree of rigidity. If we only want to tighten exit, we can tie down just the front. If we only want to free up entry, we can tie down just the rear.

And, since the car does corner a bit lower, we do get slightly reduced c.g. height. Whether it's smarter in this regard to just run the car lower statically will depend on how bumpy the turns are, compared to the straightaways, and whether we can run the car lower statically and still be legal.

It will be apparent that this isn't simple, nor a magic way to make any car faster. Using this is tricky. The suspension doesn't know if it's getting topped out or bottomed out due to a crest or dip in the road, or a bump or hole, or braking, or cornering, or power, or banking, or variation in aerodynamic loads. It doesn't know if it's getting topped out or bottomed out due to displacement in roll, pitch, heave, warp, or some combination. When it hits the limiter, it hits the limiter, and from there it's rigid, or much stiffer.

Consequently, using travel limitation of any kind to control car behavior is a strategy to be approached with great caution. Using this strategy makes it harder to get acceptable car behavior in all conditions, and tuning a car using this strategy tends to yield highly track-specific, and condition-specific, setups.

Reasons to put up with this include not only tuning entry and exit balance, but also working ground clearance rules and controlling under-car aero effects.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

OLD-SCHOOL ROCKER SUSPENSION

What is your opinion of rocker arm front suspension where the top A frame continues inboard beyond the pivot point and has the shock vertical at the inner end. Space permitting, it seems a simple and effective system that minimizes joints, levers and rods and while possibly achieving nearly 1/1 shock to wheel ratio.

In its traditional form, the layout does eliminate some parts and some wear points compared to pushrod or pullrod layouts. However, it requires a long piece loaded in bending, and doesn't lend itself to rising rate as readily as pushrods or pullrods. Also, in most of the cars using it back in the day, the coilovers and the box structures containing them were about at the driver's ankles and would trap feet and crush legs in crashes. Early designs, e.g. Lotus 25, weren't so bad in this regard, and rules enacted since then requiring feet to be further aft might address this.

Coilover accessibility and shock cooling were often not very good in these designs.

In most cases, getting a 1:1 motion ratio involves making the upper arms shorter than is really desirable from a geometry standpoint, although this will depend on the particular design.

It is possible to improve the structural efficiency of the control arm/rocker by making it a truss rather than a beam, usually with the truss structure below the pivot axis. In an open-wheel car, the car would then have similar appearance and aerodynamics to a pullrod car.

Or, if we want rising rate geometry, and if the upper arms are inclined considerably to get camber recovery in roll, and if we have room above the pivot axis, we will have a markedly V-shaped rocker in front view. We may then be able to add a compression member across the top of the V, from the upper ball joint to the upper coilover end, and have a truss structure that way.

It would even be possible to incorporate adjustability into such a design, rivaling that available from a conventional pushrod and rocker. The compression member would effectively be a pushrod, of adjustable length, and the system could then perhaps be considered a pushrod system, but with the pushrod and rocker more or less integral with the control arm rather than separate.

This might offer reduction in the number of wear points and load points compared to conventional rockers, and maybe a small weight reduction. The reduced number of load points on the frame would be particularly attractive for tube frame designs.

Whether the whole concept would be advantageous would depend on how it integrates with the overall design of the car.

DE DION TUBE DESIGN AND LOCATION

I am trying to decide on the arrangement to locate a DeDion axle in an autocross car that has a transverse engine/transmission directly in front of the axle. I plan to use a watts link for lateral location. Single trailing arms on each side with a third central link does not seem appropriate because the transverse engine would dictate a very short center link.

I am considering parallel trailing arms on each side. However, I have seen applications that converge the arms on each side to a single front mount. It seems to me that in this situation there might be bending forces applied to the arms when one wheel rises and/or the other falls. Perhaps these are not significant because the axle will have quite limited vertical movement.

Which do you consider more appropriate, parallel arms or triangular arms converging to one front mount on each side?

The simple answer is to go with parallel arms. Converging arms or hairpin-style ones with a single pivot will bind in roll, unless the DeDion tube has a swivel in the middle.

A swivel in the middle complicates the DeDion tube, but this was actually a feature of many designs when DeDion suspension was popular in F1 cars – for example the Mercedes W154. That car used a tube assembly that was rigid in bending and tension/compression, but not in torsion, and single trailing arms, with outboard brakes. Lateral location was provided by a roller in a slot machined into the back of the differential housing.

The Rover 2000/3500TC used a tube that both swiveled and telescoped: it was rigid only in bending. Lateral location was then provided by fixed-length halfshafts, eliminating the need for any plunge accommodation at the shafts, and the need for any additional lateral locating mechanism. Single trailing arms were used, and brakes were inboard.

Both of these designs offer somewhat more than 100% camber recovery in roll (ignoring tire deflection), and a roll center a bit higher than the halfshafts or the roller.

The Mercedes design afforded very ample anti-lift in braking, due to the combination of outboard brakes and single trailing arms. The Rover design does not have comparable anti-lift, despite similar side-view geometry, because the brakes are inboard.

Either of these designs could also use four trailing links, and they wouldn't have to be parallel. With outboard brakes, that would permit having any desired anti-lift in braking, without significant bump steer, which is not possible with parallel trailing links.

LATERAL LOCATION OF LIVE AXLE

I'm building a 67 Camaro street car in "pro-touring" style. I'm sure you are familiar with Pro-Touring in that cars of this build style are expected to perform well in all aspects, from straight line acceleration to handling to braking. Anyway there are several manufacturers producing 4 link or similar solid axle rear suspension kits to replace the original leaf spring rear suspension. The more basic kits usually have a Panhard bar locating the axle laterally. However, some of the more expensive kits use a Watt's link to get rid of the lateral motion of the rear axle through suspension travel that is inherent to the panhard bar design. However, I have seen a few suspensions that use a diagonal bar that connects between the front left to rear right joint (or vice-versa) of the lower links providing triangulation to locate the rear axle laterally. Since the diagonal link moves in the same plane as the other two lower links in basic bump it appears to achieve the same effect as the Watt's linkage but with a much simpler solution. I'm not sure about roll though. Does the diagonal bar in fact achieve the same effect as a Watt's linkage? Can you give a basic comparison of the three systems and when they should or should not be used?

In terms of geometry, the diagonal Panhard bar does give roughly similar results to the more complex Watt linkage. It does have the disadvantage of inducing loads in the trailing links with lateral force, which the Watt linkage does not do. The load paths in a car originally designed to just have leaf springs can be quite good; the diagonal bar can attach near one of the front spring mounts.

One problem can be keeping the diagonal bar from hitting the driveshaft in certain conditions. Also, if the problem that prompts the modification is lateral compliance of the springs and shackles, the bar induces loads that will laterally deflect those parts somewhat, if the bar is simply added to the factory leaf spring suspension.

The diagonal bar has for some time been popular for drag racing rear ends with wide slicks, tubs, and relatively closely-spaced trailing links. In these installations the diagonal bar is closer to longitudinal than to lateral, and the loads on it in hard cornering can get pretty large, as can the induced loads in the other links.

The Watt linkage, when added to a layout originally designed for leaf springs, tends to require additions to the sprung structure, to pick up the loads from the lateral links. Ideally, we would really like the linkage to tie into the existing unibody rails, without the need to add a bunch of tubes to the sprung structure.

One thing that complicates matters here is that usually the unibody rails are higher behind the axle than ahead of it. What we'd really like is a lateral linkage that accommodates this.

It is possible to do this with a Watt linkage, with a bit of imagination. It is possible to have the rocker of the Watt linkage close to horizontal in side view, but a bit higher at the rear than at the front, under one of the axle tubes, anchored to the axle housing to one side of the diff, with one lateral link ahead of the axle, and one behind. The lateral links are then of unequal length. The short one goes to the unibody rail on the side nearest the rocker, avoiding any interference with the driveshaft, and feeds its loads into the rail near the front leaf spring attachment point. The other lateral link passes behind the diff, at a bit greater height from the ground, and feeds its loads into the rail where it curves down behind the axle.

There may be some interference between the rear lateral link and the fuel tank. The best resolution of this will depend on the specific installation. Some structure may have to be added to the rail.

To get the best approximation of straight-line motion with such an arrangement, the ears on the rocker need to be of unequal length, in inverse proportion to the link lengths: long ear or rocker half with short link. This is not absolutely essential, however, if packaging considerations dictate a deviation from this relationship. It is not absolutely essential to have true straight-line motion, just something reasonably close, with good load paths and nothing running out of travel or hitting anything else.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

BANKED TURN PUZZLE

I was just reading your "Measuring Grip" when it hit me that you might be able to help answer a question with respect to calculating the maximum speed of a car in a turn of a given radius, coefficient of friction and banking. The attached Exel doc contains a formula that I found on the internet that is supposed to be useful, but the speed goes infinite at 45 degrees of bank, given a coef of friction of 1, and I maintain that is wrong. My friend who has a masters degree in engineering from GA tech, and another degreed engineer, maintain that the equation is correct. (The equation contains cos theta - mu x sine theta in the denominator, mu = coef of friction)

Can you help?

The equation in question is:

$$v = \sqrt{\left(rg\,(\sin\theta + \mu\cos\theta)/(\cos\theta - \mu\sin\theta)\right)}$$

where:

v = maximum possible linear velocity of the car, ft/sec or M/sec
r = radius of turn, gravitationally horizontal, ft or M
g = acceleration of gravity, ft/sec^2 or M/sec^2
θ = angle of banking, from gravitational horizontal
μ = coefficient of friction

It will be apparent that for $\mu = 1$, the denominator goes to zero when $\theta = 45^\circ$, and v becomes undefined.

It is counterintuitive that the car should have no limiting speed if the banking isn't vertical. It doesn't look right, but it is right. However, for a 45° banking, it's only true if μ is at least 1.

At 45°, what happens is that at 1g horizontal centripetal acceleration (a_H = 1g), no cornering force is required of the tires. When speed drops below the value corresponding to 1g horizontal, the car tries to slide down the banking and the tires must exert a negative cornering force. As the speed rises above the 1g value, the tires must exert a positive cornering force, and you'd think at some speed their grip limit will be exceeded.

But let's try some numbers, not using the spreadsheet or the equation, but just using trigonometry and our own brains.

μ = 1 is the minimum requirement for the car not to slide down the banking at a standstill.

At a_H = 0, the normal force on the tires is W/(√2), force down the banking due to gravity is W/(√2), and force up the banking is 0.

At a_H = 1g, the normal force on the tires is W/(√2) + W/(√2), or (√2)W, force down the banking due to gravity is W/(√2), and force up the banking is W/(√2).

At a_H = 2g, the normal force on the tires is W/(√2) + (√2)W, or (3/√2)W, force down the banking due to gravity is W/(√2), and force up the banking is (√2)W for a required cornering force of W/(√2). The car can do that.

At a_H = 4g, the normal force on the tires is W/(√2) + (2√2)W, or (5/√2)W, force down the banking is still W/(√2), and force up the banking is (2√2)W. The car can do that.

Now we can see the pattern that's emerging. The normal force is always greater than the induced load due to banking by W/(√2), and the net force up the banking is always equal to the induced load *minus* W/(√2). So although the ratio between the cornering force needed and the normal force asymptotically approaches 1, it never gets there. So there will be a limiting speed at some point on a 45° banking if μ is less than 1, but not if it's greater than or equal to 1.

The banking angle where there is no upper limiting speed for any μ is vertical. However, there will then be a minimum speed to keep from sliding down. As μ diminishes, that minimum speed approaches infinity.

In fact, I think we can say that when μ = 1, there is exactly one banking angle with no minimum or maximum speed. When μ < 1 or μ = 1, any banking angle less than 45° has a maximum speed, and any angle over 45° has a minimum speed. When μ > 1, there will be a band of banking angles around 45° with no maximum or minimum.

All of this of course ignores certain realities of tire behavior. μ isn't constant, and we can't just go on adding normal force without failing the tires.

It also ignores the realities of paving such a banking. No race track in the world has such a banking. Talladega is supposed to be 31 degrees. Daytona is supposed to be 30. Charlotte is 23. I have been on the Charlotte banking. It feels like walking on a roof. The track gives van ride-arounds to tourists. As part of the ride, the van stops on the banking. It feels like you're going to tip over, into the infield.

If you wanted to pave a banking at 45˚, how would you keep the paving machine from sliding down the banking? How would you even grade such a banking?

At 23˚, it is difficult to keep asphalt smooth. When the sun shines on the banking, the asphalt softens and starts to sag down the banking due to its own weight. That's why tracks that are any steeper are generally concrete.

In other words, the problem posed here is theoretical. In the world we actually live and race in, you can't just go faster without limit, on any existing track.

Still, the theoretical physics of such a hypothetical case are interesting.

So are the civil engineering aspects of trying to build a track that steep. No doubt it can be done, and even has been done, after a fashion.

I have seen very small, very steep, bowl-shaped circular board tracks at fairs, where motorcycles are ridden around at giddy angles. These are actually portable facilities. They can be taken apart and moved as the midway show travels.

Board tracks for full-sized cars enjoyed brief popularity in the 1920's in the US. Some of these were very steep. They didn't age gracefully. After a few years, boards started coming loose and impaling people, and board tracks were quickly abandoned.

However, if somebody wanted to build a really steeply banked track today, perhaps the board track concept has some ideas to offer.

I don't mean the idea of using wood as the material, but perhaps the idea of assembling the track surface out of pieces formed off-site, or at least not formed in place. It might be possible to make the surface as pre-cast sections of reinforced concrete deck, and then move them into place with a crane. These segments could rest on an earthen roadbed, or a further idea could be borrowed from the board tracks and the sections could be supported on an above-ground framework. The framework could incorporate adjustment to position the segments as things settled, and smooth out the track surface. Sections could be replaced individually. Cracks between sections could be left unsealed for drainage.

Rather like the physics of driving on very steep bankings, this is fun to think about, but I don't plan on holding my breath until somebody does it.

REVERSE-CANT REAR LEAF SPRINGS

What would the effect be if the leaf springs in a road racing vehicle were to have the cant or angle of the leaf springs reversed? ... i.e. The springs are normally closer together at the front pivot point than the rear pivot ....It is my understanding that that the leaf from the front pivot to axle simply acts as a control arm?

With respect to torque reaction, the action of the spring is more complex than just the front half acting as a control arm. In top view, as regards deflection steer with lateral force, the system does tend to mainly yield laterally at the rear portion of the spring and the shackle. As a crude approximation, the axle pivots around an instant center found by projecting the top view centerlines of the springs.

If, as is common, that instant center is ahead of the axle, the system tends to have deflection oversteer: the deflection points the rear wheels out of the turn. If the springs cant out at the front, we get deflection understeer instead.

This could be a good thing. In general, it's better to have deflection understeer in a car than deflection oversteer.

One often sees the view expressed that canting the springs increases the lateral rigidity of the system. I think that would be true if some magic force kept the axle from steering with lateral deflection. But since there is nothing preventing the axle from steering except the springs, within any practical angles, with the springs canted the system still deflects, and we get some deflection steer when it happens.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SELECTING REAR END RATIO

I have a 1966 Mustang Fastback that I have been working on for several years now with the goal of having my favorite vintage muscle car that I can enjoy performance driving through the remote winding/hilly roads throughout the NW Arkansas area as well as to take to a road track (Hallett Motor Racing Circuit - 1.8 mile track) near where I live. At some point down the road I may decide to travel to some other road racing tracks as well and who knows may make the necessary preparations to run in some Vintage Mustang events, but for now I just run in the Open Tracking events at the Hallet Motor Racing Circuit that they call their High Speed Touring Series. I have just in the past few months installed a new engine in my Mustang which has really allowed me to start pushing my car more towards its limits, and in my search to try to optimize performance have noticed that my power band with my current rear gear ratio which is 3.25:1 is not quite where I would like for it to be. I was wondering if you could provide me with a little advice in regards to which gear ratio you would recommend I install in my Mustang to meet my performance desires in the way that I utilize my car as a "pro touring" type for performance street driving through the remote roads of NW Arkansas and also for attempting to push it to the limits at the road course tracks which is my real passion. The following is some additional information about my Mustang:

I currently have the complete Total Control Products (TCP) front suspension system with power rack & pinion and coil overs with dual adjustable shocks, triangulated shock tower supports and full sub-frame cross member assembly system. My rear suspension contains Maier Racing - Race 165 leaf springs and their Panhard rod system. I have a Top Loader 4-speed (1:1 in 4th gear) transmission, a 3 1/2" aluminum drive shaft with 1350 tail shaft and pinion yokes, and a Currie 9" Ford rearend with 3.25 gears, 31 spline axles and TSD (torque sensing) limited slip differential. I am currently running Yokohama 235-45-ZR17 tires mounted on American Racing Torque Thrust II 17 x 8 wheels. My engine is a 351W block stroked to 427 cu.in. displacement producing ~ 530 hp at flywheel, which as mentioned has only been in the car a few months.

I am thinking that I could improve performance by changing my rear gear ratio to either a 3.50 or 3.70 and this is where I could sure use your advice. I have good power with my new engine, but with the 3.25 gears I seem to be lacking some in low end transfer of power to the wheels. I have plenty of top end speed and the benefit of lower rpm's at those higher speeds, but I would like to be able to get to those upper speeds quicker and thus the reason I feel that I need to change the gear ratio. I have calculated my top end speeds using my current tire diameter, my transmission in 4th gear being a 1:1 ratio and top engine speed of 6500 rpm's to be the following for the different rear gear ratios:

Gear Ratio	*Top Speed*
3.25	*150 mph*
3.50	*140 mph*
3.55	*138 mph*
3.70	*132 mph*

The top end speed of 132 mph with the 3.70 gear is just a little bit lower than I would really like, but realize it would give me more lower end power transfer to the wheels. The fastest speed that I have actually had my car up to so far is 135 mph on the longest straight stretch at the local road course track (and I have to admit a few times on about a ½ mile straight stretch of isolated/remote road near my rural home where there are no access roads, houses, businesses, etc.). The winding roads that I normally drive are through remote areas of the Ozark National Forest which is like a 100 mile long road course but doesn't have very many straight stretches with the longest being around a 1/4 mile or so in length. I am fortunate to be able to push my car pretty hard through this area as there is almost never any traffic and no houses, businesses, etc. So as you can see, currently most of my driving is performance street driving through the remote winding/hilly roads of NW Arkansas along with the Open Track (lap timed) road course racing which as mentioned is my real passion.

If I had a 5-speed or some type of overdrive for my current transmission I think it would be more of a no-brainer for me to go with a 3.70 or perhaps even higher gear ratio but with the 1:1 ratio of the 4-speed I just don't have enough experience to know if the 3.70 would meet my needs as described or not.

You don't mention the transmission ratios, other than top being direct. I assume you have the close-ratio version of the transmission, which has third gear of 1.29, second of 1.69, and first of 2.32.

Ignoring frictional losses in the indirect gears, what matters for acceleration is the product of the rear end ratio and the transmission ratio for the gear you're in. This product determines the torque multiplication from the engine to the wheels.

With a shorter rear end ratio (numerically higher, as conventionally expressed), you will have more torque multiplication in a given gear. However, you will also have a lower top speed in that gear, so at certain road speeds you will actually have less torque multiplication overall because you will have to be in a higher transmission gear.

In general, the car will be fastest on a given track if geared so you just barely run out of revs at the end of the longest straight. From the information above, that would suggest trying the 3.50 or the 3.55. Hopefully, if you do improve your acceleration, that will translate to a bit more speed at the end of that straight, and you would rather not have to lift prematurely to avoid overrevving.

At lesser speeds, you may or may not be faster with a shorter rear gear, depending on whether you are forced to upshift. In many cars, there is an advantage to being higher in the gears, because the ratios get closer as we go up. With the 2.32, 1.69, and 1.29 ratios, there isn't much such effect: the spreads are 1.37 from first to second, 1.31 from second to third, and 1.29 from third to fourth. You're not getting significantly closer ratios by being higher up in the box.

Based on the calculated top speeds in high given above, with the current 3.25 rear, you have top speeds in the indirect gears of 65, 89, and 116mph. With the 3.50, those change to 60, 83, and 109.

So, assuming of course that the car is power-limited and not traction-limited, based purely on torque multiplication: with the 3.50 you'd accelerate about 8% quicker from zero to 60mph, from 65 to 83, and from 116 to 140. However, you'd also accelerate about 22% slower (78% as fast) from 60 to 65, 18% slower (82% as fast) from 83 to 89, and 17% slower (83% as fast) from 109 to 116, because in those ranges you have to be in the next gear up.

All of this is based on a simplifying assumption that the car makes equal torque at all rpm's, which of course is not the case. But as long as you're not much past the power peak, it does basically hold true that the car accelerates faster with the shorter rear end gear, except where that gear forces you to upshift and the taller one doesn't.

TORQUE TUBE FOR FRONT-ENGINE IRS CAR

I wonder if you could consider answering this question on the reasons makers use torque tubes please?

Car in mind is a Mazda RX-7 FD model of mid nineties. They run what Mazda term a "Power Plant Frame" which is a lightweight pressed steel frame tying the gearbox extension housing to the front of the differential casing. The car has independent rear suspension by double wishbones. Why would a maker choose this means of locating the rear of the gearbox and the nose of the diff instead of using a gearbox X member and supporting the diff nose of the rar subframe? Some quite low end Opels used a torque tube, as did some of the bigger Peugeots like the old RWD 504. If one were to fabricate conventional mounts for the diff nose using a weld in multi point roll cage to tie the shell *together more rigidly, and did something similar with a different but similarly sited gearbox (thinking engine and box change to something with pistons for circuit race usage), what disadvantages can you think of, handling wise? Any? My personal take on this is that it may have been instigated to allow compliant diff mounting to reduce NVH, yet control unwanted diff nose movement to control drive train wind up and wheel hop. I can't find much on the pros and cons of this setup.*

There are basically two reasons for tying vibration-isolated components together with subframes. The first is to allow soft mounts for good isolation, without incurring undue movement of the individual isolated components. The second is to unite the isolated components as one large mass, which can then be used to achieve a measure of inertia damping by tuning the natural frequency of this mass to interfere with the natural frequencies of the sprung structure as a whole.

In the case of a torque tube or similar structure tying a front engine and transmission to a sprung differential, not only do we suppress side-view windup of the diff in response to axle torque, we also create a structure that resists front-view rotation of the engine and diff relative to each other due to driveshaft torque – the same torque that creates torque roll and torque wedge when the diff is unsprung. This relieves the frame or unibody of the need to resist this torque.

In most cases, there is a penalty in space efficiency and weight efficiency for using subframes, although to some extent this can be recovered in the main structure by either eliminating loadings as just described, or at least spreading loadings among a smaller number of more widely separated points.

When the suspension is softly damped, in pursuit of soft ride, inertial damping from major isolated components can actually help handling as well as ride, as least in terms of the car's behavior on irregular surfaces. But in race cars, we normally use stiffer shocks, which make all tuning of natural frequencies less important, and we want light weight and a stiff overall structure much more than good NVH characteristics.

Therefore, traditionally, we build race cars with little or no isolation of the engine, trans, and diff. We solid-mount everything, try to get some structural stiffness gain from the components, and deal with any NVH issues by having a loud exhaust and wearing earplugs. If the shocks are stiff enough for good handling, there won't be much oscillatory behavior from the suspension.

However, in recent years race car designers have taken a fresh look at inertia damping, as readers will know who have followed inertia damping's recent introduction, and prohibition, in F1. The reason for the renewed interest in inertia damping is that modern high-downforce race cars have such stiff springing that tire deflection becomes a very significant portion of the total "suspension" compliance: the tires deflect anywhere from half as much to just as much as the suspension proper.

The tire deflection does not displace the shock absorbers, so the shocks can't damp it. There is some internal damping in a tire, but nowhere near as much as we'd like. A very stiffly sprung car can bounce or pitch on the tires much like a tractor (although with smaller amplitude and higher frequency), and the shocks can't suppress this – hence the interest in inertia damping.

So if the car is going to have ground effects and wings, and very stiff suspension, there could be a case for trying to use compliant mounting to get some inertia damping. This could involve using a torque tube, or not. It is common practice in live-axle passenger cars to use the engine/trans assembly for inertial damping, without a torque tube. It would be quite possible to do that in a race car with a sprung diff, and solid-mount the diff. The only downside would be that the engine/trans assembly will undergo rotational displacement on its mounts when it applies torque to the propshaft, and all packaging, plumbing, wiring, and linkages will have to accommodate that movement. This can be mitigated somewhat by wide spacing of the motor mounts.

It should be mentioned that we don't necessarily get inertia damping from compliant mounting of a major mass. If the frequency is not tailored to the tire and suspension frequencies, it is quite possible to get reinforcement of suspension and tire oscillation, rather than interference. So this would not be something to be undertaken lightly. For most of us, the most prudent recommendation is still to go for rigidity, lightness, and simplicity, and solid-mount everything.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TORSION BAR TUBE HEIGHT

What is the significance, if any, of the torsion tube mounting height on the chassis of a sprint car or Northeastern DIRT modified which both use cross torsion bar springs? Sprint cars at both front and rear (usually) and modifieds in the rear. I hear a lot of talk about low rack, mid-rack, flop tube, stack tube etc. To my way of thinking the only differences that any of these configurations can make are torsion arm angles (static/dynamic) and torsion arm lengths left to right. Arm angles would determine if the torsion arm/bar has a rising or falling rate. And arm length would just effect the effective rate of the bar at the axle. Am I missing something? For that matter, we are always discussing spring mounting in general at the axle end – spring base, front or rear of axle, directly on the axle or on a control arm (introduces side view motion ratio) but what are the implications of the upper spring/coil over chassis mounting points (excluding mounting angle)? i.e. upper spring mounting point height.

When the springs, regardless of their type, act only as springs, and not as part of the locating linkage, it shouldn't matter much how high or low the springs are. What matters is how far the spring displaces, for each inch of displacement at the wheels, in the mode we are examining. Spring mounting only matters as it relates to that.

However, when the spring doubles as a locating member, a part of the suspension that positions an upright or an axle, then its location determines some aspect of the suspension geometry.

In the case of a simple beam axle on leaf springs, the leaf springs locate the axle laterally, and therefore their height determines the roll center height.

Although the front and rear torsion bar installations on a typical sprint car may look similar, they are dissimilar as regards the use of the torsion bar arms as part of the suspension linkage. At the front, the bar arms simply rest on the axle tube, and are free to slide on the tube. They do not act as suspension links. All they do is hold the car up. In some cases, a separate roller above the axle tube may be provided, but the effect is the same.

At the rear, the situation is different. The torsion bar arms serve as the leading links in a Watt linkage that locates the axle longitudinally. The leading end of the arm typically has a spherical joint. The spherical joint rides on a bolt or pin on the bottom ear of the birdcage, on which the spherical joint can slide laterally. This allows the suspension to move freely in roll. A tubular trailing link runs forward from the top ear of the birdcage to the frame, completing the Watt linkage mechanism.

Ordinarily, at static condition both the trailing link and the torsion bar arm are close to horizontal, and the axle moves very nearly vertically near static position. The system has little roll or bump steer, and little thrust anti-squat or anti-lift. It does have considerable torque anti-squat and anti-lift, from the vertical forces at the ball of the torque tube (sprint car) or front of the torque arm (DIRT Modified).

If we just raise the rear of the torsion bar arm, and leave all other points unchanged, the axle moves forward as the suspension compresses, creating thrust pro-squat and pro-lift, and roll understeer. The Watt linkage has a side-view instant center behind and above the axle.

Thus, although the roll resistance and wheel rates are not necessarily affected by torsion bar mounting height, the side-view geometry, and attendant properties, do change.

TWIN I BEAM SUSPENSION

I am designing a vehicle to be raced in the desert and I have been looking at different front suspension designs and have been considering using the twin I beam suspension seeing it allows a lot of travel and has been known to do well in the past. So I am interested if you might know how to calculate the roll center with the twin I beam, and if there is anything else I may want to look for when using the twin I beam?

For readers unfamiliar with this suspension, it is basically a modified swing axle suspension, used mainly on Ford trucks (the questioner's e-mail address suggested the vehicle is a Ford Bronco). Each wheel is on a single transverse arm, with a leading arm added for longitudinal location and torque reaction. The only difference between this and a classical swing axle is that the transverse arms extend past the middle of the vehicle, crossing over each other. Usually, they extend to the opposite frame rail.

The spindles are similar to those on a beam axle. There are no ball joints, so there is no worry about running these out of travel. This accounts for the system's compatibility with large amounts of travel.

The disadvantages are similar to the disadvantages of a swing axle suspension, only not quite as bad. The camber changes a lot as the suspension displaces, typically around a degree and a third per inch of displacement. The roll center is a bit higher than we would really like, and the system jacks noticeably when the ground-plane forces get large. It also jacks in braking; it has more than 100% anti-dive.

But if the surface is loose and slippery, and the tires are squishy and rounded in profile, and the bumps are huge, the system's drawbacks don't become too conspicuous, and its advantages – simplicity and long travel – begin to make it attractive.

It is even more attractive if the vehicle we're using comes with this suspension from the factory. In that case, using anything else requires a lot of fabrication, and using the factory suspension, with minor modifications, saves a lot of money and work.

The earliest Ford Broncos used the twin I-beam system only on the two-wheel-drive models. Four-wheel-drive versions got a beam axle on coil springs, located laterally by two transverse links anchored to the frame where the pivots would be in a two-wheel-drive version. This would have created a bind, but Ford made it work by using highly compliant rubber bushings on the links.

So if the off-road racer being contemplated really runs off-road all the time, this suspension is worth considering. If there is a prospect of competing in events with significant pavement stages, it starts to look less attractive. It can be lived with on pavement, if the driver is brave, but a short-and-long-arm system will outperform it, and so will a beam axle.

As to where the roll center is, the system can be thought of as a short-and-long-arm system with the front-view instant center at the beam pivot. To find the approximate roll center graphically, we construct a line from the contact patch center to the pivot in front view, and see where that line crosses the vehicle centerline. This is not exactly correct, but it's close. Typically, the roll center ends up somewhere around eight inches above the ground, give or take a couple of inches. That's higher than we would really like, because an independent suspension with the roll center that high will tend to jack noticeably, but this really only becomes a problem on pavement, and perhaps on tacky clay.

RELATIONSHIP OF TIRE SIZE AND WEIGHT DISTRIBUTION

Many car owners say that because their car has 50/50 weight distribution, they inherently have the best handling possible. However, many owners install wider tires at the rear to improve traction. My opinion is that by doing so, while their car may still have 50/50 weight distribution – at rest – the dynamic balance will be a car that understeers. True?

On a related note, let's say we have a mid-engine car with 40% front / 60% rear weight distribution. If the rear tires are made 50% wider than the front tires, wouldn't this negate the static weight distribution, essentially creating a car with more balanced traction front and rear (granted, on a *constant-speed circle?) I understand there may be an issue with getting enough heat into the tires to generate traction proportional to their width. This assumption aside, is this thinking correct?*

It is definitely correct that, at least up to a point, we can compensate for weight distribution with unequal size tires. We can also compensate with roll resistance distribution, camber, tire pressure, and aerodynamics.

So it is safe to say that if the car has 50% rear weight, rear drive, and bigger tires in back, it will probably understeer in steady-state cornering if camber properties and settings, overall roll resistance, tire pressure, and downforce are equal at both ends, and if speed is low enough so that the rear tires do not have to transmit a huge amount of power just to maintain steady speed. If any of these conditions are not present, all bets are off – and usually not all of these conditions are present.

We can at least say that if we take an existing car, and add tire size at the rear only, leaving all else unchanged, the tire size change will move the car toward understeer.

In many forms of racing, we do not have free choice of tire sizes; the rules impose a maximum size. This may be the same for both ends, or not. The rules often also impose restrictions on other aspects of the car's design, which limit what sort of weight distribution we can have, and what sort of downforce distribution we can have.

Nonetheless, it is interesting to consider what we should want in terms of weight distribution and tire size in a rear-drive car, given a free hand.

It is useful to note what course car evolution took in F1 and in sportscar racing, before tire sizes were limited. The cars were decidedly tail-heavy – around 60% rear – and the rear tires were about a foot and a half wide. Front tires had about 2/3 the tread width of the rears. The cars had more aero downforce at the rear than at the front.

Why would this be better than 50% rear, and equal-size tires? There are various reasons, but probably the main one is that the car doesn't just have to corner; it also has to brake and put power down. It brakes with all four wheels, and propels itself with only the rears.

Even without downforce, in straight-line limit braking about 20% of the wheel loading transfers from the rear to the front. Even with 60% static rear, the car does 60% or more of its braking with the front wheels. A car with 50% static rear and modest downforce does about 70% of its braking with the front wheels. It is possible to get the brake bias needed, but it becomes difficult to sustain it over the length of a race. The front brakes tend to overheat and go away.

As for putting power to the ground, there is no mystery as to why more rear percentage is an advantage with rear wheel drive.

Finally, in most cases it is difficult to accommodate foot-and-a-half wide tires on the front of the car, especially without power steering. The scrub radius ends up being really large, and the car gets difficult to steer.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TUNING TRANSIENT BEHAVIOR OF FRONT-WHEEL-DRIVE AUTOCROSS CAR

I'm an autocrosser driving a left heavy FWD car. I run in the SCCA Street Prepared class, allowing me spring and swaybar changes. The car as it sits right now is roughly 2500lbs with driver, 63% front, 54% left, with ~210hp. I'm running ~550lb front wheel rate, and a ~495lb rear wheel rate. I'm contemplating changing the right side springs so that they are softer then the left side springs to improve rotation and power down in right hand turns. I expect that this should give more even front tire loading under acceleration, but more uneven front tire loading under braking. I'd also expect that the car would be looser under acceleration when turning right, but tighter turning right under braking. It seems as if the drawbacks will outweigh the benefits, considering the car brakes at 1.1g's but accelerates at only .3g's. Is there a better way to get better tire loading of the inside front tire in right hand turns on a left heavy car under rules that prohibit moving mass around in the car? Perhaps preload on the swaybar?

The second question pertains to slaloms. In case it matters, my car uses a 275/35/15 Hoosier R6 front tire on a 15x10.5" wheel with a 205/50/15 Hoosier A6 rear tire on a 15x6" wheel. I'm familiar with your article on shock tuning for the various phases of a corner. However, in a high speed slalom reality seems to diverge from theory. What I mean by that is that on my car, if I need to tighten the car in a slalom, increasing front damping seems to make the car looser. Conversely, increasing rear damping makes the car tighter. I do feel the damping changes affect normal corner entry as expected, so I'm wondering if my simple analysis of a slalom being a series of corner entry events is flawed. My shocks do have independent compression and rebound adjustment, but partly due to lack of testing time I've been erring on the side of caution and changing both compression in rebound settings in equal amounts and direction (IE, add 2 clicks of rebound and compression). The other part of the reasoning comes back to the analysis of a slalom. If one thinks of it as a series of corner entry/corner exit events compressed closely together, it would seem adding front compression would tighten the entry portion and adding front rebound would tighten the exit portion. Instead it seems as if I speed up how quickly the front of the car reacts to the steering inputs and then the outside rear tire can't get to a point where it is generating sufficient grip as quickly as it needs to.

Other adjustments to tighten the car in slaloms seem to tighten the car everywhere else. I've tried rear toe in (I currently run between 1/16" to 1/8" total rear toe out), less rear tire pressure, more rear camber, softer rear bar, and softer rear springs. While effective to varying degrees, the increase in understeer on sweepers is undesirable. Since this seems to be a transitional issue, it would seem that shocks would be the proper tuning tool.

Another thought on this subject is that, since I run the rear tire on a rather narrow wheel for the express purpose of getting the rear tires to operate at a higher slip angle and thus make the car looser in sweepers, would it be possible that a wider rear wheel might help the rear tires get to the slip angle they need to generate grip sooner, thus mitigating the "snapping towel" effect I am picturing?

Taking the first part first, the questioner's analysis of the effects of left-stiff springing in a right turn are essentially correct: the car should be tighter when trailbraking, and looser under power, and the tendency toward inside front wheelspin should be reduced. In left turns, all these effects should be reversed.

Note, however, that this only applies when the car has all four wheels on the ground. Once the inside rear wheel lifts, it is no longer possible to affect wheel load distribution using the right/left pitch resistance distribution or the front/rear roll resistance distribution. The questioner does not mention whether the car runs on four wheels in right turns, or three.

A case can be made that the right/left wheel rate split should be roughly proportional to the right/left static weight split. This would make the percentagewise longitudinal load transfer on each side similar, or at least the elastic component of it. That would actually tend to keep the left percentage at each end more nearly constant in pure longitudinal acceleration than a setup with identical wheel rates right and left. It also would, in theory, give a better ride, because the car would have similar natural frequencies right and left.

If this approach were pursued by softening the right side of the car, we would shoot for a right front wheel rate around 470 lb/in and a right rear wheel rate around 425 lb/in.

It would probably also make sense to put one or two more clicks in the shocks on the left than on the right, to keep the damping ratios more similar on both sides of the car.

Now, regarding the case of a slalom: a slalom is indeed a series of corner entry events interspersed with corner exit events, usually with nothing else in between. We are adding lateral acceleration and yaw velocity in one direction, and roll displacement in the opposite direction, until we are alongside the cone. Then we are decreasing lateral acceleration and yaw velocity, and de-rolling, until we are midway between cones. Approximately midway between cones, we reach a point where roll displacement is zero; roll velocity is inward with respect to the cone behind us and outward with respect to the cone ahead of us; yaw velocity is zero; yaw acceleration is outward with respect to the cone behind us and inward with respect to the cone ahead of us; the car is at something very close to static attitude (zero suspension displacement), but it is not at zero yaw or roll velocity, nor zero yaw or roll acceleration.

Up to the point where we are beside the cone, we have a corner entry event. From there to the approximate midpoint between cones, we have a corner exit event. Then we more or less seamlessly enter the next corner entry event, for a turn in the opposite direction.

Depending on the car and the driver's style, there may be some application and release of power, just to affect the car's oversteer/understeer balance, or the whole slalom may be taken at constant throttle. Either way, speed is not going to vary a great deal, and longitudinal accelerations will be modest. Accordingly, pitch displacements and velocities will be small. Except possibly for the very first and very last cones, straightaways are not a factor, so we are not concerned with trying to get "in fast" by braking late, or "out fast" by putting a lot of power down.

In most cases, the surface will be approximately flat. In some cases there may be some humps and dips in the surface, or a bit of tilt, but it is very uncommon indeed to have banked turns in a slalom.

The whole process can be thought of as a series of roughly sinusoidal cycles, in yaw displacement, yaw velocity, and yaw acceleration, and in roll displacement, roll velocity, and roll acceleration, with not much else going on. The respective cycles of displacement, velocity, and acceleration will be offset, or phased, 90 degrees apart.

To illustrate, suppose we are in a slalom, midway between cones, approaching a cone that we will go around on its right, turning the car leftward. The cone after that will be taken on its left, turning the car rightward. After that one, we will be in a condition similar to our starting condition. That will be one cycle. For simplicity, let's suppose that the car takes one second to get from one cone to the next. A full cycle then takes two seconds. Let's call the condition where the car is parallel with the line of cones zero yaw displacement.

Let's follow the car through a cycle, and look at the velocities and accelerations. To begin with, let's ignore effects resulting from vehicle attitude angle or drift angle and from damping, for simplicity.

At t = 0, the car has roughly zero roll displacement and yaw velocity. It is in the middle of a transition from turning rightward to turning left. Its roll velocity is rightward, and at a maximum. Its roll acceleration is zero, transitioning from leftward to rightward. Yaw displacement is rightward, and at a maximum. Yaw acceleration is leftward, and at a maximum.

At t = 0.5s, the car is beside the first cone, to its right, and cornering hard leftward. Roll displacement is rightward, and at a maximum. Roll velocity is zero, transitioning from roll (rightward) to de-roll (leftward). Roll acceleration is leftward, and at a maximum. Yaw displacement is zero. Yaw velocity is leftward, and at a maximum. Yaw acceleration is zero.

At t = 1.0s, the car is midway between cones, halfway through the cycle. Roll displacement is zero, transitioning from rightward to leftward. Roll velocity is leftward, and at a maximum. Roll acceleration is zero, transitioning from leftward to rightward. Yaw displacement is leftward, and at a maximum. Yaw velocity is zero, transitioning from leftward to rightward. Yaw acceleration is rightward, and at a maximum.

At t = 1.5s, the car is beside a cone again, to its left, cornering hard rightward. Roll displacement is leftward, and at a maximum. Roll velocity is zero, transitioning from leftward to rightward. Roll acceleration is rightward, and at a maximum. Yaw displacement is zero, transitioning from leftward to rightward. Yaw velocity is rightward, and at a maximum. Yaw acceleration is zero, transitioning from rightward to leftward.

At t = 2.0s, we have completed one cycle, and are back to our starting condition, two cones down the line.

Acceleration changes velocity, and velocity in turn changes displacement. Acceleration leads velocity by 0.5s or a quarter of a cycle, and velocity leads displacement by a similar amount. Displacement and acceleration are half a cycle out of phase: they zero together, and they peak together, but in opposite directions. All the functions are approximately sinusoidal, but with phase shifts between them. This is the expected pattern for first and second derivatives of a sinusoidal function. Velocity is the first derivative of displacement with respect to time, and acceleration is the second derivative.

In general, any car wants a tighter setup for slaloms than it wants for other situations. This is partly because the magnitude of peak yaw accelerations (in the preceding example, these would be at or near t = 0, t = 1.0, and t = 2.0) increases as the vehicle's attitude or drift angle mid-turn (t = 0.5 and t = 1.5) increases. The car's yaw inertia tends to tighten it during the entry events and loosen it during the exit events. The most common problem is that the car is too loose during the exit phases ($0.5 < t < 1.0$; $1.5 < t < 2.0$), rather than too tight during the entry phases ($0 < t < 0.5$; $1.0 < t < 1.5$). Thus, any problems with the car's balance will tend to be most conspicuous during the exit events.

How does this relate to damper tuning?

First of all, everyplace except right beside the cones, we have something approximating pure roll velocity. Therefore, we have extension velocity on both shocks on one side of the car, and compression velocity on both shocks on the other side. This is true during both the exit phases and the entry phases. So we can't tune entry with compression and exit with rebound. What we can do is tune the relationship between entry and exit with the relationship of front and rear low-speed damping, meaning both compression and rebound. That is, we can stiffen the front damping in compression and/or rebound, relative to the rear, or vice versa.

The dampers generate a roll moment that is opposite in direction to roll velocity. During the entry events, they fight the roll and generate an anti-roll moment. During the exit events, they fight the de-roll and generate a pro-roll moment. Consequently, adding damping at one end of the car makes tire loading at that end more unequal during entry and less unequal during exit.

So at least in theory, adding front damping should not loosen the car in all parts of the slalom, but it should loosen it during the exit events, and that is exactly the type of balance problem that tends to be most conspicuous in a slalom.

In general, it helps a car in a slalom if we stiffen the rear damping. This does tend to hurt our ability to put power down on exit with the inside front wheel, unfortunately. One way to mitigate this is to take advantage of the fact that the car has some rearward pitch velocity as we add power during exit, meaning that the inside rear is compressing faster than the outside rear is extending. If we add our rear damping mainly in extension, and not in compression, we maximize the effect in slaloms and minimize the effect on exit drive traction around the rest of the course.

Regarding the choice of rear tires and rims, it is important to remember that balancing the car by reducing grip at the rear, or just adding slip angle at the rear, may make it feel better, but it doesn't make it faster. If the front end is what limits the car, the car goes just as fast through sweepers with stiff, sticky tires on the rear as it does with floppy, slidey ones. It has more understeer, but only because the rear has more grip, not because the front has less. And the car will be faster in the slaloms.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHEN LOTS OF FRONT ROLL STIFFNESS HELPS, AND WHEN IT DOESN'T

The following question was forwarded to me by a correspondent who said they got it from a Formula Student/Formula SAE forum, where somebody was suggesting it should be put to me.

In the UK most of the successful hillclimb and sprint cars have a monoshock front suspension which is very stiff in roll. Examples include the Gould and Force cars, and more recently Graeme Wight Jr's Raptor design. The rear suspension is often considerably softer in both bounce and in roll. As a consequence the cars often corner on three wheels, with the inside front lifting. Why is this a successful method when all of the suspension/chassis technical books suggest that weight transfer results in a loss in overall grip due to the non-linearity of the tyre?"

We have debated this topic in the amateur paddock and have come to the conclusion is that these technical sources are generally correct for tyres which are already up to temperature. However in a hillclimb or sprint the tyres are initially cold, so the grip gained by tyre temperature increase may outweigh the loss in grip through loading up the outside tyre.

Furthermore the hillclimb tracks are often very heavily cambered so cornering on 3 wheels may improve car stability and tyre contact patch. My analogy would be the stability of a table on any uneven floor; stability can be achieved by 3 good points of contact.

I have added A/R bars to my OMS in order to stiffen the car in roll, particularly at the front. My OMS 1100 has conventional independent suspension all round with Koni double adjust dampers. Attached is a photo at an Aintree sprint showing the inside front just lifting. The car feels better with the A/R bars and takes sweeping corners better than before. I am convinced that I can lean on the car much more now in corners.

First of all, it is vital to recognize that having a lot of front roll resistance, relative to the rear, does not increase overall load transfer. It increases front load transfer, and decreases rear load transfer, and the total remains essentially unchanged.

Increasing overall roll resistance does not increase load transfer either. It just reduces roll, and, accordingly, camber change. It is true that stiffer settings tend to heat tires faster, and more, because the tires experience more dramatic load changes as they go over bumps.

Whether more heat helps the tire depends on the compound. In general, tires intended for long runs are designed to heat up and get tacky. Tires intended for short events such as sprints, hillclimbs, or autocross often use compounds that are tacky at room temperature. The Hoosier tires used by Formula SAE teams I have worked with use the same A25 compound for the slicks and for the moulded rain treads: the dry tires use a rain compound, and consequently work well when moderately cold. They do get hot as they run, but they don't benefit much from it.

It probably makes good sense to use heat-insensitive compounds for short events, if for no other reason than to give the driver a car that doesn't change dramatically in the course of a run.

Whether they are heat-sensitive or not, tires generally do follow accepted theory as regards load sensitivity of the coefficient of friction, cold or hot: the coefficient of friction does decrease as normal force increases; a pair of tires make less force the more unequally they are loaded.

Consequently, a setup with more front roll resistance, relative to rear, will be tighter (have more understeer). This tends to be a good thing in sweepers, and a bad thing in tight turns. Not only do most drivers prefer a looser (more oversteering) car in tighter turns, but cars tend to get tighter as turn radius decreases, especially at really small radii, as seen in FS/FSAE events.

This is partly due to "off-tracking": the rear tires track inside of the front ones in a tight turn, even when running at a significant slip angle, and track outside of the fronts in a large-radius turn when running at a slip angle. Even at steady speed, as on a skidpad, the driver must apply substantial power in limit cornering, just to overcome tire drag. Assuming rear drive, when the rear tires are tracking inside of the fronts, the drag force from the front tires creates a pro-understeer (outward to the turn) yaw moment, and so does the propulsion thrust from the rear tires. When the rear tires track outside the fronts, the effect reverses, and the yaw moments from drag and thrust add oversteer.

Again assuming rear drive, a car also tends to get looser in conditions requiring heavier throttle application: corner exit; hill climbing; high speed turns. This results from the "traction circle" effect: the tires can't make as much lateral force when they are being asked to make longitudinal force.

Uphill, this effect is partially countered by rearward load transfer resulting from gravity, but generally a powerful car will be looser uphill than on level ground. This explains why, in the old days, people put dual rear tires on hillclimb cars.

Finally, if the car has any form of limited-slip differential, or a spool, there is a "locked axle push" effect when cornering, and this becomes more pronounced in smaller-radius turns.

This explains why autocross cars, and especially FS/FSAE cars, don't want a front-stiff setup, while hillclimb cars do, and cars for full-size road courses generally fall somewhere in between, depending on aero properties. FSAE events are run in parking lots. There is little climbing or descending involved. There are no sweepers. The courses are designed to keep speeds very low, in the interest of safety.

The last FSAE car I drove has roughly 57% static rear weight, give or take about 1% depending on driver weight. It has equal size tires front and rear. Ordinarily, one would expect a car like that to have some oversteer even if it were set up to lift a front wheel. But on the turn radii used in FSAE, that doesn't happen. The car has to be set up so rear-stiff that inside rear wheelspin is a problem. When you are cornering near the lateral force limit, you can't throttle-steer; the inside rear spins, and the outside rear stays stuck. That's with a Gleason-style limited-slip. The inside rear can't transmit enough torque to lock the Gleason. Earlier cars, with less rear percentage, would lift the inside rear off the ground on the skidpad.

Any attempt to get more inside rear loading resulted in excessive understeer.

The cars also use considerable static negative camber at the front, and considerable caster, so that the front wheels corner at more favorable camber than the rears, or at least the outside front does.

Typical skidpad sizes for testing full-size cars are 200 or 300 feet diameter. FSAE skidpads are 50 feet in diameter at the inside. A good car can pull around 1.3g lateral acceleration. The car is at its steady-state lateral force limit, at a speed of about 25mph.

These cars should be loose in a sweeper. I can't tell you from experience whether they are. We have never run an FSAE car anyplace where there was a sweeper.

Finally, we should note that this whole discussion applies to rear-drive cars, and that when we speak of greater or lesser front-stiff roll resistance distribution, we are speaking in relative terms. A given design will need a more rear-stiff setup when used for autocross than when used for hillclimbing or full-size road course work, but that does not necessarily mean that all autocross cars are more rear-stiff than all hillclimb cars, irrespective of all other factors.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHEN ONE IS A VERY SMALL ANIMAL: CONSIDERING AN RC CAR

After reading through some of the newsletters you have written about vehicle suspensions etc., I was hoping I could ask you a few questions about an rc model car. Since a year ago I have a new hobby in racing with remote control cars. These are electric motor powered by battery 4wd models with independent suspension. Perhaps you have seen them before, but just in case, I have included a few pictures in the email. At the moment I work as a structural engineer but since starting the hobby I have been catching up on my original automotive engineering degree and dusting off my vehicle dynamics books.

The car actually races "on road" on tracks especially suited to their size with proper very fine, smooth, tarmac surfaces.

The size scale is supposed to be 1/10 scale. Wheelbase is about 254mm. Track width is about 167mm in the front, with the rear usually 1-2mm narrower. The minimum weight limit is 1350 grams. The cars are constant 4-wheel drive with a locked front axle and a rear differential. Depending on the type of motor and controller used, the top speed can reach anywhere between 60km/h to 115km/h. During the races I compete in, on the track with the largest straight we reached 90km/h. So the speed is far from scale.

Perhaps these cars are completely different to their full size counterparts but I still find it fascinating to see some of the numbers that come up when I check them. For instance: roll centres below ground; sprung cg below the unsprung cg height; a locked front axle (both wheels rotate at the same speed); natural frequency of the suspension in the order of 8-12Hz; pressureless rubber tires filled with a foam insert, and tire additives to increase grip of the tires are allowed.

The unsprung mass c.g. is above the sprung mass c.g. due to the fact that the tire diameter is 63mm, with the tire, bearings, and uprights being the parts making up the unsprung at 31.5mm above ground, while the unsprung mass major components are the batteries, motor, servo and speed *control module. These are all located directly on the carbon chassis plate. With a ride height of 5mm, the sprung c.g. is located about 20-25mm above ground.*

I worked my way through determining the roll centres, basic wheel rates, load transfer in a general 1g corner, and roll stiffness distribution to better understand how the car works. I was hoping that by better understanding the engineering part I could find my advantage over the younger natural talent drivers who might not be so proficient in this area.

Now I got a little bit puzzled over a few things in the suspension. It is supposed to be a H-arm with camber link system for all four wheels. But how to find the antiforces in this type suspension? The inside point of the H-arm can be angled by shimming more (or less) on one side of the mounting points (mounting point you can seen in the rear view image). This should create anti's but also change the arm travel (wheelbase change and track width changes).

Would the changes to track width and wheelbase be more influence compared to the antis created (if antis are created as this is the original question)? Secondly, these cars frequently use droop limitation by screws in the H-arm that restrict how far the arm is allowed to travel down to the chassis plate.

I understand the droop limitation effect in the accelerating and braking condition, but what will happen in roll? How will the wheel loads respond when the inside wheel reaches its droop limitation point while the outside wheel is still suspended by the spring? When this happens is the chassis roll completely determined by the compression of the outside spring?

I hope you can give me some help in order to decipher what makes these cars work, and why some things are done this way in these cars.

I will not attempt to present myself as an expert on these cars. I have never encountered them before, much less raced one or advised anybody who raced one. But it is fascinating to consider some of the things that happen when one tries to drastically scale a mechanism up or down.

In general, small machines, or animals, will have better "scale performance" than their larger counterparts. Some effects don't change with scale at all; some change dramatically.

For example, the fact that ants can lift astonishing percentages of their body weight is primarily due to their small size. If there were giant ants as depicted in '50's horror movies, they probably couldn't walk, or even breathe adequately to support their body mass. Mice have skinny little legs and pudgy little bellies, yet can jump amazing multiples of their own height or length.

These effects result from the fact that a muscle's cross-sectional area, which largely determines its strength, increases with the second power of scale, while its volume, which largely determines its mass, increases with the third power of scale.

Very small creatures (spiders, small insects) can be dropped off a roof and not get hurt, because they reach aerodynamic terminal velocity at such a low speed that it's like they have parachutes on. If we take a given engine cylinder and scale it up, it has dramatically less valve and port area per unit of displacement, and will theoretically make less power per unit of displacement.

These effects occur because an object's surface-to-volume ratio goes up as it gets smaller. Taking a sphere as a representative example, frontal area and surface area both increase with the second power of scale, while volume (hence mass, if the sphere is solid) increases with the third power of scale.

If we could scale a car down to 1/10 of its size, without having to completely rethink the entire machine, it would have 1/100 of its previous surface area and frontal area, and only 1/1000 of its previous mass. It would have ten times as much valve area per unit of displacement. It would have ten times the brake swept area per unit of weight. If the static deflection of the tires stayed in scale, the car would have ten times the contact patch area per unit of weight. This would require a tire vertical spring rate 1/100 that of the full-size car.

Likewise, to keep the suspension's static deflection in scale (1/10 as great), the wheel rate in ride would have to be 1/100 as great, while the ¼-car sprung mass would be 1/1000 as great. Undamped natural frequency is entirely a function of absolute (not scale) static deflection. It varies inversely with the square root of static deflection. So if the static deflection is 1/10 as great, the natural frequency is greater by a factor of $\sqrt{10}$, or $10^{1/2}$, which is about 3.16. In fact, no matter what the 1/10 scale car actually weighs, if the static deflection in Earth's gravity is to scale, the natural frequency will be about 3.16 times as great as that of its full-scale counterpart.

So a 1/10-scale car with a natural frequency of 8 Hz has a scale static deflection corresponding to a full-scale car with a frequency of 2.5 Hz. A 1/10-scale car with a natural frequency of 12 Hz has a static deflection corresponding to a full-scale one with a frequency of 3.8 Hz. Assuming that the bumpiness of the track surface really is to scale, those would be downforce car numbers.

What, then, of the likely downforce characteristics of a 1/10-scale car?

The density of these cars, if they are close to the minimum weight of 1350 grams, is remarkably similar to that of a full-scale car. A full-scale one would be 1350 Kg, or 2970 lb. So the car's mass is about 1/1000 of a full-size car's, while its aerodynamic surface area is 1/100 as large. The RC car has ten times the aerodynamic surface area per unit of mass. And it goes about ½ to 1/3 as fast, meaning that its scale speed is 3 to 5 times that of the full-size car.

The Reynolds number for the RC car is only about 1/20 to 1/30 that of the full-size car, so air flow patterns may be fairly dissimilar, but just using the crude assumption that aero forces are proportional to the square of speed, and proportional to surface area, the RC car, going ½ as fast as the full-scale one should generate about $(1/2)^2 * (1/100)$, or .0025 times the downforce, while having .001 times the mass. That would mean that the RC car has two and a half times the downforce, relative to its mass. This would be comparing the RC car at 90 kph to a full-scale car at 180 kph (112 mph), or the RC car at 115 kph to a full-scale car at 230 kph (144 mph).

Or, for 1/3 the speed, the downforce is $(1/3)^2 * (1/100)$, or .0011 times as great. That would be a ratio of downforce to mass, for 90 kph or 115 kph, similar to a full-size car at 270 kph (169 mph) or 345 kph (216 mph) respectively.

This means that the RC car should have roughly 2.5 to 1.1 times the suspension displacement due to downforce, compared to the full-size car, for a given static deflection. But note: that's actual displacement, and actual static deflection, not scale! In other words, a 1/10-scale car with a model production sports car body, with just a wing and a splitter, can create enough downforce to require LMP-range scale static deflections. So that's where the need for those wheel rates comes from.

With a serious downforce body, and/or big wings, the car could probably generate some prodigious accelerations, but it would probably then need frequencies upwards of 25 Hz just to keep the floor pan off the pavement. It isn't so much that the RC car generates dramatically more downforce for its weight; it's more that the suspension deflection caused by the downforce must be so much smaller, due to the limited travel available in such a small car.

Looking at the photos, the roll centers do appear to be a bit below ground, and there appears to be little or no camber recovery in roll – perhaps even negative camber recovery. That could probably be improved upon, unless you are confined to spec parts.

If camber recovery is zero, that means that for roll and lateral load transfer analysis, you can treat the unsprung mass as sprung. It has about the same rate of y-axis (transverse) displacement with respect to the contact patch, per degree of roll, as the sprung mass has. Even if you had 100% camber recovery, and none of the unsprung mass rolled with the sprung mass (for instance, if you had a beam axle), you would not need any different method of calculation than you'd use with the sprung mass c.g. above the unsprung mass c.g.'s.

That lower arm/upper link suspension is similar to the rear suspension on an E-Type Jaguar. When the lower arm pivot axes are horizontal, all longitudinal anti's are zero. When the axes tilt in side view, the side view force line tilts at the same angle. So to add some rear anti-squat, you raise the front mount and lower the rear one. You also get slight changes in the longitudinal anti's with vehicle pitch.

I see there are no friction brakes at the wheels, so I take it that braking is done electronically: you have "inboard brakes".

So the geometric part is simple. However, to calculate percent anti-squat, anti-lift, and anti-dive, you need to know the distribution of the longitudinal ground-plane forces, and there is considerable uncertainty concerning those. You have no center differential, and no front one, but you say there's a rear one. Looking at the photos, I can't imagine how they package that rear one – that looks like a spool – but in any case, when you have fewer than three diffs and four driven wheels, the ground plane forces at the individual wheels depend heavily on how the wheels work with and against each other as grip varies and as the car goes around turns. As a default, you could just calculate percent anti's assuming equal ground plane forces at all wheels, but you should be aware that this may not model actual operation very accurately.

As long as the axes are close to horizontal, the longitudinal anti's are all close to zero anyway. Since no torque acts through the suspension linkage as a whole (although there are torsional loadings on the lower arm), you have only "thrust anti's" and therefore you inevitably do have some wheelbase change with suspension displacement when you have some anti. This isn't necessarily a problem.

You will also have some track change with suspension movement if you have any lateral anti's. For most purposes, a little track increase with suspension compression is good. When the roll center is below ground, you have a bit of track decrease with suspension compression.

As to the effect of droop limiting, it is no different than in a full-size car. When the inside wheel reaches the limiter, the suspension wheel rate goes to near-infinity, and the suspension can only roll by compressing the outside wheel suspension. I have addressed the effects of this at length in some previous newsletters (Sept. 2007, Feb. 2010). I also have an article in the works based on K&C testing of a droop-limited car at Morse Measurements in Salisbury, NC. Stay tuned to *Racecar Engineering* for that.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

October 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ROLL CENTER THEORY

Over the years I have come across a few articles that both state how important the roll center can be towards vehicle handling, and articles that simply state roll centers are overrated and they don't do anything. What's your stand on this topic? I know the basic theories of roll centers and their role in determining the car's geometric stiffness and jacking forces (which raise the car's cg = bad). But, on some cars I've seen the roll center move laterally way outside the wheels when the vehicle rolls. What happens then? Do we simply ignore a roll center's lateral movement and only concentrate on its height? What happens if the roll center's height is off the charts way below ground? When designing pure race cars you can get the roll center where you want, but when working with factory street cars with funny geometry the roll centers start acting wildly. I'm sure making the vehicle stiff enough that we don't have to analyze these movements is the easy answer but am I missing something here?

The topic of roll centers is a godsend to automotive writers. It creates endless controversy, and we'll probably be writing articles about it forever. I've written quite a few, and even produced a video that's mostly about the subject, which is still available.

Here, I will assume that the reader is at least a little familiar with the concepts of front view projected control arms, front view instant centers, and front view force lines (the lines from contact patch centers to instant centers, whose intersection is conventionally taken as the "kinematic roll center").

Quick answer to the question of whether roll centers matter is: yes, but most people don't define them properly, and many don't have a very clear understanding of what we mean by a roll center, and what we use it for when analyzing car behavior.

Slightly longer answer: actually, it's the anti-roll or pro-roll properties of the suspension linkage that matter, and also its anti-pitch or pro-pitch properties. It is potentially useful to describe these properties in terms of roll and pitch centers, but only if those accurately reflect or predict the actual moments generated by the wheel pairs at issue. Taking the force line intersection, or the height of that point, as the roll center, does not do that, except when the force line slopes are equal and opposite – i.e. when they cross in the middle of the car.

The force line slopes matter, and the ground plane forces matter, and the moments generated as a result matter. But you can't analyze the effects by looking at where the force lines cross each other.

It is entirely possible to dispense with the whole concept of roll centers, and some think that's the best approach. However, most people still use the concept.

Use it for what? Properly, as a shorthand, easy-to-visualize way of expressing the relationship between sprung mass lateral inertia force at one end of the car, and the anti-roll moment induced in the suspension linkage in response to that lateral inertia force. It is a notional coupling point between the notional two-wheel suspension system and the sprung mass, for lateral forces: an imaginary roller in a vertical slot.

Since the roller doesn't transmit vertical forces, its lateral location doesn't matter, and can just as well be considered undefined. But its height matters. Its height, H_{rc}, is correctly assigned in terms of predicting geometric anti-roll moment M_{xG} when the sprung mass lateral inertia force F_{ySM}, times H_{rc}, equals a moment equal and opposite to M_{xG}:

$$H_{rc}\,(F_{ySM}) = -\,M_{xG} \qquad (1a)$$

This can also be written:

$$H_{rc}\,(F_{ySM}) + M_{xG} = 0 \qquad (1b)$$

The second form is sort of an algebraic expression of the way a roll center is described in SAE J670e: a point at which we can exert a lateral force on the car and no roll will result. The SAE definition is somewhat problematic, but at least its current version includes a caveat that it is not to be construed as a true instantaneous center of rotation.

With independent suspension, geometric anti-roll comes from positive or negative support forces induced in the linkage upon application of transverse (y axis) force at the tire contact patches. These support forces, F_{zGR} and F_{zGL}, depend on the magnitude of the ground-plane force and the slope of the force line, and nothing else. The moment created by these forces is the difference between them, times half the distance between them – the distance between them being the track width, L_y.

$$M_{xG} = (F_{zGR} - F_{zGL})\,L_y/2 \qquad (2)$$

H_{rc} then has to be the value that makes the moments cancel. Substituting the right side of Equation (2) into Equation (1a):

$$H_{rc}\,(F_{ySM}) = -(F_{zGR} - F_{zGL})\,L_y/2 \tag{3a}$$

Or:

$$H_{rc} = -(F_{zGR} - F_{zGL})\,L_y\,/\,(2F_{ySM}) \tag{3b}$$

Or:

$$H_{rc} = -(F_{zGR}/F_{ySM} - F_{zGL}/F_{ySM})\,L_y\,/\,2 \tag{3c}$$

Note that this implies that the roll center height depends only on the relationship of the right and left induced support forces to the total lateral force, and the track width.

The sprung mass inertia force F_{ySM} is equal and opposite to the sum of the opposing transverse forces at the two contact patches, F_{ySMR} and F_{ySML}.

$$F_{ySM} = F_{ySMR} + F_{ySML} \tag{4}$$

Finally, the induced support forces and the lateral ground-plane forces at each wheel are in the same ratio as the rise and run of the force line. If the angle of elevation of the force line is θ, that ratio is $\tan\theta$. So for right and left wheels, we have $\tan\theta_R$ and $\tan\theta_L$.

$$F_{zGR} = F_{ySMR}\tan\theta_R \tag{5a}$$

$$F_{zGL} = F_{ySML}\tan\theta_L \tag{5b}$$

Substituting into Equation (3c):

$$H_{rc} = -(F_{ySMR}\tan\theta_R\,/F_{ySM} - F_{ySML}\tan\theta_L\,/F_{ySM})\,L_y\,/\,2 \tag{6a}$$

Or:

$$H_{rc} = -\big((F_{ySMR}/F_{ySM})\tan\theta_R - (F_{ySML}/F_{ySM})\tan\theta_L\big)\,L_y\,/\,2 \tag{6b}$$

Or:

$$H_{rc} = \big((F_{ySML}/F_{ySM})\tan\theta_L - (F_{ySMR}/F_{ySM})\tan\theta_R\big)\,L_y\,/\,2 \tag{6c}$$

F_{ySMR}/F_{ySM} and F_{ySML}/F_{ySM} are the portions of total transverse force for the right and left wheels, respectively.

Graphically, we can find H_{rc} in a front view by:

1. constructing a vertical line at a distance from the right wheel equal to that wheel's portion of the total lateral force times the track – I call this the resolution line;
2. finding the intercepts of the right and left wheel force lines with that resolution line;
3. averaging the heights of those intercepts.

We can also find the total jacking force for the wheel pair by adding the support forces:

$$F_{zG} = F_{zGR} + F_{zGL} \qquad (7)$$

In roughly symmetrical independent suspensions, it does hold that more geometric roll resistance implies more upward jacking. However, it is actually possible to have an independent suspension with a roll center above ground (net geometric anti-roll) and downward jacking! This occurs when the majority of the anti-roll force comes from the inside wheel, despite that wheel having lesser ground-plane force. That is possible when the inside wheel has a lot of anti-roll geometry and the outside one has little, or slight negative, anti-roll geometry. This can occur in a strut suspension in a rolled condition.

It is also possible with beam axle suspension to add upward jacking while subtracting geometric roll resistance, and vice versa. This happens in a NASCAR-style stock car when we lower just the right end of the Panhard bar. The rear jacks up in a left turn, yet geometric load transfer is less than if we raised the right end of the bar to make it level. Conversely, if we raise the right end of the bar above level, leaving the left end unchanged, we increase roll center height and geometric load transfer, yet the rear jacks down.

It is quite possible to have a force line intersection outside the car, and toward either the inside or the outside of the turn, and have net pro-roll, or a roll center below ground. It is also possible to have a force line intersection below ground and outside the car toward either the inside or outside of the turn, and have net anti-roll, or a roll center above ground. This does not imply that the geometry doesn't matter. It implies that the force line intersection cannot be taken as the roll center.

Some thinkers note that the jacking force can itself induce roll moments, in the presence of off-center sprung mass c.g.'s or spring splits, and try to come up with ways of assigning roll center height that include spring split and/or c.g. location. I do agree that jacking-induced roll and pitch moments are real, and in some cases big enough to matter. However, they act on the entire mostly-rigid sprung mass, as does elastic roll resistance. They cannot be analyzed in isolation for the front and rear, or right and left in the case of pitch. Therefore it is not analytically productive to try to incorporate them in calculation of geometric properties of the front or rear, or right or left wheel pair suspensions.

Is it possible to spring the car so stiffly that roll centers don't matter, or don't matter much? Yes. But it has to be really stiff – stiff enough so that it essentially has no suspension except the tires. Short of the point where we produce a big go-kart, we do not reduce the geometric component of the load transfer by reducing roll. We reduce roll displacement, but not elastic or geometric roll moment.

It is also possible to spring just one end of the car stiffly enough so that geometric anti-roll doesn't matter much at that end. In that case, roll center height at the other end becomes very important. An instance of that would be a stock car with a coil-bind or stiff bump rubber setup at the front. Once the front is bottomed, there isn't going to be much further elastic load transfer at the rear, because the car can't roll much, but the rear will still have increasing geometric load transfer if lateral load increases, and that will be determined by the Panhard bar height.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

HIGH AND LOW SPEED BALANCE, AND TORQUE ARM LENGTH

I am currently autocrossing a 3rd gen Camaro. I am using a Spohn torque arm. The weight balance on the car is 51% front. The torque arm is the same length as the factory. With a wheelbase of 101 inches, I can't help but feel the length of the arm should be more around 50 inches or half the wheelbase. Is there a way to determined the length? Also, is there a way to figure braking into the equation?

My experience is the stiffer the front, the better the car turns in the more sweeping turns, but in the tighter turns, the car will push. I feel that the long torque upsets the balance of the car. My experience is with Late Model dirt track racing. We used torque arm with a length of about 36-38 inches – however, no sway bars and braking is quite different.

Short answer first: the torque arm isn't the problem, nor is anything else in the suspension, in all likelihood. Most cars, especially rear-drive ones, tend to be too tight (inclined to understeer) in low-speed turns when they're well-balanced in high-speed turns, and be too loose (inclined to oversteer) in high-speed turns when they're well-balanced in low-speed turns.

The tight turn/tight car effect results from the off-tracking of the rear tires inside the fronts in tight turns and outside them in sweepers; the increased power transmitted through the rear tires to maintain steady speed as turn radii and speeds increase; and the the increase in locked axle push effect in tight turns, when any kind of limited slip diff is employed. The effect isn't desirable, but it's normal.

There is really no way to fix this using just suspension and tires. However, rules permitting, it can be fixed with aerodynamics. If the aerodynamic center of pressure can be moved far enough rearward, the car can be given good balance in both high-speed and low-speed turns. In fact, it can even be made to push in the sweepers and be loose in the hairpins, chicanes, and slaloms. All it takes is a fair amount of rear spoiler or wing. If the rules allow that, everything is easy. If the rules preclude that, we are stuck with tuning the suspension around an aero balance problem.

To some degree, we can crutch the problem with relatively stiff low-speed rear damping, but that really just makes the car loose in/tight off, perhaps more so in low-speed turns where transients are more abrupt. That isn't really the same thing as making the car looser slow/tighter fast, steady-state, but it can help the car turn in on the slower segments.

So the immediate problem at hand doesn't relate to the torque arm, but what about the torque arm? How exactly do these devices work, and what does it take to optimize them?

A torque arm is a beam that extends forward from the rear axle, usually alongside the driveshaft, used to react drive torque (M_y) acting about the axle centerline (in side view to the car). The front end of the arm is attached to the sprung structure in a manner that locates it vertically but not horizontally. The arm therefore reacts only axle torque, and locates the axle only rotationally.

The arm will also react braking torque, if the brake calipers are mounted to the axle tubes.

Other members must be provided to locate the axle longitudinally and laterally. The Camaro layout is the simplest possible: two trailing links and a Panhard bar. The trailing links are a bit below the axle, and roughly horizontal, for minimal bump steer. The layout is very compact.

The most common method of locating the front end of the torque arm is a drop link. It is also fairly common to use a rubber or plastic bushing, and simply let the end of the torque arm slide in that. In dirt Late Models, it is customary to create a compliant drop link by using a coilover for the link.

The stock Camaro layout takes simplicity to the extreme by not even having a slider or a drop link – just a really big, soft single bushing. The torque arm itself is a stamped channel-section piece. This works fairly well, but there is some undamped vertical compliance in the bushing, which can potentially lead to wheel hop. Some aftermarket torque arms, such as those from Spohn used by the questioner, substitute a nicer-looking triangulated tubular weldment for the stamped arm, with either a slider (lower-priced systems) or a drop link (higher-priced versions). These attach either to the transmission tailhousing, as the stock arm does, or to the nearby crossmember. There is little or no change in the geometric properties of the system. There is a reduction in vertical compliance at the front end of the arm, with freer motion longitudinally at that point. This reduces compliance axle wrap, which can in some cases become oscillatory and cause wheel hop.

The torque arms in dirt Late Models are deliberately made highly compliant, but the compliance is damped.

The height of a torque arm does not affect its properties, nor does its angle, but its length affects the amount of anti-squat under power and anti-lift under braking. If the longitudinal locating linkage provides near-zero bump steer, the long arm in the Camaro provides something close to 100% anti-squat: the rear suspension does not compress or extend very much under power. The shorter arms used in dirt Late Models provide considerably more than 100% anti-squat: the rear actually lifts noticeably under power. In this form of racing, torque arms are often called lift bars.

It is also possible to get comparable amounts of anti-squat using tension-compression links. One advantage of a torque arm over such options is that it is at least potentially possible to get an anti-squat support force in the suspension that is not sensitive to ride height. This potentially allows somewhat greater amounts of anti-squat to be used before wheel hop is encountered.

Dirt Late Model suspensions often have large amounts of roll oversteer. The rear wheels move forward when the suspension extends and rearward when it compresses. This causes the inside (left) rear wheel to move forward in rightward roll and the outside (right) rear to move rearward, aiming both rear wheels out of the turn. This geometry adds an additional anti-squat/anti-lift effect, sometimes called thrust anti-squat/anti-lift. This component generally does increase as the car lifts, although it is possible to make it not do this.

The Camaro layout also reacts braking torque through the arm. Since the rear brakes only do about 30% of the braking in limit straight-line braking, anti-lift in braking is in the range of thirty to forty percent in that condition. In gentler braking, the rears do a somewhat greater percentage, and the anti-lift is correspondingly greater.

More anti-lift might lead to wheel hop in braking, if the rear wheels come close to the point of lockup. Surprisingly large amounts of anti-lift can be used if rear brake percentage is limited so that the front brakes always lock well before the rears.

In dirt Late Models, it is common to use a lot of rear brake to help the car turn in. To allow this without wheel hop, the brake torque is commonly reacted in a manner that creates pro-lift, roughly countering the thrust anti-lift that comes with the roll oversteer. In the most popular layout – two trailing links and a birdcage left and right, with the calipers mounted to the birdcages – the upper trailing linkon each side is inclined steeply upward at the front, so that the side view instant center is behind the wheel and well below the wheel center.

The torque arm length does not affect cornering balance unless there are asymmetries in the car – spring or wheel rate split, a markedly off-center c.g., or the torque arm itself off-center in the car. Such asymmetries can cause the car to gain (or lose) left-turn wedge under power. They can thus be used to tune entry and exit balance in an oval track car. However, in general we cannot use these tricks in a car that has to turn both right and left, because the wheel load changes we create will affect car behavior oppositely in right and left turns.

My advice with the autocross Camaro would be to: leave the torque arm as is; use plenty of overall roll resistance, since the strut front end has poor camber recovery, especially when lowered; have enough of that roll resistance at the rear (probably use a rear anti-roll bar) so that the car has good balance in the tighter turns; add rear spoiler to stick the rear in the sweepers.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2010

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

HYBRID OR SEMI-INDEPENDENT SUSPENSIONS

I recently ran across an interesting article that you may have seen, as it was in Racecar Engineering in 2001: http://www.locost7.info/files/suspension/dax/. In your June 2008 Chassis Newsletter, you stated that "With a passive independent suspension, it is impossible to have zero camber change in both ride and roll." This system looks like it does a good job of keeping camber very constant in both modes, and I was wondering your take on it, any particular upsides aside from the obvious camber advantage, and any downsides, weight being one of course.

I've started looking into ideas for building a car for the track. I basically want a FSAE car, stretched for the higher speeds of a road course, with more power. My target is 800 lbs, 175whp with a Hayabusa engine, which should be enough to give a good thrashing to the quickest cars that show up each weekend! Without the luxuries of the large budget I'll have to minimize the CNC parts, so I'm looking to keep it simple and easy to fabricate for the build and future repairs. Chassis will likely be square tube, things like that. I'm leaning towards a monoshock on the front with a torsion bar for anti-roll, but I'm not sure if that will be up to the task on the rear until I do some numbers. I saw that camber compensating suspension and wondered if it might be worth including, as every car needs a trick up it's sleeve... I look forward to hearing your feedback on it, thanks!

As may be seen from the linked article, this system superficially resembles a conventional short-and-long-arm (SLA) suspension, and in fact is an optional fitment on the front end of a car that comes standard with SLA front suspension. Unlike the SLA system, however, it mechanically interconnects the upper control arms in a manner that, in cornering, pulls the top of the outside wheel inboard and pushes the top of the inside wheel outboard, providing more camber recovery than the SLA would be able to generate. Indeed, it can have 100% camber recovery in roll, just like a beam axle does (ignoring the component of roll due to tire deflection) – or even more than that if desired. It can, at the same time, also have no camber change at all in ride, again like a beam axle – or if desired, some moderate amount of camber change in ride.

This suspension is one of a large family of layouts that look rather like what we traditionally call independent suspensions, but they actually are not independent in the usual sense. That is, they share two characteristics that we normally associate with beam axles, and not with independent suspensions.

First, they do not have the same camber properties with respect to the sprung mass in both ride and roll. Second, when there is vertical displacement of only one wheel of the pair, and not the sprung mass, the wheel not displaced vertically undergoes camber change, and usually some lateral displacement or scrub at the contact patch.

There is no generally agreed convention regarding how such systems should be classified. They are generally not found on vehicles produced in large numbers, except perhaps if one includes twist beam suspensions in the category (I would do so). I would term this class of systems either *hybrid* or *semi-independent* suspensions. I would therefore stand by my earlier statement about the limitations of passive independent systems.

But semantics aside, what are the possibilities and limitations of these hybrid or semi-independent systems? Are they worth pursuing? Are they mere reinventions of the beam axle, with different appearance and more parts, or do they offer possibilities not available with a beam axle?

First of all, most semi-independent systems package more like independent systems than like beam axles. We don't have to make room for a beam across the car which needs a lot of space around it so it won't hit anything as it moves. Even a twist beam generally needs less space than a true beam axle. Some semi-independent designs connect the upper control arms using hydraulics rather than links. These offer the greatest packaging flexibility, at the expense of added complexity and vulnerability to leakage, cavitation, unintended damping effects, and other ills associated with hydraulics.

Unfortunately, in one packaging-related respect, semi-independents with good camber recovery in roll resemble beam axles more than independents: they need a big wheel well, assuming the car has fenders. Depending on the combination of displacement at the two wheels, the top of the tire can assume a fairly wide range of positions with respect to the body, and the wheel well must accommodate this. With an independent system, the spatial envelope needed to enclose the top of the tire is generally considerably smaller. For an open-wheeled car such as the questioner is contemplating, this is only a factor immediately inboard of the upper portion of the wheels, and it is usually easy to provide ample room there.

A beam axle generally provides zero scrub (zero track change) and zero camber change in ride; 100% camber recovery in roll due to cornering (again, ignoring effects due to tire deflection); and considerable camber change and scrub in roll due to bumps. Some variations may include an axially and/or torsionally compliant slip joint in the beam, and these may have some scrub in ride, and may also have more or less than 100% camber recovery in cornering, but such designs are fairly rare.

Many semi-independent designs can be configured to have a little camber change and scrub in both ride and roll, and lower roll centers than are commonplace with pure beam axles. Whether this actually makes the car faster or more pleasant to drive than it would be with beam axles probably depends on road or track conditions, but the design flexibility can probably be regarded as an advantage, rather than a minus.

A SEMI-INDEPENDENT DESIGN USING SUPERIMPOSED SUSPENSION ELEMENTS

Sent with this newsletter is an inventor's sketch of a semi-independent design superimposing one traditional suspension configuration on another, the two acting in series. The system that directly connects to the wheels is similar to a Ford truck twin I-beam layout (modified swing axle). The system that connects directly to the sprung mass is a sliding pillar configuration, roughly as used on Morgans. Inventor Gil Scigalski of Illinois sent me a description, and asked for a free opinion, agreeing that the material could be used for publication.

Mr. Scigalski writes:

In its general layout the design resembles a Ford Twin I Beam suspension, but with the characteristic long overlapping swing axles pivoting from unsprung rather than sprung parts of the vehicle. The unsprung parts consist of vertically aligned chassis mounted struts similar to the sliding pillars used on Morgan sports cars. In addition to transversely locating its associated swing axle, each strut also connects to and moves with its opposite axle via a slot formed therein. Each slotted connection accommodates the lateral difference between the vertical path of the strut and the arching path of the swing axle during jounce-rebound. Longitudinal location and springing-damping of the swing axles are not shown but inferred and accomplished by means known to the art.

After pondering this layout for a while, I have concluded that it is roughly equivalent to a beam axle with a Panhard bar or other lateral locating mechanism at the height of the swing axle pivots. Or more precisely, it is equivalent to a slip-jointed beam axle, with its ends hung on sliding pillars or trailing arms and a roll center at the height of the swing arm pivots.

The inventor claims that the suspension will completely eliminate roll, up to the point of tire breakaway. Actually, tire breakaway or no, it will generate geometric anti-roll, but not 100% unless the roll center is raised to unsprung mass c.g. height. As with a beam axle, this will result in large amounts of scrub on one-wheel bumps. The inventor is correct that the system eliminates the jacking that would be encountered when trying to achieve anything like 100% geometric anti-roll using pure swing axles or any other independent system. But again, this merely duplicates what a beam axle can do.

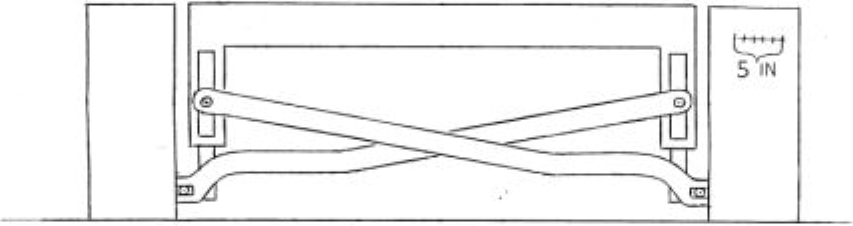
5 IN


GSX
HOLE SHOT
RACING

CAMBER Nectar

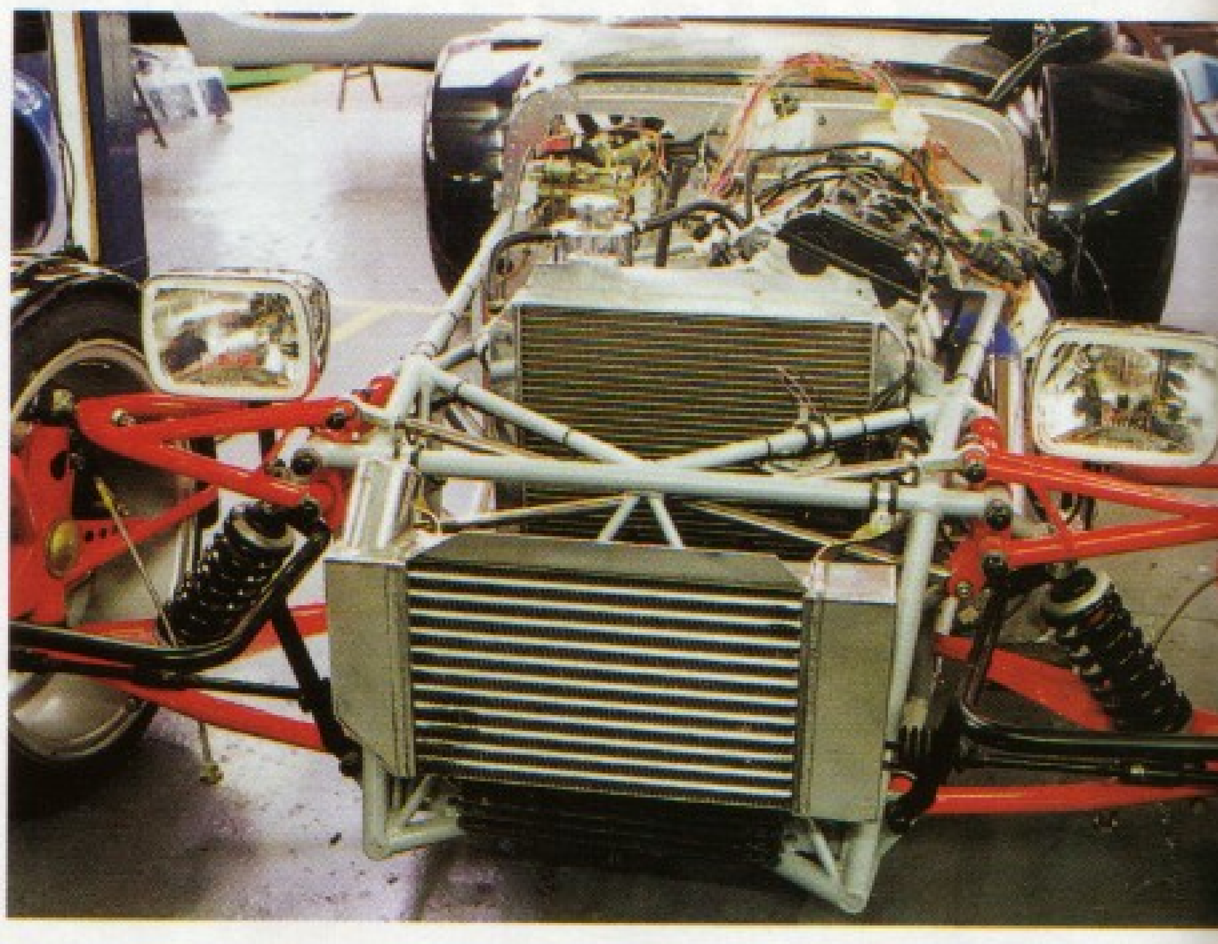

Kit car company, Dax, has developed an imaginative approach to combating camber change on wishbone independent suspension

Words/Photos	Simon McBeath
Illustrations	Jim Bamber

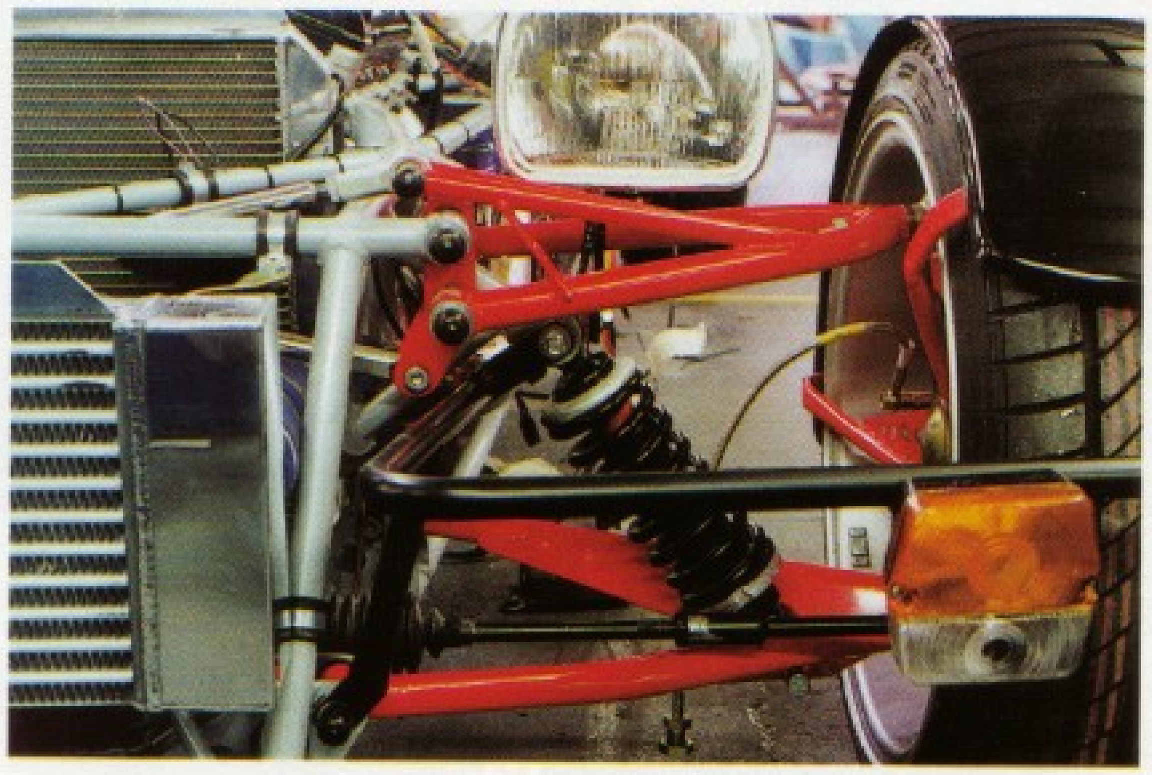

A new suspension system that contributed to a driver securing eight pole positions in eight races and five race wins, as well as substantially reducing lap times, surely has to be worth a second look. Andrew Sterling's Dax Rush Rover V8 competes in Class A (engines from 1661cc to 3651cc with two carburettor chokes) of the 750 Motor Club's Kit Car Championship. After starting the 2000 season with reasonably good places his performances were transformed once Dax designer Peter Walker's new suspension system, for which world-wide patent applications have been filed, was fitted.

The Kit Car championship is for front-engined, fully road-legal kit cars running on road tyres. The popular tyres tend, not surprisingly, to be soft compound, low aspect ratio types. Peter Walker asserts that 'with wider, flatter tyres it's even more important to control camber

Top: The suspension as fitted to a Dax Rush showing the diagonal cross-links

Above: the suspension looks conventional apart from the upper wishbone and its mounting

Right: The upper wishbone pivots are mounted on a pendant link hanging from the chassis mounts

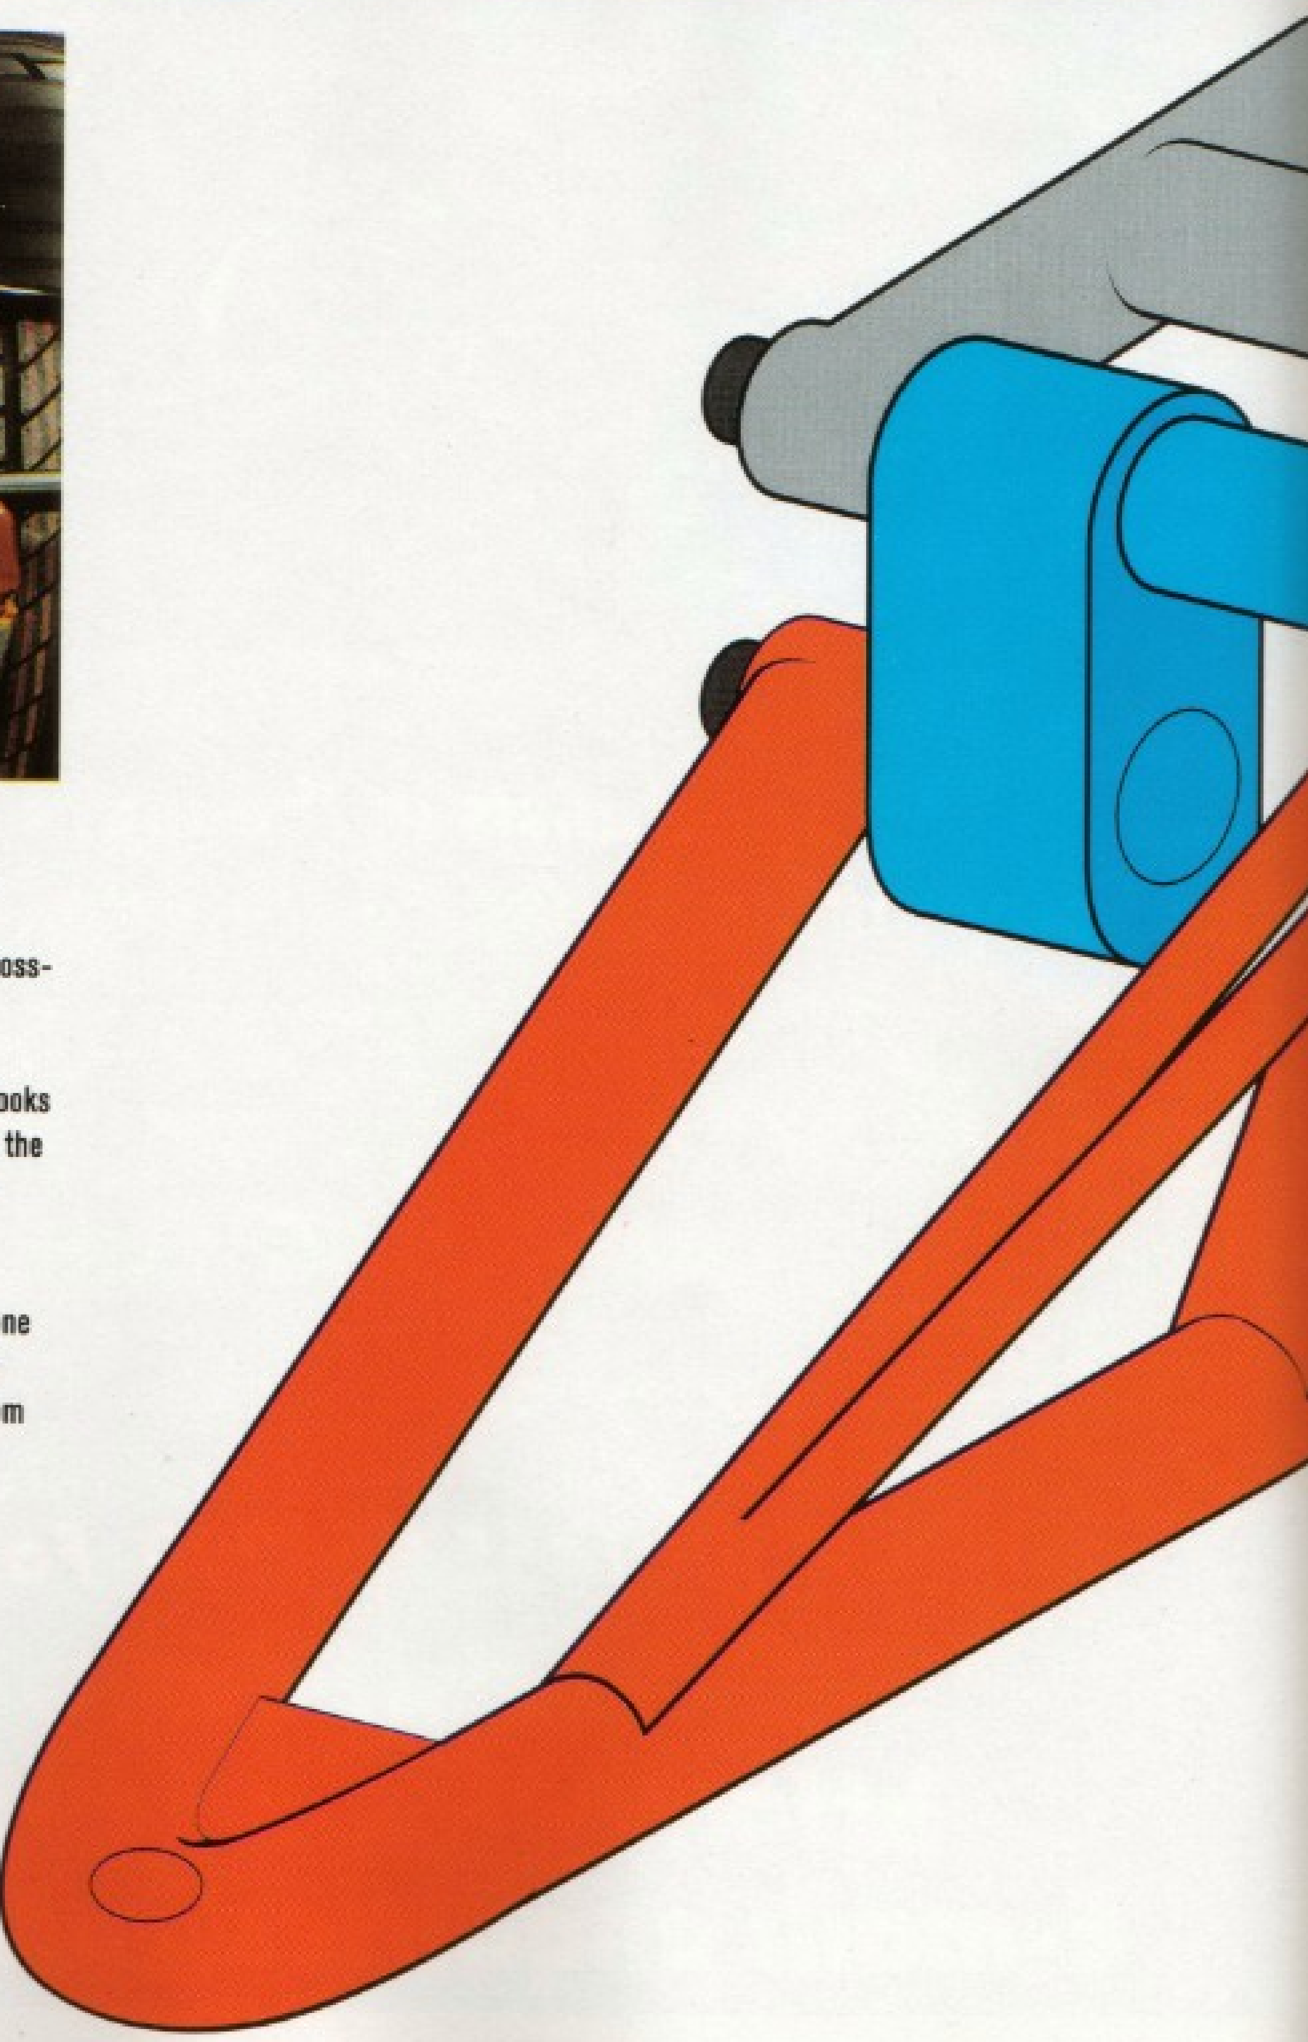

> "The cross-links govern the changing positions of the wishbone pivot points during suspension movement"

(change). It's almost impossible to achieve complete camber compensation with conventional unequal wishbones. You can change camber to largely compensate for roll, but this will then act against you in bump or jounce. This leads to uneven tyre temperatures and friction levels across the tread and can mean changing, and often less than optimum, grip levels during the transition from, say, cornering to acceleration.'

Peter continues: 'Whilst on holiday around three years ago I was pondering the inherent problems with unequal length double wishbone designs and gaining the maximum benefit from wide tyres, and the solution just popped into my head. I first thought of a system using the lower links, for which I have been granted a patent. But using the top links allows rubber-mounted lower arms, which is better for noise, vibration and harshness in the road-car case. The top links are all mounted in PU bushes.'

The new front suspension on the Dax Rush utilises unequal length, convergent wishbones, and 'around 50% of the camber compensation in roll comes from this means' remarks Peter, 'while the other 50% is provided by the action of the system's links'. It's the inboard upper wishbone mounts which are not conventional, and instead pivot halfway down 'pendant links' which themselves hinge from the chassis mounting points. There are also lever arms connected to the wishbones above and slightly inboard of the chassis pivots that couple via cross-links to the bottom of the →

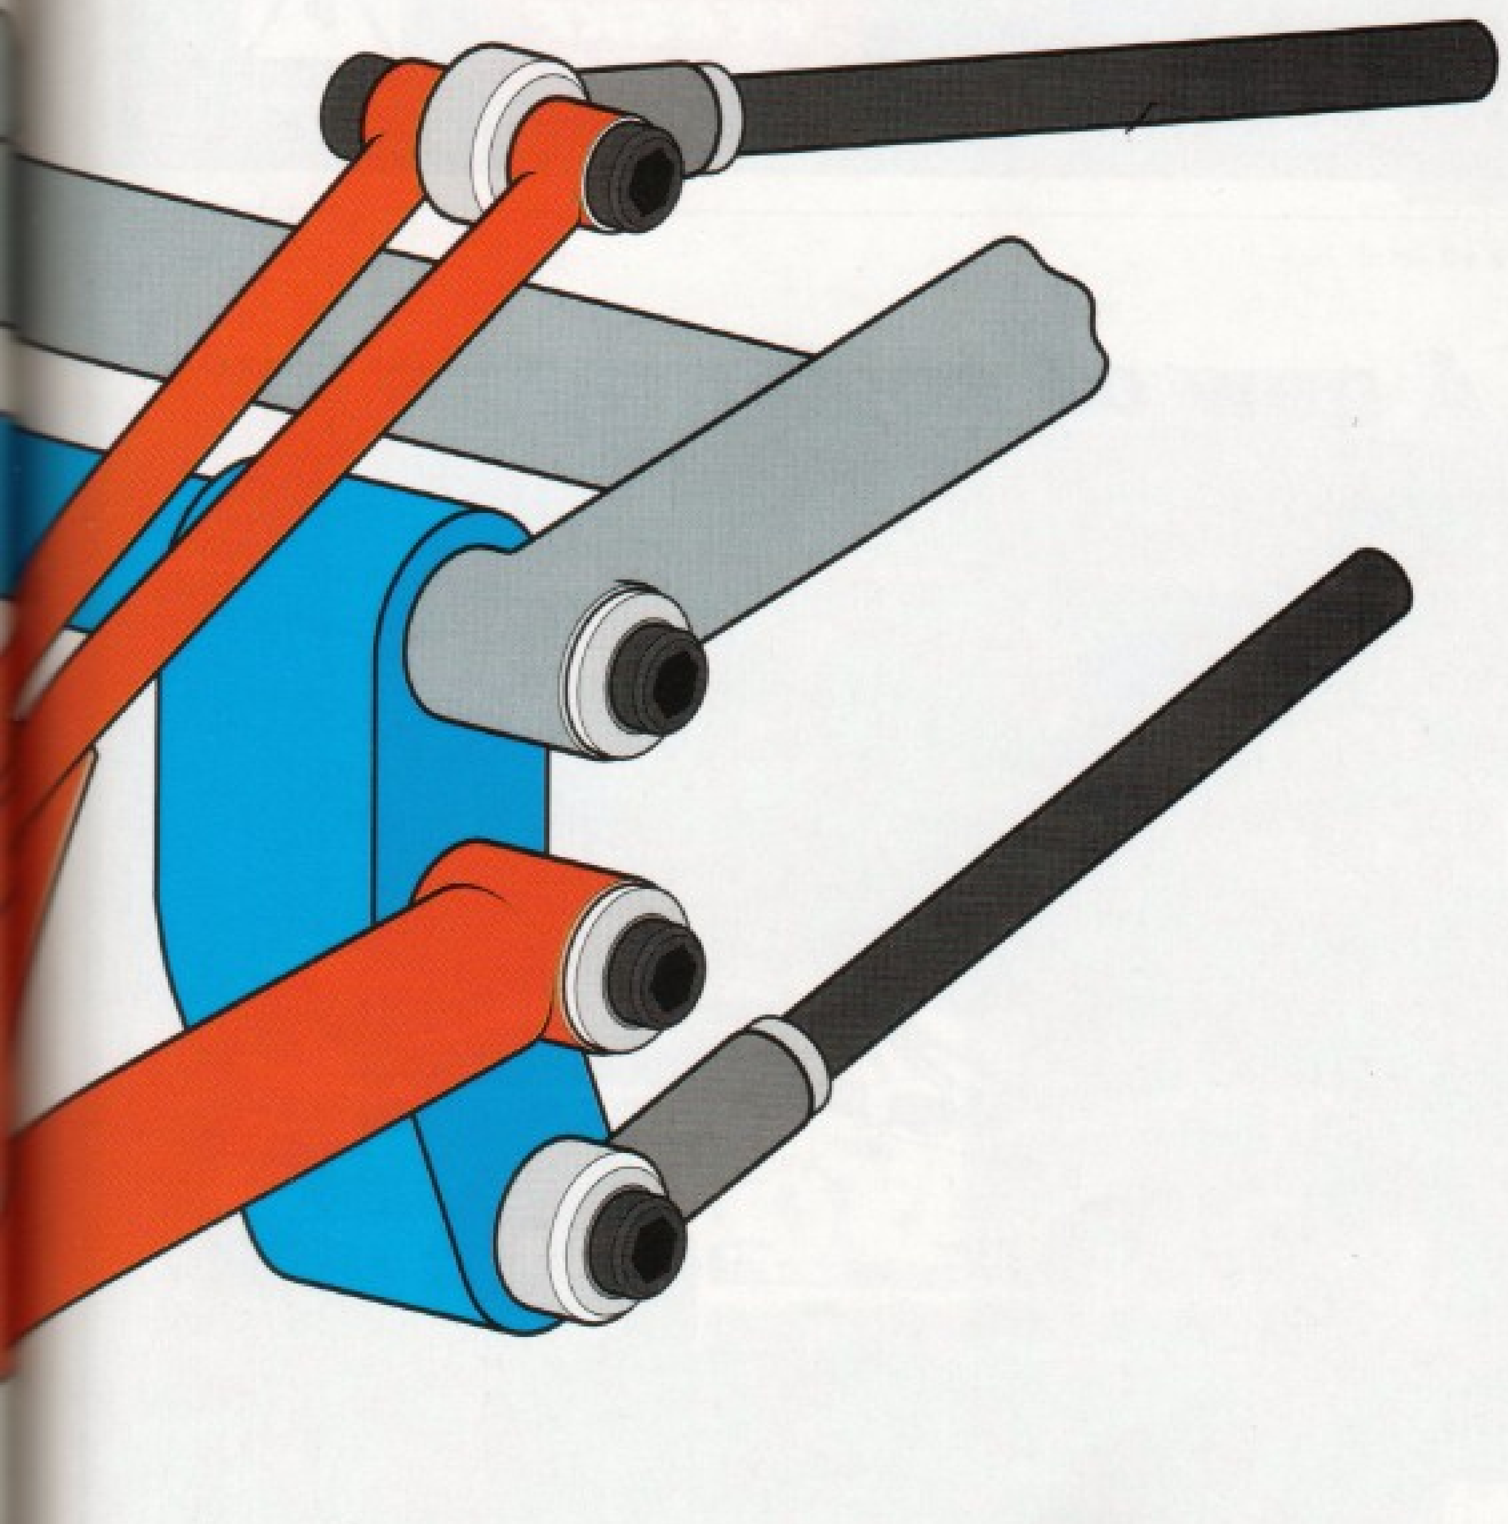

Bennett's System

The problem of maintaining wheel camber control in bump, jounce and roll is practically as old as independent suspension itself. For example, in the mid-1960s a novel design was produced by Torix Bennett on the Fairthorpe TX1 that attempted to combat the problem. Tyres may have been a lot different in those days, but Bennett's basic criterion was the same as it is now – that a loaded wheel in cornering should not be allowed to take up positive camber but, instead, should be kept vertical or with slight negative camber.

His design bore physical similarities to Peter Walker's system, as used on the Dax, in that diagonal cross-links transferred wheel camber movement across the car. Furthermore, the lower cross-link connections were hung on pendant links hinging from the trailing arms in a manner reminiscent of the links on the Dax upper wishbones. The basic layout differed on the Fairthorpe in that simple lower trailing arms were used rather than wishbones, and the top of the upright on one side connected via the cross-links to the trailing arm on the other. Adjustment of link lengths and pivot positions provided the Fairthorpe scheme with adaptability that enabled camber movement to be controlled to varying degrees (no pun intended).

The geometry, like that of the Walker's Dax layout, was also said to offer intrinsic anti-roll forces and this, together with a high roll centre, minimised roll angles. Drivers of the car commented on the excellent grip and traction offered by the system.

One of the inherent disadvantages of Bennett's design was the space required by the cross links and the scheme may also have been suspect in one-wheel bump because vertical deflection of one trailing arm would, via the cross-links, have induced positive camber in the opposite wheel.

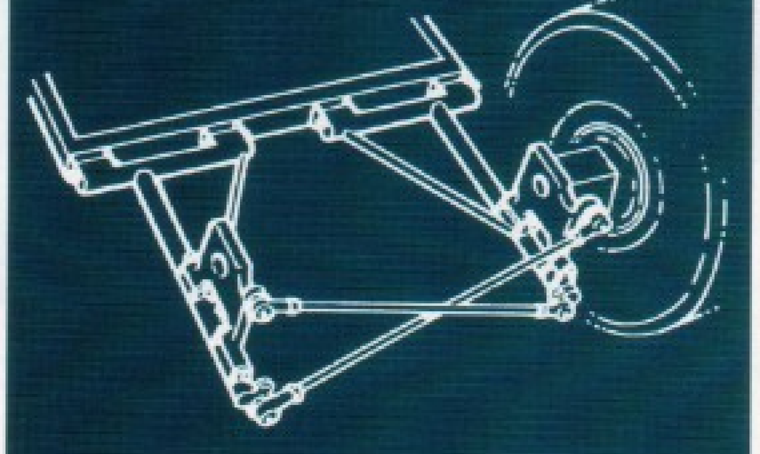

A caption goes here like this A caption goes here like

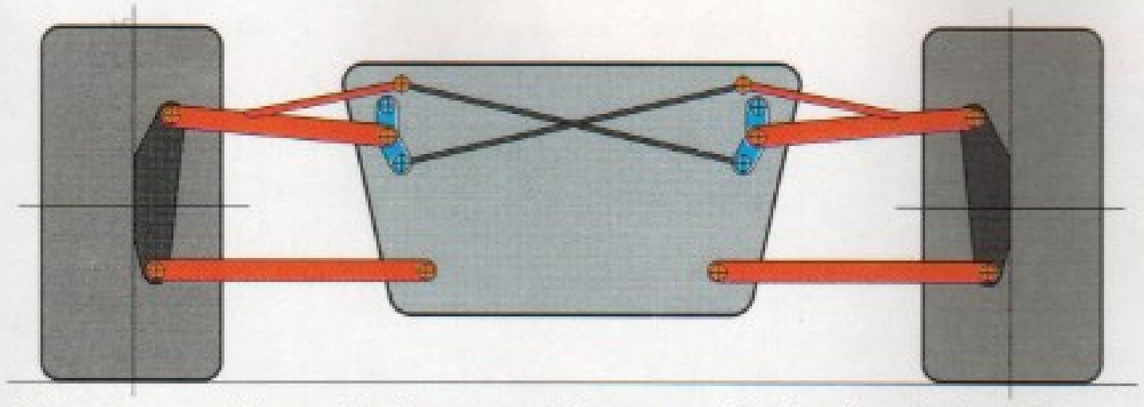

Static, the cross-links position the upper wishbone mounts to keep the wheels vertical

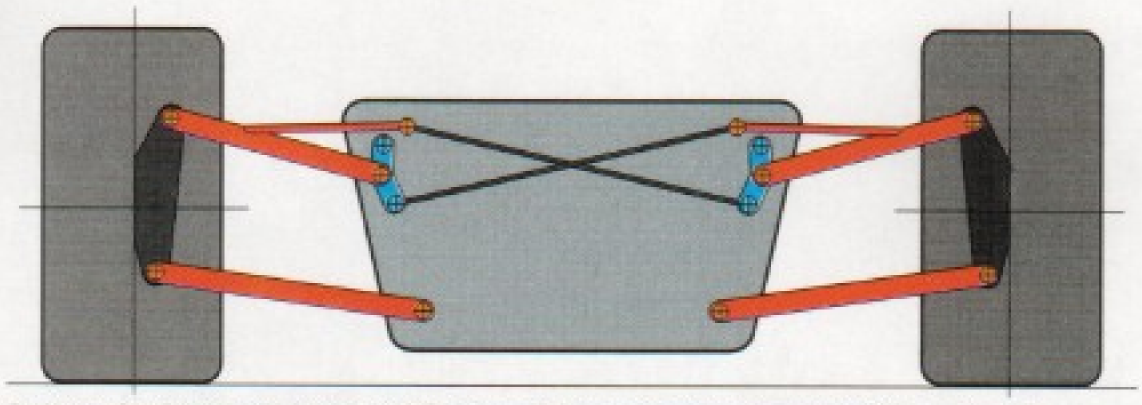

In two-wheel bump the upper wishbones are pushed outwards preventing camber change

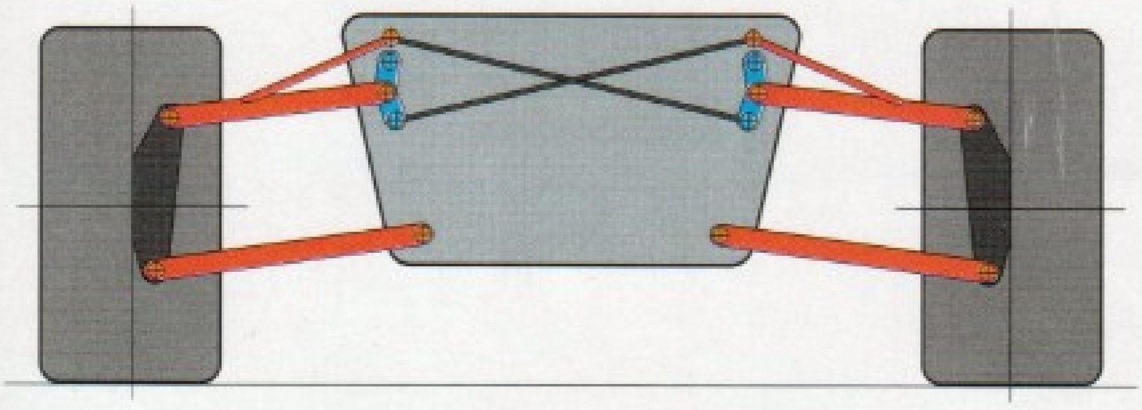

Likewise in droop wishbones are pulled inwards again preventing camber change

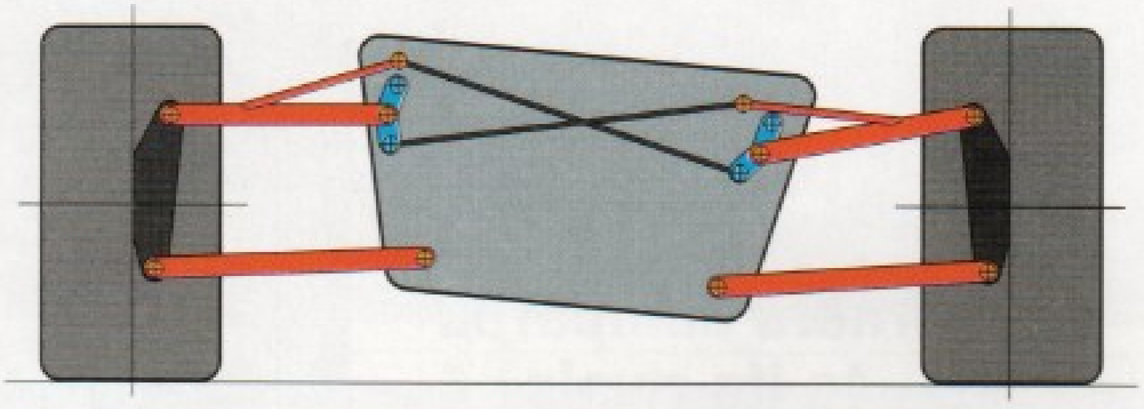

In roll the wishbones have opposite effects on each other still maintaining zero camber

"Whilst on holiday three years ago I was pondering the problems"

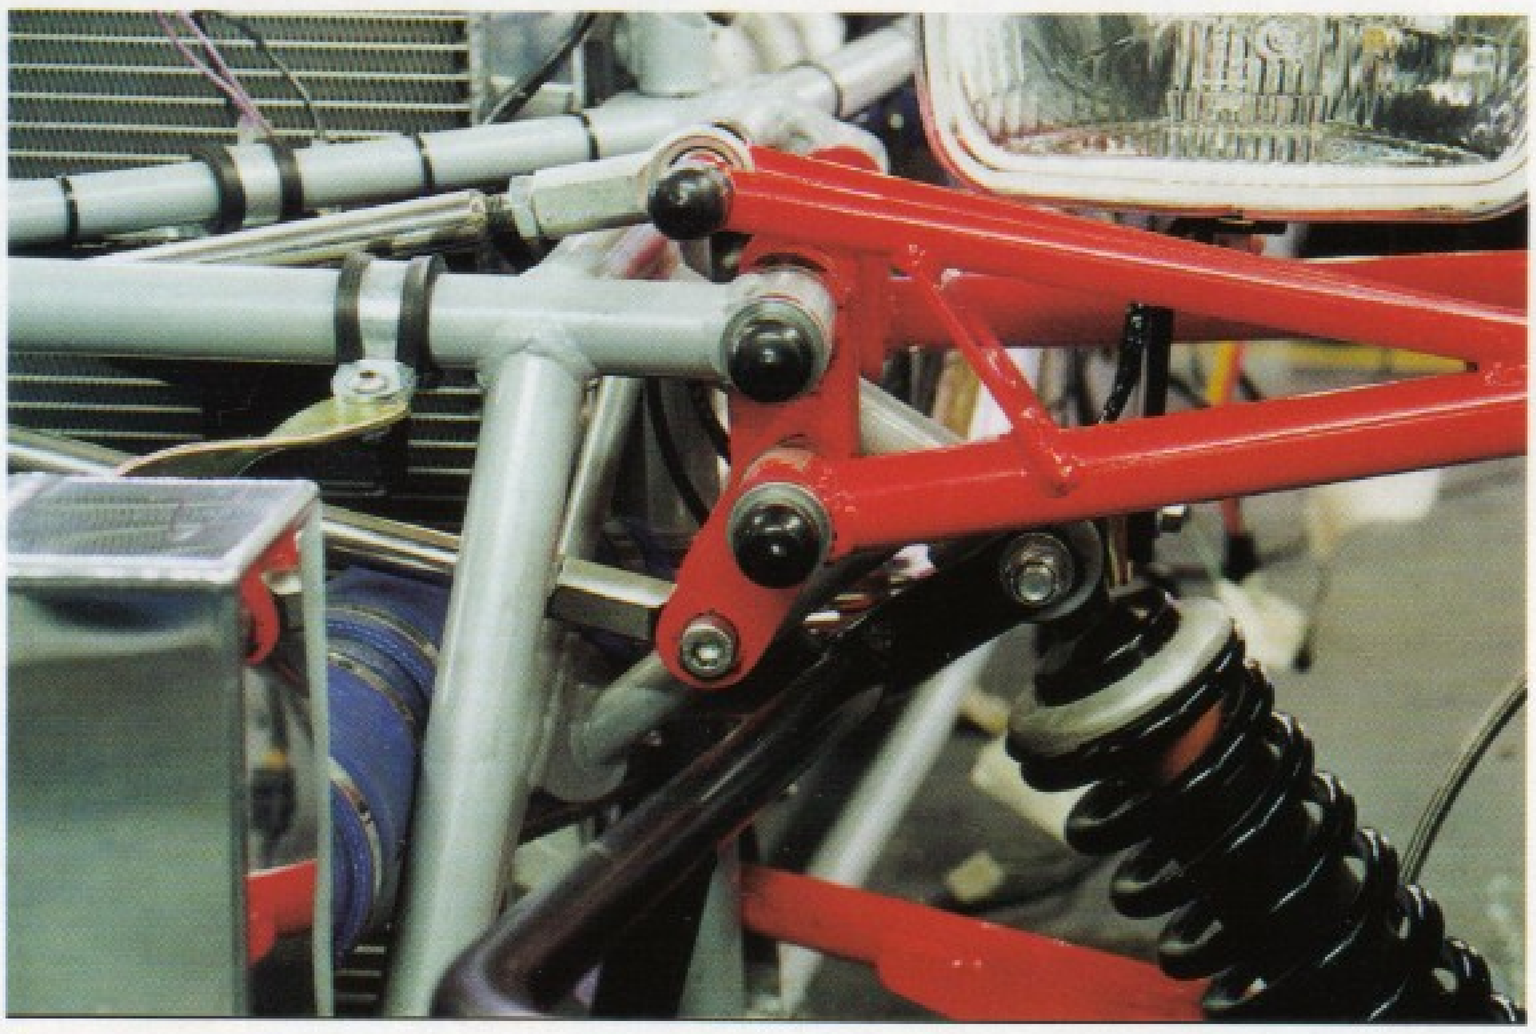

System is simple in operation and picks up neatly on chassis mounts similarly positioned to a conventional suspension. However, the steering has to be in the same plane as the lower wishbones which required special uprights on the Dax

pendant links on the opposite side of the chassis. So the system uses moveable top wishbone pivots, and also couples one side of the suspension to the other. The precise geometry was evolved by Peter through calculations and hardboard models which enabled him to establish the requisite link lengths, angles and pivot points.

The cross-links govern the changing positions of the wishbone pivot points during suspension movement. Furthermore, the way the geometry is set means that only 50% of the camber change normally encountered is induced anyway. Thus, in the two-wheel bump case when conventional unequal length, convergent wishbones would be inducing (increased) negative camber, the cross-links counteract the usual change by moving the top wishbones outwards.

Although, as stated, a conventional system can control camber change in roll up to a point, this new system counteracts any tendency to change to positive camber in roll. The cross-links pull the inboard top wishbone pivot (on the loaded side of the car) inwards, which pulls the top of the upright inwards in a calculated way to keep the wheel vertical (or at a datum angle). Meanwhile, the cross-linked partially unweighted inner wheel also remains vertical.

Compared to a conventional wishbone set up, which Peter says might give 'three to four degrees of camber change over the range of suspension travel, the new system keeps the wheels vertical or very nearly vertical over five inches of travel in roll and two-wheel bump or jounce, and in all transient cases in between.'

The only situation where the system might be considered less able than a conventional fully independent system is the one-wheel bump case. Similar to the roll situation, in one-wheel bump the wheel that rides the bump is forced upwards, inducing some negative camber which transmits to the opposite side of the car as positive camber. Whilst this may not be a disadvantage in straight line running, one could envisage a situation on a race track when an inside wheel runs over a kerb whilst cornering. Assuming reasonably compliant springing, the resultant positive camber on the outside (more loaded) wheel could cause a loss of grip. Peter counters that saying 'this is an instantaneous case, in which, because of the way the system works, only 50% of the camber change normally associated with suspension movement would be triggered. In reality it would probably not be much different to a stiff racing suspension which may cause the whole car to tip in reaction to the wheel hitting the kerb.'

"Overall the car is said to look very 'flat' in corners compared to its rivals"

The system is not restricted to camber compensation. According to Peter his system allows a higher roll centre than with a normal double wishbone set-up 'without the normally unavoidable increase in scrub' (he states that the roll centre is at the same height as the lower wishbone). This in itself helps to reduce roll angles for the same cornering load because the roll centre is nearer the centre of gravity height. Peter also states that 'the cross-links transfer an inherent anti-roll force which occurs when a side load is applied to each tyre. As cornering forces try to make the car lean, the reaction through the cross-links opposes that effect and keeps the car upright.'

Overall the car is said to look very 'flat' in corners compared to its rivals with conventional anti-roll systems and, although no quantitative measurements of tyre temperature profiles are available, the tyres apparently wear very evenly across the tread.

When applied to the front end of a vehicle, the camber change induced by the system requires the steering track rod to be in the same plane and the same length as the lower wishbone. This ensures that there are no toe changes during suspension movement and, in the case of the Dax car, dictated a new design for the front uprights.

Peter says his design could be applicable to any car but that it is largely aimed at performance cars on low profile tyres where keeping the wheels upright will be to the potential betterment of cornering, braking and traction. Good ride quality and a lack of roll can also be maintained. The steering on any car employing the system on its front suspension would obviously need to comply with the parameters mentioned above, but it was in independent rear suspension format that the system was first developed. It is doubtful however that racecars with limited suspension movement would have the travel to gain significant benefit from camber compensation, and packaging would be difficult in many installations.

It seems there is already interest from 'elsewhere' in Peter's design and 'Camber compensation and Anti-roll suspension' is now being offered as an option on the Dax Rush, of which the company produces about 60 units annually. Based on Ford Sierra running gear, the addition of Peter Walker's suspension system is said to provide the Dax Rush with a 'nice overall handling compromise'.

Good camber control in roll, bump and jounce has always been something of a Holy Grail to suspension designers. However, this system does seem to do an excellent job over a large range of suspension movement.

DAX's Sterling Work

Andy Sterling had achieved race wins in his potent Dax Rush prior to the fitment of the camber-compensation and anti-roll suspension. However he states unequivocally that it is much easier to drive the car on the limit with the new system.

'The car used to be quite a handful on the limit' he says, 'which was partly because we had to run very stiff springs previously. The new suspension has allowed us to run much softer springs, yet the car now has greatly improved grip. The benefits are particularly good on tight, twisty circuits and the car is quicker where there are rapid changes of direction. I would say the system is worth around one to two seconds a lap on tighter circuits.'

The Dax uses the popular Yokohama A032-R rubber, successor to the A008-R and Andy confirms 'The tyres do now wear very evenly across the tread, where before you would see excess wear on the shoulders. And grip in the wet is also phenomenal. I honestly think Peter Walker's design is an exceptional bit of equipment. There are not many modifications that you can make that would produce the same gain in lap time. And it has made it relatively easy to go quickly. I'd won races before, but I was able to set pole positions and fastest laps with the new suspension. The design really does work.'

High praise indeed and, when you consider that he has finished second in class in the Kit Car Championship for the last two years, Andy Sterling's confidence heading into the new season, with the camber compensation system now well sorted, is high.

The Dax camber compensation suspension has been proving itself in racing and real performance benefits are claimed

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

COIL-BIND AND BUMP STOP SETUPS IN STOCK CARS

1) When Ron Hornaday's Camping World (Craftsman) Truck broke a swaybar arm at one of the races this season, the left side of the truck was noticeably higher than before. This could indicate a setup that included a highly preloaded, very-stiff front swaybar. Is there an advantage to a highly preloaded swaybar? Is it done to increase LF spring loading for a LF coilbind setup?

2) With today's front soft-spring Late Model setups, some racers use front bump stops for the main purpose of keeping the chassis from bottoming out on the track (using a stiff front bar), while other racers are hard onto their front bump stops as their primary suspension setup (using a smaller front bar, using various methods for swaybar preload). A lot of these setups seem to be trial and error, without knowing what's really physically happening. There's not much information in the literature discussing the vehicle dynamics, load transfer, tuning methodology, etc. for bump stops systems. Seems it's not well understood why it sometimes works as well as does, nor what happens during a race with these setups as the tire-to-track friction coefficient changes, or the racing groove changes forcing the use of a different bank angle. What's your explanation of the physics behind the use of bump stop setups: a) for the urethane rubbers; b) for spring coilbind systems.

The primary purpose of these setups in stock cars is to work the ground clearance rules. It has been a longstanding tradition that stock car tech inspection includes a requirement that the car pass over a barrier that serves as a ground clearance gauge. This comes from the days when stock cars were much closer to stock than they are now, and there was a desire to have the racing vehicles bear some resemblance to actual road cars.

Without the ground clearance rules, there would be no reason to use bump stop setups.

One could argue that sanctioning bodies should simply abandon the ground clearance rules, or just set the requirement very low. This would definitely simplify setup, and probably lead to closer competition. However, there is also a case for leaving the rules as they are, to make computer-aided simulation and modeling so difficult that most racers will have to rely on brain-aided engineering instead.

The exact height of the ground clearance gauge varies among classes and sanctioning bodies, but it is generally considerably higher than we want to run the car on the track, particularly at the front. We would like the valance or splitter to barely clear the track, at all times. If we can't get that condition at all times, we would at least like to obtain it in the turns.

The idea, then, is to have a car that is very soft in ride at the front, but hits some form of stop when it gets to the height we would actually like it to run. We then use some combination of aerodynamic downforce, hold-down shock valving, and turn banking to keep the suspension compressed to the stops, at least when it counts most.

We thus have a car whose wheel loads and wheel rates are dramatically different when it's down on the stops than when we set it up in the shop. We can set the car at the on-the-stops ride heights in the shop, and see what cambers and alignment settings we get, but we can't faithfully reproduce the aerodynamic loadings, banking loadings, and damper forces. If we are a well-funded team with very good engineers and adequate wind tunnel time, we can simulate the loadings at a particular speed, turn radius, and banking angle. We can use the calculated or recorded aero forces and x, y, and z accelerations as inputs for kinematics and compliance testing, and get fairly realistic wheel loading measurements, provided we also do a decent job of estimating tire ground plane force distribution. Even this doesn't fully capture the damper forces, but it's useful, especially for comparing effects of changes everywhere but the dampers.

However, if we are a Saturday-night Late Model team, we generally won't have any budget for wind tunnel or K&C testing. We will be relying on qualitative understanding of the car's dynamics, and dialing the car in by degrees using that qualitative understanding. Fortunately, we will also generally be running one track repeatedly, or a small number of tracks, and will also be allowed free test-and-tune days and relatively cheap track rental during the week, and there will be no rule limiting how much we can test. That means we can do lot with qualitative understanding, if it's reasonably good qualitative understanding.

It is useful to think of the car on the bump stops as if it were a static setting, with some added loads involved which we don't precisely know.

Returning to the question posed first, yes I would say the observation described does imply a preloaded front sway bar, assuming nothing else got broken or damaged or moved. And the questioner iscorrect in supposing that this has an effect on the behavior of a coil-bind setup.

For a given set of static ride heights and wheel loads, a setup with a preloaded bar will have the left front spring compressed more at static, and the right front spring less compressed. This means that as the car comes down to coil bind, the left spring will hit coil bind first, or at least earlier than it otherwise would. The car will have a more de-wedged condition at the point of right-side coil bind than it otherwise would. This should translate to a freer car.

If the front bar is quite stiff, it is possible that the car might enter and exit the turns with only the left front actually coil-bound, and the bar holding the right front up. This would allow the car to have more roll compliance in the front end than it would if coil-bound on both sides, making it turn better in the middle. It would also be relatively left-stiff in pitch, which helps it put power down on exit.

We should note that the same will normally not apply to a bump rubber setup, where the springs do not coil-bind. The left spring will be compressed more at static, just as in the coil-bind case, but the shock will not. Therefore, the bump rubber will be no closer to bottoming out than it would be without bar preload. If we want a comparable effect with bump rubbers, we need to lengthen or pack the left rubber.

As to the pros and cons of stiff rubbers versus soft ones, that's a bit like the tradeoff between stiff and soft setups generally. Stiff ones give better control of aerodynamics but poorer absorption of bumps, and very high elastic roll resistance at the front, which often leads to a mid-turn push.

One sometimes hears it said that the front of a bump-stop car has to be stiff in roll to keep the left front corner down and the splitter or valance the proper distance from the track on both sides. Actually, it is only necessary that the car as a whole have adequate overall roll resistance. We can keep the front of the body from rolling using the rear suspension. The car doesn't have a swivel in the middle. The front needs to be soft, and then stiff beyond a certain displacement, in ride, but it doesn't need to have any more stiffness in roll than is required to meet the ride requirement. The front/rear roll stiffness distribution can be whatever it needs to be to make the car turn in the middle and come off strong.

Am I suggesting a rear anti-roll bar? Yes, if the rules allow it. It is allowed in the upper NASCAR divisions, and on Late Models by many sanctioning bodies. However, NASCAR does not allow it on Late Models. When a rear anti-roll bar is not permitted, one has to get roll resistance with the springs and the Panhard bar. Doing it with the springs has the advantage of keeping the tail a bit higher, which adds aerodynamic downforce. It has the disadvantage of making the rear end stiff in ride, hurting ability to ride bumps.

On tracks that are banked steeply enough to make the left rear compress in the turns, it is common to use a right-stiff spring split at the rear. This does have the advantage of de-wedging the car as it compresses on the banking, but it has the disadvantage of making the car right-stiff in pitch. That makes it harder to get the car tight enough on exit, when it's free enough in the middle. In some cases, however, this approach may be indicated, as a way of getting desired wheel loads mid-turn, without a ridiculously stiff wheel rate in ride at the rear.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

SWAY BAR/SPRING RATE EQUIVALENCY

Do swaybars reduce mechanical grip?

From what I understand, theoretically a swaybar is supposed to act just like a spring but only reduce body roll when a car experiences lateral forces. I also believe most swaybars are linear in rate. Is this true? Therefore, if we compare two identical cars in all aspects except:

Car A has:
200lb/in front springs and a 200lb/in swaybar rate
300lb/in rear springs and a 300lb/in swaybar rate

Car B has:
400lb/in front springs and no swaybar
600lb/in rear springs and no swaybar

Since both Car A and B have the same amount of roll resistance and the same ratio of front/rear roll stiffness, both cars should handle identically and achieve the same lateral acceleration on a smooth surface correct? I ask this because I've come across a post made by the 2010 SCCA DP Champion (he has one other years in various classes but I don't remember them) and he states that swaybars do reduce mechanical grip. If this is true, how would they do this and where would the static load go under cornering? Here's his post, and here is the link to the message board: *http://www.mr2oc.com/showthread.php?t=298272&page=2&pp=30*

"Installment #4:

"Now that we understand the function of the springs and how they interact with the body to determine ride quality and body roll we can now directly address the swaybars. To achieve our target body roll and roll ratio we have chosen specific wheel rates. However, the rates we have chosen would result in a very stiff ride if we used only springs to achieve that wheel rate. To soften

the ride we need to lower the spring rates, especially in the front (to achieve a positive favored speed) however, that will upset our chosen roll ratio so we add back wheel rate by supplementing spring rate with swaybar rate.

"From installment #2 we had chosen 550 lbs/in front and 400 lbs/in rear spring rates. Its easy to now choose a front spring rate that will produce a more comfortable ride. Say 300 lbs/in front springs would give us a positive favored speed and move the center of suspension closer to the CG of the car, just aft of the center point between the two axles. We can then add back the spring rate by installing a swaybar with a rate of 250 lbs/in.

250 lbs/in front bar + 300 lbs/in spring = total rate of 550 lbs/in.

"That looks pretty easy. However, it's hard to find an off the shelf swaybar that has exactly the rate you want. In reality, unless you want to fabricate a custom bar every time you want to test a different setup, you need to calculate the rates of the available bars and then determine how much spring rate you need to achieve the desired total rate. So lets say you measure your existing swaybar and find it has a rate of 200 lbs/in. The spring rate you would want is determined by subtracting the existing bar rate from the target spring rate.

550 lbs/in target rate – 200 lbs/in bar = 350 lbs/in springs.

"Easy enough. But we have a problem here. Swaybars are not the dynamic equivalent of springs. A swaybar transfers load from the inside tire to the outside tire and thus reduce mechanical grip as they add spring rate. And the effect is not linear. The stiffer the bar in comparison to the springs, the greater the effect (loss of mechanical grip). To give an example of the effect, if we set our proposed STS2 car up using the target data we have assumed above using the target spring rates without any swaybar, the car should have good balance. However, if we achieved that same target spring rate using a front swaybar, the car would tend to understeer more than if we used only springs and no swaybar. And the greater percentage of the total front spring rate the bar accounted for, the more the car would understeer. I noted as much early in the thread.

"We now must choose how much bar we want to use for our final setup. As the reader may know, I don't use swaybars on my DP car. Nor did I use swaybars on my previous racecar, a DSP X1/9. For me it is far easier to manage the setup of the car without swaybars. I also prefer the feel of the car without swaybars. It has long been my thought that; because a swaybar reduces mechanical grip, why would you want to put anything on the car that reduces mechanical grip?

"With this simple method, it is easy to compare the effect of the front bar by comparing the same total spring rate using just springs and no front bar to the same total spring rate incorporating a front bar. I have done extensive testing and have proven to my own satisfaction that the theory is in fact accurate. The same total front spring rate achieved using a front swaybar will produce more understeer than the same total front spring rate achieved using springs only. In addition, the effect of the front bar changes based on the level of grip the surface offers. As a result, the car does not have

consistent balance from surface to surface or even from run to run as the tires heat up from and the surface cleans and heat up throughout the day. I have found that my no-swaybar setup is very consistent on different surfaces and conditions seldom if ever requiring any changes to setup. At most, a pound or two of air pressure is all that is needed to tune the balance even in the most extreme of conditions. In fact, I don't even change the setup for rain. All I have to do is bolt on the rain tires and the car is fine.

"If one chooses to use a front swaybar, the effect resulting from the loss of mechanical grip will have to be accounted for by softening the front springs enough to bring the balance back to neutral. Choosing the amount of swaybar to use is now easy and dependant on driver taste. If the driver prefers a smoother/softer ride, use a very stiff front swaybar and subtract the front bar rate from the total spring rate to determine the required front spring rate. Testing can then determine how much less spring rate is necessary to bring the handling balance back to neutral. One could also compromise and use a very soft front bar, thus minimizing the loss of mechanical grip.

Sway bars (anti-roll bars) are in fact dynamically equivalent in roll to springs; they are generally linear; and they do not transfer weight or wheel load any differently than ride springs.

However, many people get confused about this because of a subtlety in measuring and expressing rate of a sway bar, and evaluating its roll resistance equivalency to ride springs. This causes people to think they are adding one amount of roll resistance with the bar, when in fact they are adding twice as much as they think. Of course the effect on the car reflects the difference, and they then erroneously conclude that the bar must work in some fundamentally different manner in roll than the ride springs do.

Ordinarily, a ride spring has one end (often the top) that is fixed with respect to the frame, and one end (often the bottom) that moves with respect to the frame. An inch of movement or displacement of the spring is simply an inch of movement at the end that moves. Simple. No way to get confused about that.

But a sway bar has a middle portion that is fixed with respect to the frame, and two ends that move. In roll, the ends move oppositionally. So, what then is an inch of movement for the device as a whole? Is it an inch of movement at each end relative to the frame, which is two inches of relative movement between the ends? Or is it one inch of relative movement between the ends, which is half an inch of movement at each end relative to the frame? Both definitions make semantic sense, and neither is right or wrong. However, the two methods produce rate numbers that differ by a factor of two.

A typical sway bar rate testing fixture measures the force produced at the moving end of one bar arm when that arm is moved an inch. That is the bar's rate in pounds per inch per end pair. The bar's rate in pounds per inch per end is twice that value.

When we multiply the bar rate in lb/in/end pair by the square of the bar-end-to-contact-patch motion ratio, we have the bar's contribution to wheel rate in roll in terms of pounds per inch per wheel pair. An inch per wheel pair is half an inch per wheel.

When we multiply the bar rate in lb/in/end by the square of the bar-end-to-contact-patch motion ratio, we have the bar's contribution to wheel rate in roll in terms of pounds per inch per wheel. An inch per wheel is two inches per wheel pair, and that displacement results in twice as much force change as an inch per wheel pair does.

We evaluate wheel rate in ride, and roll resistance contribution from the ride springs, in lb/in/wheel. Therefore, to have an equivalent value for the bar's contribution we need to also use lb/in/wheel, which is double the number we get if we simply multiply the number from the bar rater by the square of the motion ratio.

So if we take a front bar that contributes 200 lb/in/wheel pair and add it to ride springs that give a wheel rate of 350 lb/in, we don't have a wheel rate in roll of 550 lb/in; we have 750 lb/in in roll. So of course the car has more understeer than it has with 550 lb/in wheel rate in both ride and roll. If we used a bar that contributed 100 lb/in/wheel pair, then we'd have an equivalent setup, and very similar car behavior.

WHAT WHEEL RATE FOR DAMPER CALCULATIONS?

Normally when we calculate the damping rates we calculate the Ccr = 2 Kwmsm and then we apply the damping ratio. By theory the damping ratio for best grip should be about 0,4 at high speeds. In all the books and articles that I read the way of calculating the Kw is $Kw = Ks/MR^2$ where Kw = Wheelrate (N/mm) Ks = Spring rate (N/mm) MR = Motion Ratio (wheel/spring travel).

When we need the best grip is when we are cornering. At this moment we have lots of weight transfer and we have the roll bars acting like springs and at some cases, like at the front of a modern Formula Ford, we have a roll bar that is several times stiffer than the spring itself.

My question is: when we calculate the damping ratios for the best grip should we enter with the total rolling stiffness and with the weight transfer?

Or making the question by another way: when we increase significantly the stiffness of a roll bar should we increase the damping rates?

I would not advise using something other than sprung mass per wheel and wheel rate in ride for calculating damping ratios. However, I do think it's advisable to bear in mind that this is only a crude way of modeling car behavior, and getting the best performance from an actual car may require us to deviate.

In practice, a car with a lot of anti-roll bar relative to its springs generally will work best with somewhat more damping than those springs would require with less bar. However, there is no simple way to calculate an optimal damping coefficient for such cases.

"Weight transfer" in cornering is really wheel load transfer. The mass is what counts when examining oscillatory behavior, and the sprung mass per wheel doesn't change when we corner. However, the car's equivalent mass per wheel in roll is different, and generally smaller, for roll motion than for ride. This means that even with no anti-roll bars, the car has a higher natural frequency in roll than in ride. This is especially true when the sprung structure is vertically and laterally compact, as in a formula car.

So the answer is yes, generally we need more damping to go with more bar, but exactly how much, and in what shaft velocity range, we have to determine by test. In general, the optimum will depend some on the driver as well as the car.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
March 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

FOLLOWUP INFORMATION TO JANUARY ISSUE ON STOCK CAR A/R BAR PRELOAD AND RELATED ASYMMETRIES

In the January newsletter, a questioner mentioned Ron Hornaday's truck being higher on the left after breaking a sway bar arm, and wondered if that would imply that the bar was preloaded. I said it seemed to me that it would.

Actually, that would be so if the truck was higher on the left when sitting still, and if no pit stops including chassis adjustments occurred. If it was only higher on the left on-track, at speed, there is also another possibility.

I owe a tip of the hat here to Mike Keena-Levin of Morse Measurements in Salisbury, NC on this. Morse Measurements does kinematics and compliance (K&C) testing. They do a lot more work with top-division NASCAR teams than I do (mostly Cup and Nationwide; not so many trucks, so they tell me). What they divulged to me about current bar setups is evidently so commonplace nowadays that they consider it no breach of confidentiality to share it with me, and for me to share it with my readers.

It is illegal in the uppermost divisions to preload the sway bar at static, so nobody does that. However, they do use highly unequal sway bar motion ratios on the right and left. This produces some interesting effects, especially when the bar is very stiff and the springs are very soft.

The unequal motion ratios are achieved by using an arm on the left end of the bar that applies its force to the lower control arm at a point further inboard than the right one does. These cars have what is known as a soft link on the left front. Instead of a conventional drop link connecting the bar arm and the control arm, there is a pad that allows the bar arm to lift the control arm, but not draw it down. One reason for this is to keep the bar from de-wedging the car so much if the driver puts a wheel on the apron. Another reason is to facilitate enforcement of the prohibition against bar preloading.

With a bar arm/control arm connection like this, it is easy to have a selection of bar arms with different bends, creating different top-view offsets, that produce a variety of contact points and bar-end-to-tire-contact-patch motion ratios.

If the motion ratio is dramatically less on the left than on the right, some remarkable things happen when the vehicle is on the K&C rig and the front end is displaced in pure ride (equal displacements right and left; car held level in roll but moved up and down). The wheel rate in ride for the left front wheel can be negative: the load at the contact patch can decrease as the front end compresses. In extreme cases, the left front tire will actually lift into the air, even as the fender above it is coming down and the overall front tire pair loading is increasing.

When a setup like this is running on the track, and subjected to aero or banking loading, with only the relatively meager rear elastic roll resistance to work against, the car leans to the left under such loading, and gains wedge, due to the action of the bar. If something in the bar system breaks, and the bar becomes ineffective, the car will run higher on the left as it goes around the track. However, it will not sit higher on the left when stationary, if in fact there was no bar preload in that condition.

Another thing that could make the truck sit higher on the left in such a situation, even when stationary, is the adjustment to the rear jacking screws that the crew would very likely make to compensate for the broken sway bar arm. The crew would need to put more wedge or diagonal percentage into the vehicle, and they would in most cases use the rear jacking screws to do this, since the front ones are less accessible. They would need to raise the left rear corner and perhaps also lower the right rear, and this would tilt the vehicle to the right. This would have to happen on a pit stop. If the arm broke late in the race, and the truck did not pit after that, and it was higher on the left when sitting still, that would suggest a preloaded bar.

FOLLOWUP QUESTION TO FEBRUARY ISSUE ON A/R BARS

I have a followup question to your comments on sway bars in your February newsletter.

It seems to me that at full load, where all the car's weight has transferred to the outside wheel, springs and bars would have the same effect. But while in transition, while the outside is compressing, the way the bar resists that compression is by exerting upward force on the inside wheel. Would that not decrease grip on the inside tire, as long as it would otherwise still have corner weight (downward force) pressing the tire to the surface?

Here's a scenario:
- 1700 lbs of weight on the front end, 100 lbs/corner of unsprung weight
- the springs are 500 lbs/in, and there is 1200 lbs of swaybar resistance (for every inch of chassis roll, or 2" of total swaybar flex)
- the car should roll 1" when all 1700 lbs have transferred

- at 1/2" of roll, approx 425 lbs will still be on the inside corner, and 1275 lbs on the outside
- inside corner should be exerting 425 lbs of force on the inside tire against the ground, giving it grip

- *the spring is being compressed by 325 lbs of corner weight, and so it will be compressed .65" by that. The swaybar will be exerting 600 lbs of force in the other direction, and so will compress the spring another .83"*

  Is the amount of force exerted on the tire reduced by the upward pressure of the swaybar, therefore reducing the traction on that inside tire? Is it correspondingly placing load on the outside tire then?

The questioner is confused on a number of points. I will try to sort things out.

First of all, springs, sway bars, and all other interconnective springing devices are purely displacement-sensitive. None of them have any different effect due to the suspension having a roll velocity. They are only sensitive to roll displacement. Dampers create forces that affect load transfer when the suspension has roll velocity, but sway bars do not, nor do springs.

The scenario posited has a number of problems. Conventionally, the 100 lb/wheel of unsprung weight is not treated as transferring through the suspension, but as discussed in previous newsletters, if there is zero camber recovery in roll, the unsprung masses do create a roll moment that the suspension must resist. So for simplicity, let's suppose that the effective mass acting on the suspension in this half-car model really is 1700 lb.

If that's so, the weight or load transferred due to cornering is not 1700 lb. It is, at most, the load on the inside wheel. If the half-car is assumed to be symmetrical, that's half of 1700 lb, or 850 lb transferred. At that point, the inside tire is at the point of impending lift.

If the wheel rate in roll is 1700 lb/in, 500 from the spring and 1200 from the bar, the half-car has a displacement of only half an inch per wheel at 100% load transfer. Any further roll moment will lift the inside wheel.

At half of that load transfer, 425 lb, roll displacement is ¼" per wheel.

In either case, it would not matter if there were no bar and the spring rate at the wheel was 1700 lb/in instead. The half-car would act exactly the same.

The tires do not know where the roll resistance comes from. They only respond to how much of it there is in total, at a particular instant. The roll resistance may be elastic (from bars and springs), frictional (from intentional and unintentional damping), or geometric (from linkage-induced support forces). But wherever it comes from, it can only hold the car upright by exerting force on the ground, through the tires. There is a simple, inexorable relationship between roll resisting moment, load transfer, and track width:

- Load transfer through the suspension (i.e. less unsprung component) times track width equals roll resisting moment for the wheel pair.
- Roll resisting moment for the wheel pair, divided by track, equals load transfer through the suspension.

This is true regardless of what part or characteristic of the suspension generates what portion of the resistance.

The total roll resisting moment from the front and rear wheel pairs together always equals the roll moment created by sprung mass inertia in response to acceleration (and gravity, which nowadays is sometimes considered an acceleration). The relative roll resistance of the front and rear suspensions controls the front/rear apportionment of the total, but not the overall magnitude of the total. Adding roll resistance only at the front increases front load transfer but not the total load transfer for the vehicle. It follows that rear load transfer must be less.

So yes, the bar does unload the inside front wheel, and it does reduce front grip, and correspondingly increases rear grip, compared to the same setup without the bar. However, it does not do this any more or less than any other method of obtaining the same front roll resistance.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

April 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHY NEGATIVE CAMBER ONLY ON THE FRONT WHEELS, OR MORE ON THE FRONTS THAN ON THE REARS, ON SOME CARS?

I would like to ask you about negative camber. I have an MG Midget that runs on 8in wide slicks. The ones used before (crossplies) are no longer available and some that fit are radials. I have been told that if I run radials I will need more negative camber. I notice that most classes from Australian V8s to F1 have a lot of negative camber on the front. An RX7 owner I know runs about -7deg for road racing. I also note that the negative is only on the front; the backs run flat or minimal negativeve. If this is required to get cornering grip with radials, why is it not also applied to the back. As the backs always have (almost) 0 camber should this not give severe oversteer as the front would have much more grip than the back? How is it that the backs can grip with minimal negative but apparently the fronts can't?

In a fair number of cases, the reason for not having negative camber on the rear wheels, or having very little, is simply that it's not mechanically possible to run much negative on the rears because the rear of the car has beam axle suspension.

If the car does not have a full-floater axle (rear wheels mount to hubs that turn on bearings running on snouts at the ends of the housing tubes; axle shafts have splines on both inboard and outboard end), but rather a conventional style axle (wheel attaches to a flange on the end of the axle shaft; axle shaft has splines only at the inboard end), we can only get a small amount of negative camber. We have to bend the middle of the axle housing down to do this, and we are limited by friction and wear at the splines, and in some cases by inability to insert the axle into the diff at all.

With a full-floater axle, we can do this and can also typically get half to three quarters of a degree at the outboard ends, with ordinary straight splines. Additionally, where legal, we can often obtain special axles with barrel-shaped splines, and special cambered snouts. But even with this hardware, it is uncommon to see more than about two degrees.

With independent rear suspension, it is, at least in theory, easy to get all the negative camber we want, although in production cars we may be limited by the adjustment range afforded by the stock hardware.

Negative camber is a mixed blessing, with any tire. For cornering, what we really want is some amount of inclination into the turn, meaning negative camber on the outside wheel and positive camber on the inside one. Exactly how much we want depends on the tire, the rim width, the inflation pressure, the grippiness of the road surface, and the normal force on the tire. The reason negative camber on both sides of the car is helpful is that, up to a point, it's worth giving up some cornering power on the inside tire to get an increase on the outside tire, since the outside tire is more heavily loaded and consequently more important: the gain on the outside tire is greater than the loss on the inside one. Ordinarily, with passive suspension we can't get greater than 100% camber recovery in roll, or even anything approaching 100%, so we sacrifice inside tire inclination to improve outside tire inclination.

The more load transfer we have, the greater the gain on the outside tire becomes, relative to the loss on the inside tire. Taking an extreme case, if there is no load at all on the inside tire, its camber doesn't matter at all.

But of course we don't just want lateral force from our tires – we need them to make longitudinal force as well. And for that, we want them straight up: zero camber. So camber settings are always a compromise between lateral grip and longitudinal grip. At the front, running a lot of negative camber will help front cornering power up to a point, but this will come at the expense of braking ability. At the rear, both braking and propulsion will be adversely affected (assuming rear wheel drive).

When we see really extreme static camber settings, that usually means the suspension's camber control properties are less than optimal. All Mazda RX7's have strut front ends, which have very poor camber recovery in roll when lowered for racing. The first version of the RX7 has a beam axle rear end, with flanged axles, so it isn't mechanically possible to get a lot of negative camber at the rear, but the camber recovery in roll is much better than at the front. Later ones have a form of semi-trailing arm independent that can be set with quite a lot of static negative camber, but has better camber recovery in roll than the front suspension has.

When the front end has much poorer camber control in roll than the rear has, we can't necessarily assume that front wheel inclination is more favorable than rear when cornering, even on the outside front wheel, just from the static camber. It may be, or it may not be. And if the front end does have more favorable cornering camber than the rear, it does not necessarily follow that we will have oversteer. The camber does create a tendency toward oversteer in such a case, compared to a different camber picture, but other factors also enter in. Those factors include front/rear roll resistance distribution, front/rear weight distribution, and the front/rear tire size relationship.

If we do manage to achieve better cornering camber at the front than at the rear, that means we can use more spring and/or anti-roll bar at the front, and get balanced cornering with slightly greater limit adhesion, plus better ability to put power down while cornering, due to greater inside rear wheel loading. So we usually want to set the front camber for best front cornering power, provided we don't hurt braking too much. If the car is too loose, we just add front spring and/or bar. We could balance the cornering by taking front camber out instead, but that results in lower limit lateral acceleration, and poorer ability to put power down on exit.

With a spridget, we at least have a front end that benefits from being lowered, in terms of camber recovery in roll, rather than deteriorating with lowering as a strut suspension does. With radials or bias-plies, give the front end as much negative camber as will improve front cornering grip, tempered by attention to braking, and then add front roll resistance as needed to eliminate any oversteer.

ROLL CENTER/FRONT VIEW GEOMETRY WITH DUAL BALL JOINTS

On cars with double ball joint strut suspension (Hyundai Genesis Coupe, Pontiac G8) I understand that the steering axis operates through the 'virtual' joint at the intersection of the two control arms when viewed from above. How do you calculate the roll centre position for a suspension like this? Do we use the same 'virtual' joint in the same way as regular strut suspension? Or is there some other point?

With any strut suspension, there is a virtual upper control arm plane that is perpendicular to the strut axis. The strut axis is often not the same as the steering axis, even with single ball joints. Often, the ball joint lies somewhat outboard of the strut axis. The steering axis is then a line through the ball joint center of rotation and the upper strut mount center of rotation, and the steering axis inclination is greater than the strut inclination.

With dual ball joints, the virtual ball joint can be even further outboard, but that only affects the steering axis, not the strut axis, and not the virtual upper control arm plane.

If the two lower links are in a common plane, or very nearly so, that common plane can be taken as the lower control arm plane. The front-view projected control arms are then the lines where the upper and lower control arm planes intersect the front axle plane (vertical, transverse plane containing the front wheel centers). The front-view instant centers are the points in that axle plane where the upper and lower front-view projected control arms meet.

The front-view force lines are then definable as lines from the contact patch center to the instant center. The slopes of these lines, in combination with the ground plane forces at the contact patches, determine the linkage-induced support forces in the system, and in turn the anti-roll or pro-roll moment produced by these, and accordingly the roll center height.

The front-view force lines are always instantaneous perpendiculars to the front-view projected motion paths of the contact patch centers. The motion path's instantaneous inclination from vertical, or the dy/dz of the contact patch as the suspension moves, is the same as the instantaneous slope from horizontal, or dz/dy, of the force line.

It is possible to create multi-link suspensions that do not have pairs of links in common planes, or even close to common planes. In such cases, it becomes problematic to try to define projected control arms. It may only be possible to find the motion path of the contact patch with a 3-D computer program. But even then, the front-view and side-view projected force lines are instantaneous perpendiculars to the contact patch motion paths.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

May 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

WHY USE SWAY BARS AT ALL?

In the last few issues, you have been talking about anti-roll bars and how they affect mechanical traction.

My question is why do we use a sway bar at all? What effect will removing front and rear sway bars have on the behavior of the car? Sure the car will roll on the corners, but what about traction on the bends and overall performance?

I would definitely not advise taking the sway bars off an existing car, and then taking it for a lap at speed. On most cars today, the front/rear roll stiffness distribution is tuned with the bars, and the car will not have the same oversteer/understeer balance if the bars are removed.

It is quite true that for vehicles with the sprung mass c.g. heights commonly encountered in cars, roll has only a small effect on overall load transfer. It is also true that we can achieve any roll gradient (how much the car rolls per unit of lateral acceleration) using only the ride springs. And finally, it is true that we can get any front/rear distribution of load transfer using only the ride springs. So why not just do that?

In fact, many cars have been built with no sway bars (or anti-roll bars – I use both terms interchangeably). I'm not sure when the first anti-roll bar appeared, but before World War II they were practically unknown, both on passenger cars and on race cars. But cars without any anti-roll bars basically fall into three categories:

- ones that roll a lot;
- ones that ride really hard and don't absorb bumps well;
- ones with beam axles at both ends.

Some cars are in more than one of these categories at once.

Anti-roll bars and other interconnective springing devices offer the following advantages:

- they let us control wheel rates in the four modes of suspension motion – roll, pitch, heave, and warp – independently, or at least independently within certain constraints;
- they let us achieve better control of roll for a given ride quality;
- they afford us a way to readily adjust front/rear load transfer distribution, even from the driver's seat, with a minimum of effect on other things.

What's wrong with just letting the car roll?

With independent suspension, the wheels inescapably lean with the sprung mass to some degree. We can compensate for this by using geometry that makes camber go toward negative in compression in ride, but we can't get 100% camber recovery in roll without excessive camber change in ride. Generally, we have to accept 50% camber recovery or less: the wheels lean at least half as much as the body. So roll hurts camber, and consequently reduces grip, with any independent suspension. This is not true with beam axles, but there are other disadvantages to roll even with beam axles.

Roll uses up suspension travel and, on the outside, uses up ground clearance. Depending on the amount of travel available, roll can bottom out or top out the suspension, or lead to the bodywork or skirting dragging, or create a situation where undetected bottoming out or topping out occurs on bumps. In cars with ground effects or very low front wings, roll can disrupt the under-car aerodynamics enough to cause problems. These disadvantages occur even with beam axles.

Roll does increase lateral load transfer, directly and indirectly, although the amount of this is quite variable. There is a small lateral migration of the car's c.g. with roll, relative to the tire contact patches. In tall vehicles, this can become significant. In cars, it is relatively small, but still present, and undesirable. In addition to this direct effect, there can be an indirect increase in load transfer if roll forces us to run the car at greater ride height to avoid bottoming.

The whole question of whether using anti-roll bars (or stiffer ones, or various blends of spring versus bar) causes the car to run higher or lower is fairly complex, and will be addressed further below.

So far, we have been discussing effects of roll in steady-state cornering. In abrupt transient maneuvers, greater roll displacements imply greater roll velocities and accelerations. Roll acceleration will add to lateral load transfer when the direction of roll acceleration is into the turn. This occurs as the vehicle approaches maximum roll displacement, and roll is outward, and still increasing, but slowing. In taller vehicles, this added component of load transfer can sometimes be enough to make the difference between the vehicle staying upright and having all wheels on the ground, or overturning or lifting wheels. This phenomenon explains why some SUV's will lift wheels or roll over in a lane change test, even though they will stay right side up and slide controllably in a skid pad test. Reducing roll displacements in such a vehicle will improve overturning resistance in abrupt maneuvers.

We generally will not encounter this in racing, but the same effects are present, in lesser magnitude. We can often observe that softer setups favor smoother drivers, and vice versa.

For road use, and generally for real-road competition, ground clearance requirements are mainly established by the need to clear driveways, snow, and obstacles at low speed. Bottoming of the suspension or the underside of the car on bumps and dips is usually not the main constraint. We will also want to choose ride stiffnesses and front/rear natural frequency relationships that will make the car ride well and take bumps well. We will need to balance these objectives against a need to keep roll within reasonable limits. For this, we will do well to get between 30% and 60% of the overall elastic roll resistance from the bars, and the rest from the springs.

For a pure racing car, the constraints are different. We will generally be prepared to put up with a fair amount of hassle loading and unloading the car, and having it unable to negotiate driveways, if that's what it takes to win races. We will then be concerned about how low we can run the car without having the underside hit the track (much). We may also need to limit heave and pitch displacements very tightly to control under-car aerodynamics.

Those constraints will dictate that we have the car very stiff in the heave and pitch modes. When we want the car that stiff in heave and pitch, adequate roll stiffness almost comes with the package automatically. Even then, however, we will generally run some anti-roll bar, for the following reasons:

- to provide fine adjustment of roll resistance;
- to provide quick adjustment of roll resistance;
- to provide driver adjustment of roll resistance;
- to allow steeply rising ride rate, or pitch and heave rate, via a third spring, while maintaining linearity or whatever rate curve is desired, for the roll and warp rate.

With a bar and a third spring, we can actually have a setup where the wheel rate in roll exceeds the wheel rate in ride at or near static ride height, while the ride rate is the higher of the two once the suspension is compressed a bit.

There are peculiar circumstances, generally imposed by peculiar racing rules, that may favor setups with really soft springs and stiff bars, or with a stiff wheel rate in ride and little or no wheel rate in roll.

Big bar/soft spring front end setups have been popular for some time in American stock car racing, although this trend is fading somewhat. This approach makes sense, at least up to a point, for stock car front ends, because the rules impose a minimum ground clearance requirement that is considerably higher than where we'd like the front of the car to be on-track, and the turns are banked. Soft springs let the front end compress due to banking and aero forces, while the bar maintains enough roll resistance to keep the car from being loose.

The opposite extreme exists at the rear of a Formula Vee car, where the rules require swing axle suspension, which has an extremely high roll center and a tendency to jack up in cornering. Here, it works best to use springing that only works in ride, and have a wheel rate of zero in roll. The last thing we'd want for this application is an anti-roll bar.

Cars with beam axles at both ends can have very ample geometric anti-roll (high roll centers) at both ends, and can consequently have moderate roll gradients without anti-roll bars. Any beam axle suspension without an anti-roll bar has a smaller wheel rate in roll than in ride, because of the springs inevitably being inboard of the wheels. However, high roll centers carry some penalties of their own, even with beam axles, in the form of large lateral movement of the contact patches with respect to the sprung mass on one-wheel bumps. Beam axle suspensions generally perform much better without anti-roll bars than independent suspensions do, but there is nothing wrong with using an anti-roll bar with a beam axle either, and indeed this is very common on roadgoing beam axles today, and on racing ones as well, where the rules don't prohibit anti-roll bars.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

June 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

EFFECTS OF CASTER AND SCRUB RADIUS

I continue to enjoy your newsletter and would really like to see some of your insights into the effects and limitations of high caster angle, preferably in conjunction with both large and small scrub radii. It seems to me double-digit caster angles would provide favourable camber-into-the-corner of both front tires during steer and that significant scrub radius would add a jacking effect that would usefully load the inside front tire, but perhaps it's not quite that simple...

The effects described do indeed happen, along with a few other things. More caster does make the front tires lean in the direction of steer. And this combined with more scrub radius (or front-view steering offset) does add load to the inside front and outside rear tires, while correspondingly unloading the outside front and inside rear. This is sometimes called de-wedging, and as a rule it adds oversteer, or reduces understeer. But it’s not simply a case of more tire inclination and de-wedging with steer being better without limit.

First of all, it is possible to have too much tire inclination. Tire inclination improves lateral grip up to a point, but beyond that optimum, the tire runs too much on its inside edge, and grip starts to be lost.

Lots of caster and a big scrub radius will result in the car de-wedging a lot with steer, and also rolling outward as the inside front is lifted and the outside front is let down. The roll somewhat reduces the tire inclination gain, and will also tilt the rear tires out of the turn some, with independent rear suspension.

The amount of de-wedging also depends on the wheel rates in warp throughout the suspension system, and the torsional stiffness of the sprung structure.

We will not in all cases want the wheel loads to change a lot with steer. On ovals, it will generally be unnecessary. On really high speed ovals, it will tend to make the car twitchy and make it harder for the driver to be smooth. More moderate caster and scrub radius are preferable for that application. Applications where we want to go the other way include autocross, street circuits, tight road courses, chicanes, twisty hill climbs, and so on – situations where we need the car to respond fast to abrupt maneuvers, and stability and smoothness are not paramount.

Where the vehicle has a spool or locked axle, and runs on a road course, it is particularly common to employ lots of caster jacking. When the turns go both ways, we can't use tire stagger to prevent locked-axle push, and unloading the inside rear becomes crucial to making the car turn, especially on the tighter radii.

Examples of racing vehicles using unusual amounts of caster jacking to overcome locked axle push would include go-karts and Australian Super Cars, and perhaps even Legends cars to a smaller degree. We don't have much freedom on spindle or upright design in a Legends car, but we do have some caster adjustment. On karts, it is common to vary the scrub radius using spindle shims as a way of tuning the vehicle. Sometimes, parts of the frame are made removable as a way of adjusting torsional stiffness as well.

Scrub radius and caster definitely affect steering feel. It is important to remember that a race car is a tool for a human being, and drivers' preferences in steering feel differ. It is also important to note that if we are designing our own spindles or uprights, there are additional parameters to consider, which interrelate with scrub radius and caster. This means that there is more than one way to vary caster, and more than one way to vary scrub radius.

If we are adjusting an existing car, or if we are committed to a single upright design, adding caster adds trail. This increases steering handwheel torque per unit of lateral force at the front tires. This translates to higher control effort, especially in hard cornering, and also increases the car's tendency to follow lateral road slope. When we have more freedom on our upright design, we can employ some pin lead to decrease trail, and have any combination of caster and trail that we want.

Even when we have less trail due to pin lead, increasing caster still increases caster jacking.

If we have design freedom, we can vary scrub radius by varying front view steering axis inclination (SAI), or by placing the ball joints closer to the wheel centerplane or further inboard. These two approaches have different effects on steering feel.

If we had zero SAI and positive caster, and some scrub radius, the front end of the car would not only roll oppositely to steer (with corresponding de-wedging of the car), it would also drop with steer. This would cause the steering to try to go away from center at low speeds, and at a standstill. If, conversely, we had no caster and some SAI, the front of the car would be lifted with steer, and the steering would seek center with respect to vehicle centerline.

When we have some SAI and some caster, what happens at a standstill or at parking speed is that near center, the steering seeks center, but at some amount of steer, it tries to go away from center.

SAI also affects camber change with steer. It causes both front wheels to go toward positive camber with steer. This means that when we have a combination of SAI and caster, the outside wheel will gain inclination with steer, but at a decreasing rate, and at some point will start to lose inclination, while the inside wheel will gain inclination at an increasing rate. This may not be too bad, particularly if the front wheels have some static negative camber.

Finally, large amounts of caster and scrub radius can in some cases produce really disagreeable behavior in the steering, in the form of various types of oscillation. Any runout in the tires or pulsation in the brakes will be felt more in the steering. I have had clients running large caster settings on stock cars report steering shimmy at lower speeds. People's willingness to live with such effects in the pursuit of more speed will vary.

CAMBER GAIN RECOMMENDATION

I have built 2 7s type cars and am busy with an exoskeletal car. This car is out and out for track racing.

Do I go for the unequal double wishbone design (top wishbone 2/3 of the bottom)? What should the camber gain be that I design into my uprights?

I do not know how much the car is going to roll.

Does one strive to get the camber gain front and rear the same? (independent rear suspension)

Unequal double wishbone is good. Top wishbone 2/3 of bottom is a reasonable rule-of-thumb recommendation for uprights of typical dimensions. You don't need to worry too much about straying from the 2:3 ratio if packaging or structural requirements dictate.

"Camber gain" typically refers to the rate of camber change per unit of suspension displacement, as measured statically in the shop, moving the wheel or upright up and down with a jack and measuring camber change as the displacement changes.

My standard recommendation for camber gain is 0.6 to 0.9 degrees per inch, which would be roughly 0.02 to 0.03 degrees per millimeter. With typical track widths, this gives camber recovery rates between a bit less than 50% and a bit better than 25%. That is, ignoring roll due to tire deflection, and ignoring other compliances, the wheels lean a bit more than half as much as the body in cornering, but less than ¾ as much, and they don't experience any really huge camber changes due to acceleration, braking, and bumps.

We also need to avoid excessive jacking. That is, we want the roll centers fairly low: modest geometric anti-roll. There is no hard and fast relationship between camber gain and geometric anti-roll when we have total freedom to move all points, but in the majority of cases, where we face packaging constraints and have committed to some point locations, more camber gain will produce more lateral anti, or higher roll centers.

My general recommendation is that at static ride height, the force line – the line from the contact patch center to the front view instant center – should slope upward toward the center of the car, at around four degrees. In no case should the slope be more than eight degrees (roll center around 4" above ground). This means that the coordinates of the front view instant center, using the contact patch center as a local origin, as in Mitchell software, should have z (vertical) and y (transverse) coordinates with a ratio around .07 z/y, or in no case more than .14.

You will want to make sure you stay within these limits at the highest ride height you design for. Generally, race cars have their ride height tuned to the circuit. With a smoother surface, we can run the car lower, and we should. Generally, that lower ride height brings lower roll centers. That's okay.

It's not bad if the geometric anti-roll is very small, or even negative (roll center low, or below ground), but don't let it get too high. When there is little geometric anti-roll (roll center near the ground), the lateral location of the force line intersection will move all over the place with suspension movement. Some people will tell you that the force line intersection is the roll center, and/or that you can model car behavior by taking moments about that point, and that minimizing lateral (or vertical) migration of that point is vital to good car behavior. It isn't true. Don't worry about it. Cars work fine with little geometric anti-roll. Just don't have a lot, with independent suspension. That does create problems.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

July 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MUMFORD LINKAGE

I would also like to find more information on the Mumford link as it seems to be a potentially useful modification to my Bugeye Sprite. Can you point me to any sources of information on how it works? There doesn't seem to be a great deal of info readily available, perhaps it's not as popular as a Watts linkage of Panhard rod because of the complexity and difficulty in determining link sizes, pivot locations, etc. Any leads you can give me will be greatly appreciated.

The Mumford link, or linkage (since it is a linkage: a system for motion control containing more than one link, together with other elements), is actually a family of possible layouts, commonly used for lateral location of beam axles. This family has at least four possible members.

There are two other families. The better known Watt linkage consists of a central rocker, and two links attaching to the ends of a centrally pivoted rocker, extending in opposite directions from the rocker. The WOB linkage has a rocker with its pivot at one end, and two links different distances from the pivot, extending in the same direction from the rocker.

The Mumford linkage uses two rockers and three links. All three systems can give very nearly straight-line axle location in ride, while allowing roll about a reasonably well defined point, which serves as the roll center, or as one end of an axle axis of rotation whose axle plane intercept defines the roll center.

The Mumford linkage is thus a more complex means of accomplishing a simple task than its rivals are, so the only way it can be justified is if it either provides better packaging or load paths in a particular situation, or if it permits geometrical properties unobtainable otherwise.

The usual appeal of the Mumford linkage is that it is able to provide a roll center below any part of the hardware. The roll center can be lower than the ground clearance of the car. It can be lower than a smooth floor pan under the rear axle, transitioning smoothly into a diffuser.

Now, why do we want an unusually low roll center with a rear beam axle? Not to prevent jacking. Maybe to minimize lateral scrub on one-wheel bumps, particularly with stiff tire sidewalls. But more commonly, we want a low rear roll center with a live rear axle so we can have a lot of elastic rear roll stiffness, via springs and/or antiroll bars, relative to the front elastic roll resistance. The advantage of this, for a road racing car, is that it minimizes torque wedge – the change in diagonal percentage due to driveshaft torque.

With appropriate design, there are better ways to minimize torque wedge. But these may or may not be understood by the designer, and they may or may not be allowed by the rules.

It should be noted that there is also a way to get a roll center below any portion of the linkage using a Watt linkage. More on this later.

There isn't any standard nomenclature for the parts of a Mumford linkage, so I am inventing some. The mechanism has two rockers, each with two link attachment points and one pivot. One link from each rocker connects to the axle (if the rockers are attached to the frame or sprung structure) or the frame (if the rockers are on the axle). I will call these the side links. The third link connects the two rockers to each other, and constrains their movement so they can only pivot equal amounts, in opposite directions. I will call this the center link.

There are four basic configurations of the Mumford linkage. The rockers can be near the center of the vehicle, with the side links extending outward from them, and a short center link, or they can be out toward the sides of the vehicle, with the side links extending inboard from them, and a long center link. Additionally, the rockers can be on the frame or on the axle. Thus, there are four permutations:

- Rockers unsprung (on axle), inboard (RUI)
- Rockers unsprung, outboard (RUO)
- Rockers sprung (on frame or sprung structure), inboard (RSI)
- Rockers sprung, outboard (RSO)

All of these options can give a roll center below any point on the linkage. However, they differ with respect to what I call the Mitchell index: how the roll center height varies as the suspension moves in ride.

When the rockers are unsprung, the roll center only moves a little with ride, relative to ground. It moves oppositely to the sprung mass: a Mitchell index that is negative, with a small absolute value.

When the rockers are sprung, the roll center moves with the sprung mass in ride, a bit more than the sprung mass: a Mitchell index a bit greater than one.

If the system is paired with an independent suspension at the front, the front suspension will in most cases have a Mitchell index of one or greater. It is difficult to get a Mitchell index near or less than zero with independent suspension.

It is desirable to have similar Mitchell indices at the front and rear of the car. This way, the roll axis inclination will not change drastically as the car negotiates humps and dips while cornering: the relationship of front geometric anti-roll to rear geometric anti-roll will not vary a great deal with heave displacement. The car’s oversteer/understeer balance should therefore be more constant, particularly in undulating sweepers. That lets the driver press closer to the limit of adhesion, with less worry that the car will get away from him.

However, many cars race fairly effectively in total disregard of this. It is not the end of the world if the front and rear Mitchell indices are dissimilar.

I mentioned that it is possible to get a very low roll center with a Watt linkage. This involves laying the rocker flat, under the axle, and having it angled in top view, rather than straight front to back. The two links are then made higher at their outboard ends than at their inboard ends, where they connect to the rocker. Care has to be taken to make sure the joints don’t run out of angular travel as the system moves. Packaging permitting, the rocker can be mounted to the sprung mass, and the outer ends of the links can be attached to the axle tubes. In that case, the Mitchell index is a bit greater than one: better for compatibility with independent front suspension, but not essential. Alternatively, the rocker can be attached to the axle, and the outer ends of the links can attach to the sprung mass. In that case, the Mitchell index is negative, and small in absolute value.

The questioner is interested in applying the system to a Mk. 1 or Bugeye Sprite. The stock suspension on that car consists of trailing quarter-elliptical leaf springs a bit below axle center, and rubber-bushed trailing links above those. The springs provide lateral location. There is no additional mechanism for lateral location. The roll center is then at spring height, and the system appears to be fairly compliant laterally.

American production sports car rules (e.g. SCCA) would require that the leaf springs be retained, but would allow addition of lateral locating devices, and also allow changing bushing design.

If the objective is merely to take lateral compliance out of the system, probably the best approach would be a diagonal Panhard bar running approximately from where one of the leaf springs connects to the unibody to where the opposite spring connects to the axle. This would work with stock bushings, or with upper links having rod ends and adjustable length. It would not lower the roll center. The bar might need to have a bend, and pass over or under the driveshaft.

If a lower roll center is desired, the attachment of the springs to the axle has to be modified so that the springs no longer locate the axle laterally. Or, rules permitting, the springs might be eliminated entirely, and replaced with Heim-jointed links. In any case, the axle must be free to move laterally with respect to the springs, or the system will bind in roll.

If the roll center is lowered, the car will need more rear spring or anti-roll bar for similar steady-state cornering behavior. The main advantage to going this way will be that the car will put power down a bit better exiting right turns.

TOE OUT FOR TURN IN

Popular wisdom always suggests that for a faster turn-in, toe-out is the way to go. I've never read a satisfactory explanation into the possible reasons and it seems counter-intuitive to me. Thinking about what happens at the tires, if you have toe-out on turn-in the outside tire has to pass "over center" before it begins building grip in the direction of the turn. Thinking that the outside tire is carrying most of the load and therefore creating the majority of the grip I would think that toe-in would provide a faster turn-in as you would be building grip faster as the more loaded tire would already have a slight slip angle before you even turn the wheel.

I've never done much on-track experimentation to see if I personally believe toe-out improves turn-in, but it's hard to believe such a prevalent opinion came out of nowhere. I have a theory that maybe the cause could be the change in relative wheel heights as the scrub radius/caster kingpin/trail cause the inside front wheel to move down in relation to the chassis and the outside to move up.

It would seem that with a fast enough turn in that the inside was temporarily the heavier loaded tire and its steered angle would create more grip than the outside until the point that the load transfers to the outside tire.

If this is the case, it would seem that in softly sprung vehicles with lots of steering inclination, toe-out would be the way to go, at least on tight courses requiring fast turn-in.

What are your thoughts on this age old myth/truth?

Cars do generally exhibit quicker initial turn-in with static toe-out. My analysis is that this does not have to do with the lateral (y axis, per SAE convention) forces from the front tires, but rather the longitudinal (x axis) forces, which can also produce yaw moments.

When the car is running straight, and the front tires have either toe-in or toe-out, the tires are both running at a slight slip angle, and accordingly generating both some lateral forces and some drag forces. The drag forces are roughly equal, and additive. The lateral forces are roughly equal, and opposite in direction, so they approximately cancel.

When the handwheel is turned just a tiny bit, one front wheel will be running straight, and the other will be turned into the corner, generating a bit of drag, and some lateral force into the corner.

If the car has toe-in, it will be the inside front wheel that's running straight, and the outside one that has some slip angle. In this condition, the lateral force creates a yaw moment into the turn, but the drag force creates a yaw moment out of the turn: the two yaw moment components are subtractive.

If the car has toe-out, it will be the outside front wheel that's running straight, and the inside one that's making lateral force and drag force. Now the lateral force and the drag force both create yaw moments into the corner, and are additive. Consequently, the net yaw moment is greater, and the car experiences a greater yaw acceleration: it turns in quicker.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

August 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TOE, ACKERMANN, CORNERING FORCE, TURN-IN AND RELATED PUZZLEMENTS

Following on from your latest chassis newsletter regarding the advantages of static toe out to aid turn in, what are your views on using anti-Ackermann geometry to put more steering angle onto the tyre with the greater slip angle?

Last month I mentioned that a car gets quicker initial turn-in when it has some static toe-out at the front wheels. I'll stand by that, but I think I need to clarify a bit what quicker initial turn-in actually means in this context. It doesn't necessarily mean that the car negotiates that segment at a higher speed, or that the lap time will be lower. It means that it takes a smaller handwheel movement to produce the desired initial yaw acceleration to make the car follow a desired line, and the car feels more responsive.

This may improve lap time, when a very rapid increase in yaw velocity is desired, as in a chicane or a sharp turn in autocross. But in such cases it will matter at least as much what the toe condition is at large handwheel displacements.

In most oval track and road course turns, it is easy for the driver to make too abrupt a yaw input, with any front end settings. Getting the car to respond quickly isn't really the key to good segment time, as much as being smooth is. The beginning driver usually needs to work on having slower hands.

Let's review a bit of terminology. Anti-Ackermann, or negative Ackermann, is steering geometry that makes the front wheels toe in with respect to each other when steered away from center. Zero Ackermann is parallel steer: front wheels maintain static wheel-to-wheel toe setting as they steer. Positive Ackermann is steering geometry that creates wheel-to-wheel toe-out with steer. 100% Ackermann is steering geometry that gives an amount of toe-out that produces scuff-free operation in low-speed turning.

The turn center is the point about which the car's c.g. or origin is instantaneously moving when the car is traveling in a curved path – in other words, the instantaneous center of curvature of the car's path.

The car's origin is a point chosen as its center for modeling purposes. It may be at the car's c.g., or it may be at the midpoint of its wheelbase and tracks. I tend to prefer the latter, at least for the purposes at hand.

Slip angle is the angular difference between the direction a wheel or other object is aimed (its heading) and the direction it's actually traveling (its bearing).

A tire that's cornering hard runs at a slip angle. On a given surface, at any given load (F_z), pressure, temperature, camber, and so on, the tire will have a characteristic relationship between slip angle and lateral or y-axis force or F_y. When graphed with slip angle on the horizontal axis and F_y on the vertical axis, the function displays as a concave-down curve that rises, peaks, and then descends. The peak occurs at the slip angle where the tire generates the greatest F_y under those conditions.

That peak occurs at a greater slip angle when the load on the tire is greater. Consequently, a pair of unequally loaded tires will generate the greatest total F_y when the more heavily loaded one is running at a somewhat greater slip angle than the less heavily loaded one. This leads some to conclude that the wheels need to be toed in a bit when cornering, to achieve that slip angle relationship, and therefore a race car should have negative Ackermann.

Unfortunately, the relationships involved are far from simple. In particular, there is not a simple relationship between toe and slip angle, contrary to what one might suppose.

At low speed, with all steering done by the front wheels, and little or no slip angle on any tire, as seen from above the car, the turn center lies on the rear axle line. For the front tires to be at zero slip, as seen from above their axes need to pass through the turn center. This means that the outside front wheel must steer less than the inside front wheel, and the axis of the outer front wheel must lie rearward of the inner front wheel's axis. We may then say that the inside front wheel is leading the outside front wheel through the turn (or the outer is trailing the inner), and neither rear wheel is leading the other.

In this case, the front wheels must have toe-out to have equal slip angles. If they are parallel, the outside one will have a positive slip angle and the inside one will have a negative slip angle. Both will be scuffing, in opposite directions. Interestingly, the car as a whole, at its origin, has a negative slip angle: the body points out of the turn, and the nose tracks outside of the tail. The rear wheels track inside the fronts.

If we start adding speed, the car as a whole will assume a drift or attitude angle, and all the tires will start to assume slip angles, if they had none at very low speed. If the car has neutral handling, or an understeer gradient of zero, the driver will not need to move the handwheel to keep the car on course, but the car will change attitude relative to its direction of travel. The tail will run further out, and the nose will run further in. The turn center will no longer be on the rear axle line, but will be forward of it.

As we keep increasing speed, and the turn center keeps moving forward with respect to the car, we will reach a point where the turn center lies on a line perpendicular to the car's centerline, and passing through the origin. At this point, the body has a slip angle of zero. The outside rear wheel leads the inside rear wheel. For the rear wheels to have equal slip angles, they need to have toe-in. The inside front wheel still leads the outside front, but by less than at very low speed. For the front wheels to have equal slip angles, they need to have toe-out, but less than at very low speed.

The rear wheels will now be making tracks that lie approximately on top of the front tire tracks, assuming front and rear track widths are close to equal.

Since the car is now making some lateral acceleration, it will now have some lateral load transfer. This will make the outside tires more heavily loaded than the inside ones, and there will be a case for having less toe-out at the front, and more toe-in at the rear, because of that. However, the exact amount of this that we want will depend on the exact nature of the tire properties we're dealing with. It won't be the same for all cars, or all tires.

As we add more speed, the slip angles will increase further, and the turn center will continue to move forward with respect to the car. The car as a whole will now have a positive slip angle, and the rear tires will be tracking outside the fronts. At some point, we will reach a condition where the turn center lies on the front axle line. In this condition, we finally have the relationship we might have imagined between front toe and front slip angle: toe-in definitely implies greater slip angle on the outside tire; toe-out definitely implies greater slip angle on the inside one; parallel front wheels actually do have equal slip angles.

The rear tires are now tracking well outside the fronts. The outside rear is leading the inside rear by more than before. The rear wheels now need more toe-in than before, to have equal slip angles.

Finally, as we add more speed, the turn center can move forward of the front axle line. In this situation, the outside front wheel leads the inside front, and the inside front has the greater slip angle if the front wheels are parallel. The front wheels now need toe-in to have equal slip angles.

We can thus define three turn center zones:

- Zone 1: Turn center between rear axle and origin. Rear wheels track inside fronts. Fronts need toe-out to have equal slip angles.
- Zone 2: Turn center between origin and front axle. Rear wheels track outside fronts. Fronts need toe-out to have equal slip angles, but not as much as for Zone 1.
- Zone 3: Turn center ahead of front axle. Rear wheels track outside fronts, more than in Zone 2. Front wheels need toe-in to have equal slip angles.

Note that because we have assumed a zero understeer gradient, the handwheel and the front wheels (with respect to the car) are at the same angle through all of this. In the real world, the understeer gradient is usually positive, or in some cases may be negative – it is seldom truly zero. If the car has the front wheels steered, say, an average of four degrees to the left, it could be understeering through a left-hand sweeper with the rear wheels at a very small slip angle (turn center in Zone 1), or it could be on a smaller radius left turn with the rears sliding considerably (turn center in Zone 2), or it could be powersliding around a left hairpin (turn center in Zone 3). The car could even be turning right, in a state of pronounced oversteer, with the driver countersteering (turn center in Zone 3, and possibly to the right of the car's centerline).

The car could be cornering, and have the front wheels pointed straight (zero handwheel displacement). This implies some amount of oversteer, and implies that the rear tires are tracking outside the fronts and the turn center is in Zone 3. The turn could be a sweeper, with the rear wheels sliding just a little, or it could be a tighter turn, with the rear wheels sliding a lot. In any such case, if the handwheel displacement is zero, Ackermann will have no effect at all, and the toe setting will be about the same as static, altered only slightly by compliances and roll steer.

Each of these situations requires a different front toe condition for equal slip angles, or for a given amount of slip angle inequality. Yet any conventional steering system can only provide one fixed relationship between toe and average steer. Therefore, there is no way to create steering geometry that is optimal for all conditions.

However, we can make some useful generalizations, at least for cars running on pavement and having small understeer gradients:

- Very tight turns will be taken with the front wheels steered hard into the turn, and the rear wheels will track inside the fronts, unless the driver is severely horsing the car with the throttle. The inside front will lead, and the turn center will be in Zone 1. Even if we want the outside front tire to operate at somewhat greater slip angle than the inside front, we will need some toe-out. Since the front wheels are steered a lot, Ackermann will have a lot of influence on toe.
- Sweepers will be taken with the front wheels steered very little, and the rear wheels will track well outside the fronts. The outside front will lead, and the turn center will be in Zone 3. We will need slight toe-in, even to have equal slip angles on the front wheels. To get greater slip angle on the outside front, we will need some added amount of toe-in. Unfortunately, it will be hard to get this using Ackermann, because we have very little handwheel displacement. It will have to come mainly from our static setting.
- There will be some intermediate range of turn radii where the handwheel displacement is moderate, but enough so that Ackermann has noticeable influence, and the turn center is in Zone 2 – meaning we need a little toe-out for equal front slip angles, or roughly zero toe or a little toe-in if we want more slip angle on the outside front.

How do we get a compromise that at least roughly matches these requirements? By using the combination that is standard practice for road cars: a bit of static toe-in, and some positive

Ackermann. To tailor the same car for autocross or hillclimb use, we might use some static toe-out instead.

But it isn't really crucial to get this perfect. Lots of races have been won with negative Ackermann, and static front toe-out. Negative Ackermann is advantageous when catching oversteer slides. In many front-steer cars, packaging constraints confine us to negative Ackermann, and we have no choice but to make the best of it.

Finally, there is some reason not to want the inside and outside tires to reach peak cornering force exactly together. That gives us the highest peak cornering force, but also may make breakaway a bit more sudden. If the inside and outside tires optimize at slightly different lateral accelerations, that may spread the cornering power curve and make the car more forgiving. This is somewhat analogous to tuning an engine's intake and exhaust systems for somewhat different rpm, to spread the engine's power curve.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

September 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANTI-ROLL BARS ON FORMULA VEES REVISITED

It has been brought to my attention that Formula Vees as raced in Europe, including the UK and Ireland, do use anti-roll bars at the rear. In my May 2011 newsletter, I was responding to a questioner who wondered why we use anti-roll bars at all. I said that when there is ample geometric roll resistance, and a problem with jacking, as in the swing axle suspension at the rear of a Formula Vee, an anti-roll bar is undesirable. When that newsletter appeared as my column in the August *Racecar Engineering*, it was accompanied by an illustration showing a European-style Formula Vee that had a rear anti-roll bar, with a caption pointing out that this is standard practice over there. The car also has an external front anti-roll bar, and rear coilovers arranged to act in both ride and roll, as in older US-spec Vees.

In the US, one never sees this on Vees. Modern ones all have rear springing that acts only in ride. One of the multi-leaf torsion bars at the front is replaced with a solid anti-roll bar, but that is still inside the tube. I was under the impression that external bars at the front were illegal in SCCA.

It turns out that that's not quite true. Here's what the 2011 General Competition Rules say:

<u>9.1.1.C.3 Suspension</u>

A. The front suspension and steering shall be standard VW Sedan as defined herein, or an exact replica of the same material and dimensionally identical. The following modifications are allowed:

1. Removal or modification of spring packs including the use of ride height adjusters incorporated into the front beam provided they are not adjustable from the cockpit. At least one spring pack shall be retained as the primary spring media for the front suspension.
2. The use of any anti sway bar(s), mounting hardware, and trailing arm spacers.
3. The use of any direct acting, tube type shock absorber(s) mounted in a longitudinal, vertical plane and acting through the standard mounting points. Spring shocks and linkage activated shocks are prohibited.

> 9. Caster, camber, and toe in/out settings are unrestricted. Clearancing of carrier or trailing arm to eliminate binding is permitted. Offset suspension bushings and alternate locating spacers are permitted.

There is a front track limit of 52.5" maximum. Track is defined as the distance between rim centers. No height is specified, so this would have to mean the center vertically and longitudinally, as well as laterally. That would appear to mean that the front wheels can have considerable negative camber, without running afoul of the track width rule by being too far apart at ground level.

So it actually would be legal to add an external front bar, and make it adjustable as well, if desired. The front bar on the European-spec car in the illustration is non-adjustable, far as I can see in the photos. The rear one is adjustable, via a series of holes for the drop links to attach to.

The rear bar on the Euro-Vee is very slender and willowy, with fairly long arms, and thus serves as a fine-tuning device. The front bar is visibly thicker, with shorter arms, and therefore inevitably higher in rate.

The Euro-Vee is also visibly longer than a US-spec Vee, and has rack and pinion steering, and pushrod-and-rocker actuated front shocks. Those shocks would be illegal in US Formula Vee, and there would be no room for inboard shocks anyway. SCCA rules call for an 81.5" minimum and an 83.5" maximum wheelbase. That's a bit longer than an Austin Healey Sprite wheelbase, and a bit shorter than a Triumph Spitfire's: pretty short for a mid-engine car. The Euro-Vee driver still sits back near the engine, so the longer wheelbase implies less front percentage on the Euro-Vee.

The Euro-Vee has different tires than the US-spec car. It runs on treaded radials, and they are the same size front and rear. At least in dry to moderately wet conditions, US cars run on bias-ply slicks (radial slicks are prohibited), and the rears are wider than the fronts.

This means the US car has more rear tire relative to front, and more front weight relative to rear, compared with the Euro-Vee. That would explain why the Euro-Vee needs some extra front bar. If the rules allow the front bar to be adjustable, it's hard to see the need for the rear bar. On the other hand, if the front bar has to be non-adjustable, it starts to make sense to have a soft rear bar that is adjustable.

The Euro-Vee also does have a rear Z-bar, and it is apparently stiffer than the anti-roll bar. The Z-bar runs inside a frame tube, so it is impossible to see how fat it is, but the arms are short compared to the a/r bar arms, and the a/r bar is really slim. The car has three rear springing systems: a stiff one that acts in ride only; a soft one that acts in roll only; and a soft one that acts in both ride and roll. Really, a car only needs two of those systems, one would think. It doesn't appear that there is any significant non-linearity designed into any of the three systems.

Would that combination be legal in SCCA? As I read the rules, no. Here's the relevant wording:

B. The rear axle assembly shall be standard VW sedan as defined herein with axle location provided by a single locating arm on each axle.

1. The rear axle tube may be rotated about its axis.
2. Coil spring(s) shall provide the primary springing medium, with telescopic shock absorber(s) mounted inside the spring(s). Cables, straps, or other positive stops may be used to limit positive camber. An anti-roll bar or camber control device may also be used. When said anti-roll bar or camber control device is removed, the required coil springs shall continue to perform functionally.
3. The shock absorber mounts may be modified.

A "camber control device", as used here, is a device which is functionally equivalent to the "camber compensators" sold in the early 1960's for swing axle suspensions: devices that limit jacking by increasing wheel rate in ride without increasing wheel rate in roll. A Z-bar qualifies as such a device. I have seen Formula Vee rear suspensions where a torsional Z-bar was the only springing device. Apparently that is no longer legal. There has to be at least one coil spring; it has to hold the car up; and it has to have a shock concentric with it. But there can be just one; you don't have to have two. I have seen a Formula Vee rear suspension with two additional shocks to damp, but not spring, roll. Apparently that is legal.

So you can have a single rear ride spring, which gives you the same effect as a "camber compensator" – except it's not one legally, because it's the required coilover that holds the car up. You can then add a "camber control device" OR an anti-roll bar – but not both, because the rule is an "or" statement. You cannot have the same combination as a Euro-Vee, but you can have a coilover that acts only in ride, additional shocks that act only in roll, and an additional torsion bar that acts only in roll.

Under what conditions would the anti-roll bar be desirable? It becomes desirable when it is possible to put enough elastic roll resistance on the front to make the car understeer.

Current US-spec Vees corner reasonably neutrally, but their front camber is not optimal. The front grip would benefit from more static negative camber, and/or less roll. The trailing arm front suspension has zero camber recovery in roll. If the outside front tire can be kept more upright, more load transfer at the front can be allowed. That gives us less load transfer at the rear, and therefore more cornering power at the rear as well. If we encounter either inside front wheel lifting or more understeer than we want, we then might benefit from a rear anti-roll bar.

The objective should be to run the front wheels at optimal camber, not degrade their cornering power with poor static camber; minimize roll, as permitted by track roughness, to minimize front camber loss due to roll; then add rear roll resistance as needed to keep the inside front wheel at the point of incipient lift most of the time, or as close to that as front grip will allow without excessive understeer. At least on smooth surfaces, a setup with lots of roll stiffness will have more overall grip due to improved front camber, and some of the front grip can then be traded away as needed to improve rear grip.

However, even though such a setup might involve using a rear anti-roll bar, the rear suspension will still have a substantially greater wheel rate in ride than in roll, unlike other independent suspensions with anti-roll bars.

Taking the whole thing a step further, if the car is not lifting a front wheel when front camber is optimized and the understeer gradient is to our liking, it might be worthwhile to try a set of rear wheel width slicks on the front. They're a bit wide for the required 4" rims, and the steering probably would feel less precise, but the front would probably have more cornering power than with the narrow tires, which could once again be partially redistributed to the rear tires by adding more front roll resistance. Whether this would be worth the drag penalty in such a power-limited class would probably depend on the track.

GENERAL ADVICE ON SETTING DAMPERS

I race a stock car in the U.K. on short quarter mile left hand flat oval asphalt tracks. The car weighs 1440lbs. and currently sprung as follows: i/f 250lbs; o/f 275lbs; i/r200lbs; o/r200lbs. The dampers are double adjusting (bump and rebound). Please can you give me any advice on setting these dampers?

Not sure what a "stock car" in the UK looks like, per your rules. 1440 pounds would be something smaller than what we call a four-cylinder mini-stock over here. I've seen pictures of things from the UK that were called stock cars that were more like what we call modifieds in the US: no fenders on the front, and rear bodywork that only partly covers the rear tires.

Judging by the spring rates, however, this car must have coilovers. The springing would be considered on the stiff side for an oval track car that light.

I cannot offer much specific advice from so little information, but I will offer some general suggestions.

First would be to understand how the adjustments work. Not all double-adjusting shocks are the same. In some cases the "rebound" adjustment is a bleed that actually affects both directions, but affects the rebound more noticeably because the rebound shim stack is stiffer. You really want to get a good set of shock dyno traces for your shocks. Preferably, you should have your own actual shocks dynoed. Failing that, maybe you can get results from the builder showing what they're supposed to act like.

Next general rule would be to avoid using the shocks as a way of tuning handling balance or understeer gradient. You want to do that with springs, tire stagger, and alignment settings, and have the shocks valved to minimize load variation at the contact patches – in other words, to keep the wheels on the ground. For this, you want the shocks as soft as possible without making the driver uncomfortable with the car, or having uncontrolled oscillation over any washboard ripples the track may have. This means you set the shocks so the bump, especially at high shaft velocity, is softer than the rebound (this may not mean fewer clicks on that side – that's one reason you need dyno traces) by a factor in the range of about 1.5 to 3.0, and then you uniformly soften them until the car does something you don't like, or that the driver doesn't like, then back up a little.

The Mark Ortiz Automotive
CHASSIS NEWSLETTER
PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY
October 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

THE DELTAWING CAR

What are your comments on the Ben Bowlby DeltaWing concept as it now appears that a version of this configuration will be built for LeMans in 2012. I have read the article in the August 2011 "Racecar Engineering" and really have a difficult time accepting Mr. Bowlby's concepts as something that will be a workable race car. It appears that the design's method of making the car turn in is an electrically driven differential drive that will provide additional torque at some programmed amount to the outside drive wheel, adjustable stagger if you will. Is this a workable concept?? It certainly appears that the front wheels will have very little effect of turning the car at race speed.

Racecar Engineering has actually carried at least two articles about this concept. The first appeared in the April 2010 issue, and related to the project in its original form, as a proposed spec chassis for Indy car competition.

Do I think the car can win at Le Mans? Probably not, overall. Can it win its class? Undoubtedly, depending on what you call a win. It's in a demonstration class, all by itself. It has no competition. It can't lose. Whether that's winning, I will leave to the reader.

Do I think it can be made to turn in acceptably? Probably it can. In fact, I think the designers are correct that it won't even need the trick diff for that. I think they are correct that the front tires have a long enough moment arm about the c.g. so that they will make enough lateral force to produce acceptable yaw acceleration. The car will be a bit deficient with respect to yaw moment created by toe-out at very small handwheel inputs, compared to one with a wide track, but that won't necessarily slow it down.

The hard thing will be to get it to put power down with the inside rear wheel, while exiting turns. I am sure the designers are correct that traction will be good in a straight line. However, even with a rear track that's a bit on the wide side, the car has a mean track considerably narrower than a conventional car. Overall lateral load transfer will accordingly be greater than in a conventional car, and the front wheels will not be able to contribute a share of the roll resistance moment that is larger than their static weight percentage, as they would in a conventional rear-engined car. This means the inside rear wheel will unload considerably more.

Merely using a limited-slip diff or locker to deal with this would tend to cause a push. So where the trick diff comes in is overcoming the problem without this drawback.

Applying the inside rear brake would also work, but the designers are trying to create an unusually energy-efficient car, so they are looking for a different way. I must confess I don't fully understand from the published descriptions how their diff is supposed to work, so I can't comment on its likely efficacy. However, if it involves power augmentation to the outside wheel, that should have the desired effect on yaw moment, provided it can be managed appropriately and be reliable. If it can, it could have potential for conventional cars too. It would probably require a rule change, because it would be a power augmentation.

It might be possible to link the two output shafts with a motor whose armature spun and was connected to one output shaft, while the field was connected to the opposite shaft. Maybe that's what the trick diff does. I can't tell from published information.

Of course, we will have to see if all this really does work once there is a running car, but theory suggests that the car can be made driveable. Outright tricycles can be made to corner acceptably, provided the wheels are spread out and the center of mass is low and toward the wide end.

But, does that mean the concept makes sense or is an improvement over existing cars? That's a somewhat different question.

To address this broader question, we need to examine the broader nature of this effort, because it is highly unusual in this regard. The whole thing has been a spec car design exercise from the outset, albeit with some unspecified variations permitted, and a sort of "open source" rule to allow multiple suppliers for any given part. It was never conceived as something to run under existing rules, and beat other cars designed to those rules.

Ordinarily, when one sets out to build a race car, the first step is to tentatively pick an existing class, and read the rule book. The second step is to read it again. One studies existing cars in the chosen class, and tries to identify their strengths and weaknesses. One studies the tracks. One tries to figure out ways to work the existing rules. When one gets a bright idea, one tries to predict whether it would immediately be outlawed. One tries to predict what changes to the rules others might make. Other than this, one does not contemplate changes to the rules. One accepts the existing rules, and tries to build a car that can win races under those rules.

The DeltaWing project started out as a response to the need for a new spec chassis for the Indianapolis 500. The event was a given, and so was the track, but the car didn't have to beat any existing cars, or comply with any existing car specifications. The project was to design a vehicle that could put on a good show in the Indy 500, be as safe as possible, be reasonably cost-effective, and address to some degree the environmental issues that racing raises. The vehicle didn't have to beat any conventional vehicles, built by others to a common set of rules. It was to be the only design allowed. Modifications were to be permitted, but all related parts were to be "open source" – documented by drawings published on a class website for all competitors to see and reproduce if desired.

The designers thus faced the kind of challenge hypothetically addressed in some of Paul Van Valkenburgh's later columns in *Racecar Engineering*: if you wanted to design a car to solve the existing problems in racing, and you could write the rulebook too, not just design the car, what would you do?

Three problems that might be addressed would be:

1. The cars use a lot of fossil fuel.
2. The engines are exotic and expensive.
3. It's hard to pass, which makes the races dull.

The DeltaWing car attempts to address the first two concerns by having low drag, so that sufficient speed can be attained to make for interesting racing on a very large, high-speed track, with considerably lower engine power than in current cars. It seeks to address the passing issue by not relying on its front end for downforce, having minimal upwash at the rear, and having small plan view area at the front so that the front will not produce much positive lift either, when it does encounter upwash.

I do think the design looks like it should accomplish those things. However, I am not convinced that the narrow front track is really necessary for this. Rather, I think the objectives could easily be achieved with more conventional wheel positioning – and probably with greater safety, at least with respect to discouraging rollover.

Existing Indycars have little plan view area at the front, and little upwash at the rear, if you just take the wings off.

The narrow front track does look different. This is good for getting media attention, as has already been demonstrated. But if such a layout were to be adopted as a spec configuration, media attention to it would evaporate in short order. A field of look-alike "DeltaWing" cars is no more interesting than a look-alike field of some other sort of cars.

I don't know if the requirement to lap Indianapolis at about 200mph was part of an RFP to which the design was originally a response, but really it isn't necessary for cars to go that fast to put on a good show at Indy. What is necessary is that the cars must need to slow down for the turns and take them at the limit of adhesion, and be able to pass each other and run in close proximity. This is true for other tracks as well. It also helps if there is a diversity of car designs, yet sufficiently close competition to make the outcome unpredictable.

The first year I paid serious attention to the Indy 500 was 1960. As I recall, a lap just over 140mph was good enough for pole. A 150mph lap was considered a distant dream. The race was still an interesting show, and drew a good crowd. Most of the cars, including the winner, were similar in design, but there were a few oddballs to spice things up. All but the Novis had less than 500 horsepower, and the Novis seldom finished.

Wide tires and rear engines raised lap speeds into the 160's. Then came wings and turbos, then ground effects. That was what it took to raise speeds above 200mph. But was the show more entertaining? The wings needed clean air to work well, yet the cars created large wakes that denied a following car clean air. This meant that rather than speeding up in another car's draft, a car slowed down, at least in the turns. That discouraged passing.

The hope is that the DeltaWing can be made to work without wings, yet still corner quite fast, by relying almost entirely on tunnels or undercar venturis. Past attempts at this, with more conventional wheel positioning, have not been very successful. The problem is not that tunnels don't work in the wake of front suspension components. On the contrary, they make lots of downforce, and do it very efficiently. The problem is that it has proved difficult to control the location of the center of lift (downforce) with sufficient precision and adjustability, without adding front and rear wings. Also, a rear wing can help drive tunnels or a diffuser, so there is a synergistic gain that's hard to pass up. And once there's a rear wing, a front wing to balance it becomes highly desirable. Together, the wings can serve to adjust both the front/rear downforce distribution and the overall amount of downforce and drag. Different wing packages can be used for faster and slower tracks.

So wings are highly attractive. They just present problems when two or more cars are running in close proximity. When cars are nose to tail, the front downforce generating element of one is in the wake of the rear downforce generating element of another. When wing packages are regulated by rules that assume that front and rear wings are inevitable, the stage is set for races with reduced passing.

The DeltaWing team are therefore onto something by keeping the downforce generating elements toward the middle of the car. This achieves roughly a car length separation between downforce devices, instead of a few feet or inches. How much difference this can make would best be established by comprehensive testing and/or CFD, but we can get some indication from the example of winged sprint cars.

These cars are allowed large wings, but they have to be on top of the roll cage. They also have little front ones, but they're on top of the nose, and backed by the bluff portion of the nose that covers the induction system, so most of the downforce comes from the top wing. Top wings are adjustable both for attack angle and for front/rear position. The tails are required to look like old-fashioned race car tails, from the days before people caught on to the importance of downforce. They do not generate upwash. The wings do disturb the air going to a following car, but at least some distance between wing and wing is maintained.

The cars are able to run in close proximity. There's lots of passing. The show is generally entertaining.

So, if we want downforce devices amidships, should they be wings? Tunnels? Both?

Tunnels have the advantage of being more efficient in terms of lift/drag ratio. Arguably, they have an aesthetic advantage. However, wings are more adjustable and adaptable.

Both wings and tunnels depend on having the car running more or less straight: they are yaw-sensitive. That is not good. Probably tunnels are worse than wings in this regard, but in both cases, if you get the car sideways, the downforce largely goes away, and if you spin the car and it's going backwards, you can get lift. The downforce devices are your friend until something goes a bit wrong or you make a marginal mistake, and then they abandon you or even turn on you and bite you. And with the higher speeds enabled by the downforce devices, the ensuing crash is a harder hit. So why would you race with wings or tunnels? Because if they are allowed, you can't win without them.

Therefore, there is a strong case for prohibiting wings, tunnels, and other downforce devices relying on free air, altogether. Can this actually be done? It is done, pretty much, in Formula Ford, in midgets, and in unwinged sprint cars. It really is possible to regulate bodywork tightly enough to largely eliminate downforce. The quality of racing in these classes speaks for itself.

Is there any advantage to the narrow front track at all? In terms of passing, it's actually a minus. For comparable roll stability, a triangular-plan car has to be wider at its widest point than a rectangular-plan car. The car's ability to fit through a hole is governed by its width at the wide end. Faired wheels do reduce drag, but rectangular-plan cars can have faired wheels too. The structures required can be outrigger or pannier fairings or fenders, and can add side impact protection.

Triangular-plan cars, with three wheels, do have some interesting possibilities for road use, but only because they are considered motorcycles, and do not have to meet automobile crash standards. This plus the reduced need for torsional rigidity allow a street-legal trike to be considerably lighter than a street-legal 4-wheeler. However, it is likely that the licensing and regulatory status of trikes would change if they became more popular, while their disadvantages in terms of space utilization and roll stability would remain.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

November 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

TEN YEARS IN *RACECAR ENGINEERING*

With the publication of the November 2011 issue of *Racecar Engineering*, I have now had a regular column in the magazine, based on this newsletter, for ten years. The first column was in the December 2001 issue.

FIRST-GENERATION MAZDA RX7 FOUR-LINK

Over here in Australia we have a race class called Improved Production. A popular car is the Gen 1 Mazda RX7 with a live rear axle.

The suspension must retain all the original links and all bushes must be elastomeric -- no spherical bearings. Additional longitudinal links are permitted and lateral location devises are free. Conventional wisdom says that the original links are prone to binding. They consist of parallel lower trailing arms and shorter upper arms about 10in long and angled inwards towards the front. Most people have very soft bushes in the top arms and add another set of parallel top trailing arms.

My question is this---

The track is about 60in. So if the outside suspension is compressed say 3in in a corner and the inside extends 3in the total difference is 6in giving an axle angle of about 6deg compared to the chassis. I accept that this would twist the suspension arms and perhaps cause binding but would the degree of twist not be the same in the new parallel arms as well? ie if the original arms were to bind, would not the new arms also be at risk of binding to the same degree? Does the fact that the original top arms are angled make the situation any worse?

For readers unfamiliar with this suspension, it has four trailing links for longitudinal location and torque reaction, and a Watt linkage for lateral location. The upper trailing links do splay out a bit at their trailing ends, mainly just to clear the springs. This is not a triangulated four-link, where the link angles in plan view are sufficient to allow the four links to provide lateral location without any other mechanism.

The upper links are much shorter than the lowers. This is partly for packaging, but also it makes the side view geometry (and accordingly, the anti-squat and anti-lift) roughly in accordance with Olley's Rule: link lengths inversely to proportional to their height above ground. This makes the longitudinal anti's roughly consistent regardless of ride height.

The problem is that the side view swing arm length fluctuates wildly with suspension displacement. Depending on where the suspension is in its travel range, the longitudinal links try to rotate the axle housing forward with compression, or rearward, or neither, at a varying rate. In ride, no problem results, but in combinations of ride and roll, the longitudinal links on each side of the car try to rotate the axle housing different amounts, and/or in different directions. Any time the links try to create differing rotational displacements of the housing, they effectively turn the housing into an anti-roll bar.

This might be lived with if the effect were consistent (the car has a rear anti-roll bar, so eliminating roll resistance is not the point), but the problem is that this component of roll resistance varies erratically in different combinations of ride and roll displacement. In a straight line, the system behaves well, even if the surface is undulating. In cornering, it behaves well as long as the surface is smooth. However, when presented with hard cornering and a bumpy or undulating surface at the same time, the system produces capricious variations in roll resistance distribution.

It is true that any cylindrical bushing only allows torsional movement by deforming, and that is inescapable if sphericals are prohibited. But that is not the main problem with the stock layout.

In the US, at least under SCCA rules, bushing material or construction is entirely free, but the links still have to be there. It is permitted to add "traction control devices". The most common approach is to use either urethane bushings or sphericals in the lower links, replace the bushings in the upper links with foam rubber, and add a long central upper link, bent to clear the driveshaft tunnel. This effectively converts the suspension to a three-link design with a Watt linkage for lateral location.

That does make the system consistent. It doesn't eliminate torque roll, which should be an objective in live axle suspension design, but it solves the most urgent problem.

Packaging and load paths permitting, my approach might be to react braking torque through a centered top link that acts only in compression, perhaps a shock absorber with a snubber on the shaft, and react drive torque with a link or chain or cable that acts only in tension, offset to the right and angled down at the front.

UNIFIED THEORY OF SUSPENSION SETUP FOR PARTICULAR TRACK?

I am an avid BMW CCA club racer. I have progressed quite far and have even won my championship a few times. However, there are still a couple of guys that I just can't catch. Other than money, I think one of the reasons they are beating me is setup. They are very good at nailing a setup for a given track.

I have a good baseline which was given to me. I soften it a bit for bumpy tracks like Mosport or Sebring. But I pretty much stick to it. Once in a while, all the planets align and I set a lap record. But it's always a surprise to me. I think this is my little setup lottery.

I am an engineer by schooling (electrical). I have come up through the BMW and Porsche clubs where I have been instructing for years. I understand all the underlying principles and have read all the books. But I have no idea how to unite them all into a unified theory of suspension setup. In this respect, I am quite typical of club racers and probably a good chunk of your readers.

Do you think you can give me some guidance toward developing these skills? Where should I look?

My car is a BMW E36 M3, 410hp, 2200lbs, Moton 3-ways, slicks, etc...

I don't see any way to reduce something that complex to a brief set of principles or rules, especially if those have to apply to a wide variety of vehicles. For production-based cars with tightly regulated bodywork, it does get a little simpler than for winged cars, because at least we aren't dealing dramatic variations in the aero properties, but even then it isn't simple.

One universally applicable principle would be to know why you are trying a given setup change: what are the physics? What exactly are you trying to make the car do or not do?

The basic idea of having a baseline setup and softening it a bit if the track is rough makes sense. Depending on the car and the rules, it is often necessary to use higher static ride heights for the softer setup.

In addition to general roughness of the track, factors that may influence setup include:

- Dominant turn direction. Some tracks have most of their turns in the same direction.
- One or two turns being more important to lap time or passing opportunities than the rest. Most commonly, such turns are at the beginning of major straightaways.
- Predominance of tight turns or sweepers.

Issues such as these get particularly interesting if we have ballast we can move. Production-based cars may or may not have ballast, depending on the rules.

When one turn direction is dominant, or when one turn is unusually important, sometimes it pays to move ballast toward the inside of such turns. It may cost us elsewhere in the lap, but we may come out ahead overall.

When there are lots of short-duration tight turns, braking before those turns and forward acceleration after them become more important, and mid-turn speed becomes less important. Assuming rear wheel drive, more rear percentage will help the car put power down, and will also shorten braking distances.

When there are a lot of sweeping turns, and few long straights and tight turns, it may pay to move ballast forward a bit more, to work the tires more evenly in steady-state cornering.

With live-axle rear suspension, it often helps to use a bit less than 50% diagonal (meaning LR + RF), to compensate for torque wedge. If the turn before the most important straight is a tight right-hander, there is a case for being a bit more aggressive with the diagonal than we might otherwise.

But again, in all of this it is of paramount importance to know what one is trying to achieve, and have a clear understanding of how one expects a setup change to accomplish that.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

December 2011

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE/LESS ANTISQUAT – OR WHAT? – TO FIX WHEELSPIN/REAR TIRE WEAR

I race asphalt super late models. The car is good on the short run but I'm having a big problem keeping the rear tires under me during long runs, although it seems decent on restarts when the tires cool. When I check the tires after the race, the rear tires end up almost 3/4" smaller in circumference than the fronts when they started the race approximately the same size. I'm not sure if I'm burning the rear tires up because I am:

(A) overloading them (I used to have so much antisquat that the rear suspension lifted on acceleration, keeping the CG higher, and transferring more weight... I've since backed off on this until now where the car squats a little bit - I have a camera in the RR to see all this)

or

(B) not loading them enough (because suspension squat drops the CG and so less weight is transferred) and this lack of load is causing wheelspin in the long run. Is there any way I can figure out which scenario is going on?

I've thought a lot about it and here's how I break it down:

(To fix A) Benefits of less antisquat (more rear suspension travel) for forward bite:
1. more reverse rake in panhard bar (more load transferred to LR off of the RR)
2. lower rear roll center (decreasing rear roll stiffness)
3. more rear steer to the left (based on where I have my trailing arm angles)
4. and possibly a longer response time in weight transfer (since the weight must go through the springs rather than instantaneously be transferred to the tires through the suspension links, letting the tires transfer from mostly cornering to mostly accelerating force without going outside the traction circle)

(To fix B) Benefits of more antisquat for forward bite:
1. maintaining or increasing rear chassis height to get higher CG and more weight transfer to back tires...usually good for bite.

What do you think? I know trial and error is sometimes best but I'm away at school most of the year so my chances to race and try things are rare. It may also be helpful to know that I run a spring-loaded third link at about 10° angle (1200lb. spring, preloaded ~12 turns) and the car is 2750 lbs. and 51% front weight.

First of all, anti-squat does not have a big effect on rear tire loading. To see just how small the effect is, let's do some "back of the napkin" numbers.

Suppose we have a 108" wheelbase car where, at 0.8g forward acceleration and with 100% anti-squat (no rear suspension ride deflection), the c.g. height is 15". The rearward load transfer is then 0.8 (15/108) of the car's weight, or 11.1%. If the car had 49% rear statically, it will then have 60.1% rear dynamically.

Now suppose we add enough anti-squat so the rear lifts about an inch under the same power application, which will raise the c.g. about half an inch. Rearward load transfer becomes 0.8 (15.5/108) of the car's weight, or 11.5%. That gives us 60.5% rear dynamically – not much change. Or, if the rear squats an inch, we have 0.8 (14.5/108), or 10.7%, and 59.7% rear dynamically.

So there is an effect, but the difference between hardly any anti-squat and a really large amount (in pavement car terms anyway) is less than a percentage point of dynamic rear percent.

Just for grins, let's also consider a dirt Late Model that jacks the rear end up 3" under power, has the same 108" wheelbase, 55% rear statically, and a c.g. height under power with 100% anti-squat of 17". If the dirt is really tacky, maybe this car can also attain 0.8g forward acceleration. With no anti-squat, rearward load transfer is 0.8 (17/108), or 12.6%, for a dynamic rear percentage of 67.6%. With 3" of jacking, the c.g. height is around 18.5". Rearward load transfer at 0.8g forward is then 0.8 (18.5/108) of the car's weight, or 13.7%, for a dynamic rear percentage of 68.7%. That's three inches of rear jacking, delivering just over one more percentage point of dynamic rear, on unusually tacky dirt. Again, enough difference to be worth a little, but not enough to make a difference between burning up the rear tires with wheelspin, and not.

The questioner doesn't mention what kind of Super Late Model this is. I know it's a pavement car because there is no class by that name for dirt cars. However, there are a number of sanctioning bodies that have such classes, and their rules vary. At the nearest short track to me, Concord (NC) Motorsports Park, the rules are NASCAR's. There are three Late Model classes: Limited Late Model, Late Model, and Super Late Model. They all look similar, but the engine rules are different, and maybe the tire rules as well. The Super Late Models are a touring series and don't run there every week. The speed difference between them and the rest is dramatic – easily visible from the stands. Their events are also much longer: typically the Limited Late Model feature is 30 or 35 laps, the Late Model feature 50 or 60 laps, and the Super Late Model feature 100 or even 150 laps, on a half-mile tri-oval.

Then there are other sanctioning bodies, mainly in the Midwest, that have Super Late Model classes with much looser rules, which allow, among other things, three-link rear ends (NASCAR requires truck arms) and bigger rear spoilers. The questioner indicates he is running a 3-link, so I know he's not running NASCAR.

Conversation with fellow competitors should quickly reveal whether the amount of rear tire wear and heating described here is actually abnormal in this class. ¾" less circumference at the end of the race translates to about 1/8" more tread thickness worn off the rears than the fronts. For a powerful car on a short track, over a fairly long event, that may not be anomalous.

From the standpoint of theory, there are a number of things that may make the rear tires wear faster than the fronts, even if the car is not loose and even if the static rear percentage is not greater than the front. Wheelspin is one possibility, and if the driver experiences wheelspin, then we know that's a factor. However, suppose the car has considerably more aerodynamic downforce at the rear than at the front. That will tighten it (add understeer) beyond what we would expect from the static weight balance. We can tune the setup to compensate for that, mainly by statically and/or dynamically de-wedging the car, or adding rear tire stagger. However, when the rear tires are helped by the aero balance and then loaded more unequally to make up for that, they are working harder than the front tires, and will heat and wear more, even ignoring the work they have to do propelling the car.

Ideally, in such a situation we'd add front downforce, but the bodywork rules may prevent that. Pavement Late Model rules generally allow only a valance at the front, with varying requirements for front overhang and static ground clearance, but a fairly large spoiler at the rear. It may actually be that in such a case it would be possible to create a setup that would run at its own limit for a longer time without taxing the rear tires so heavily, by reducing rear downforce and adjusting the setup to re-balance the car. However, chances are such a setup would have lower limits, i.e. be slower – so it makes more sense to tune for better speed and accept the fact that the rear tires will go away more than the fronts on a long run.

I mentioned stagger. That definitely enters into this, especially when running a spool. (NASCAR Late Models use lockers, but spools are allowed, and common, elsewhere.) Most spool setups use more stagger than theoretically necessary for the turn radius. Even the theoretically correct stagger for the turn radius causes the rear tires to fight each other down the straightaways: the left one drags and the right one has to slip more and provide all the propulsion. That will definitely heat and wear the rear tires. The front tires may also have some added work compensating for the yaw moment the stagger induces, but generally stagger will not add as much workload for the fronts as for the rears.

So one way to help rear tire life, and consistency on long runs, is to minimize rear stagger, especially with a spool. This does have a price, however. A setup with ample stagger is easier to drive because it has more controllable rear breakaway. A setup with minimal rear stagger is twitchier at the point of rear breakaway. So to some extent there is a tradeoff between short-run driveability and consistency over long runs.

Notwithstanding all of these complexities, one thing jumps out at me from the question. I am wondering why the questioner is running 51% front, if there is a wheelspin problem. Late Models typically have quite a bit of ballast, which makes it possible to get more rear percentage. If the track has tight turns and long straights, the car should be faster with more static rear percentage. If the track is really bowl-shaped, then maybe not. But if wheelspin is an issue, that suggests that putting power down out of turns and down straightaways is an issue. More static rear percentage will help that, much more than anti-squat.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

January 2012

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

EQUAL WHEEL RATES RIGHT AND LEFT, OR EQUAL FREQUENCIES?

I race a front wheel drive car. In stock form the corner weights are really skewed with both the engine and driver located in the front left of the car. There are 2 schools of thought as far as the best way to address when building a race car.

A) Run equal spring rates on each axle pair. Then corner balance the car as best as possible and live with unequal static ride heights.
B) Set the ride heights equal left and right. Then run different spring rates on each corner based on the achieving equal suspension frequencies per axle pair.

What's your opinion?

If the suspension is adjustable for ride height, by any means at all, it should be possible to get any desired ride heights regardless of spring rates or static wheel loads. Or, more precisely, it should be possible to get any right/left tilt, and any fore/aft rake, and any average height, with any diagonal percentage.

But in any case, when the car is significantly left-heavy, and it turns both ways, is it better to have similar wheel rates left and right, or similar frequencies right and left?

We might mention first of all that there are sprung mass frequencies and unsprung mass frequencies. If the unsprung masses are similar on both sides of the car, but the sprung masses are different, then there is no way to get both the sprung and unsprung frequencies identical on both sides. But the questioner is referring to the sprung mass frequencies.

In terms of ride dynamics, it is desirable to have the right and left sprung mass natural frequencies similar to each other. In extreme cases, a lightly damped car may "gimbal", or experience some roll and yaw with pitch, when subjected to sequential bumps at the front and rear, or when subjected to lateral, longitudinal, or yaw jerk (abrupt acceleration change). With the relatively heavy damping used in race cars, tuning sprung mass frequencies is less important, but it's still desirable to spring for flat ride when possible, and that means having similar frequencies right and left.

There is another reason for going stiffer on the heavy side, and it has nothing to do with getting the frequencies right for flat ride. When the car is left-heavy, it tends to turn right under braking and left under power. That means that when trail braking into a left turn, the car tends to be tight (understeer), and when powering out, it tends to be loose (oversteer). In a right turn, it's the opposite: loose in, tight out, compared to a left turn.

When the car is stiffer on the left, that means the left side sees more load transfer than the right in longitudinal acceleration. Under braking, the car de-wedges (loses inside rear + outside front percentage) for a left turn and gains wedge for a right turn. Under power on exit, the car gains wedge in a left turn and loses wedge in a right turn. There is no guarantee that this effect will accurately compensate for all effects of left-heaviness when the wheel rates give equal sprung mass frequencies right and left, but at least the compensation is in the needed direction.

In the real world, we will usually be constrained by available spring rates, so even achieving truly equal sprung mass frequencies right and left will not be possible – we will just try to come fairly close, using springs we can buy. If we want to be really fastidious, we will test individual springs to determine their actual rate, rather than relying on the advertised rate. In any case, I would not be afraid to go a bit stiffer on the heavy side of the car.

TWO-WHEELER THAT DOESN'T LEAN, BUT SHIFTS WEIGHT LATERALLY?

An idea that one of my friends came up with, and has been investigating for a while, is a "bike" with race-car slicks. For his machine he was thinking of utilising a pair of very wide low-profile tyres. They would not camber like regular motorcycle tyres do. They would stay perpendicular to the road, always upright. His idea is that the main body of the "bike" would not lean but would laterally translate relative to the wheels. For example, in a left turn the body of the "bike" (engine and "rider" and all) would be displaced to the left far enough to exactly counter the rightwards overturning tendency. He's "designed" (I should say sketched) his suspension system. It relies on power hydraulics and electronics to operate.

My feeling is there may be a better, more "natural" way to achieve the desired result. Instead of shifting the body laterally in response to the front wheels steering (after they have steered), thus allowing cornering forces to already be generated, it would be better if the body could be moved laterally just before the front wheels steer.

This idea is similar in a sense to what Phillip James' FTC three wheeler is doing. See the site http://www.tiltingvehicle.net/technical.html and also patents WO2005075278 and US2008238005- both cover the same material wherein the steering of the front tyre/s is the result of the vehicle *leaning (shifting mass). Note: the steering does not initiate the lean, it is the leaning which causes the wheels to steer.*

This is much like the bicycle servo-steer mechanism outlined in the excellent Pierre Ethier SAE paper and also on website http://www.clevislauzon.qc.ca/Professeurs/Mecanique/ethierp/2-wheels/index.htm ... Again, the mass is moved (leaned or tilted), THEN the steering of the front wheel develops as the result.

I am thinking that something similar could be achieved by sliding the body sideways and having the front wheel steer occurring as the result of that action. I suspect that it may be possible that the correlation between lateral translation of the vehicle body and front wheel steering could be arranged to occur by utilising a mechanical linkage system. If this is possible it would negate the need for much of the electronic sophistication my friend wants to rely on. He'd likely still need the power hydraulics if the vehicle weight was high though.

Any comments or ideas on this?

The second paper cited alludes to nine different theories of how the rider of a two-wheeler initiates a turn. I have not studied this rather complex question to the degree that others have, and I certainly have not taken videos of motorcyclists, and slowed them down and analyzed the motions, but I will attempt some reasonably intelligent commentary nonetheless.

The case of a narrow three-wheeler or four-wheeler that is leaned simply to avoid rollover is different from the case of a two-wheeler. Also, it matters a lot whether the wheels lean, especially with a two-wheeler. And in all of these cases, the needed lean angle or lateral c.g. shift of the vehicle is not a simple function of front wheel steer angle.

In the case of a two-wheeler, the vehicle is unstable in roll. That is, it will not stand up on its own when not in motion. Colin Chapman reputedly said, “Silly things, motorcycles. If you let go of them, they fall right over.” In the case of a narrow three-wheeler or four-wheeler, even with suspension that’s very soft in roll, the vehicle does not fall over unless the line of action of the vector sum of all x, y, and z forces acting on the c.g. intersects the ground plane outside the polygon described by the contact patches. When the vehicle is not single-track, the suspension can roll the vehicle within certain limits, because it has a base of support of some width to push against. When the vehicle is single-track, everything depends on a complex inertial balance. We can’t just shift the c.g. laterally, without something else happening.

Let’s first consider what keeps a bike upright when it’s moving. The bike is always trying to tip one way or the other, because its center of mass is above the line of support described by the two contact patches. But as soon as the bike develops a roll velocity, gyroscopic precession induces a yaw acceleration that tends to bring the contact patches back under the c.g. This doesn’t make the vehicle hold itself upright for any distance without rider intervention, but it does make it rideable by most humans. If we push a trike and just let it roll, it will go sort of straight for a while, probably turn to one side eventually, and stop without falling over. If we push a bike and just let it roll, it will go a short distance, then tip to one side, turn to that side but not enough to catch itself, and fall over. The bike will have an increasing roll velocity, and an increasing yaw velocity, but the yaw won't keep up with the roll to the point where the vehicle will right itself.

The fact that the bike is unstable in roll makes it easy to initiate roll acceleration. This can be done in more than one way. If the roll acceleration needed is fairly small, just letting the bike tilt the desired way is sufficient. The rider's job is not so much to make the bike lean or turn as to arrest the lean and turn as needed to keep the vehicle upright and on course.

In all cases, however, the dynamics of the vehicle are highly dependent on precession. When the wheels are leaned, they try to yaw in the direction they are leaned. When the wheels are yawed without being leaned, they try to lean opposite to the direction of yaw. This means that if we lean the wheels into the turn, the bike will yaw in the desired direction. But if we try to steer the front wheel and keep the wheels upright, precession will roll the bike out of the turn.

The two main ways that a person initiates a turn on a bike are to roll the bike into the turn with the hips, and to use front wheel precession to roll the bike into the turn by momentarily steering the handlebars out of the turn. The former method is used most of the time. The latter method is used when an abrupt turn is required, usually to dodge an obstacle.

A bike can steer like a two-track vehicle at very low speed. The front wheel is steered a lot, the bike leans very little, and the rider balances the bike as best he can. The maneuver is only slightly less difficult than a track stand (balancing the bike essentially motionless). If we try to steer a bike like that at speed, gyroscopic precession will make the bike lean out of the turn and steer the opposite way than we intended. Indeed, the main element of learning to ride a two-wheeler is to learn not to try to steer it that way.

It may be that some genius can devise a way to make a two-wheeler turn at speed without leaning the wheels into the turn, but I would not bid on a contract to develop the algorithm, and I certainly would not try to do it with any mechanism that mechanically shifts ballast or sprung mass simply in relation to front wheel steer angle.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

February 2012

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

ANTI-DIVE, AND THE LOTUS REACTIVE RIDE HEIGHT SYSTEM

Current F1 news includes reactive anti-dive front suspension being declared illegal. Why don't they simply employ enough (100%?) anti-dive geometry to do the job. In other words, what are the negative side effects of running high levels of anti-dive geometry in road racing?

Also, can you please verify for me if this is an accurate statement:

With anti-dive the axes of the upper and lower control arms are not parallel. When the suspension compresses, the upper ball-joint moves further negative on the X axis of the car than the lower ball-joint. This leans the spindle rearward relative to the chassis (adds caster). The torque of the brake caliper tries to rotate the spindle forward relative to the chassis (reduce caster) which results in the suspension extending. These countering forces are what makes anti-dive work: compressing suspension adding caster versus caliper torque trying to reduce caster.

Taking the last part first – what makes anti-dive as we know it work – what the questioner is describing is sometimes referred to as torque anti-dive. It is indeed related to the instantaneous rate of caster change with respect to suspension displacement. There is also thrust anti-dive, which relates to longitudinal, or x-axis displacement of the hub or wheel center.

This conceptual framework applies if we think of the forces as acting at the wheel center. I find it simpler to think of the forces as acting at the contact patch. In that case, all anti-dive relates to the instantaneous rate of longitudinal displacement with respect to suspension displacement, for the contact patch center, with the brake locked. This rule covers all cases, even including an inboard brake with drop gears in the upright.

The rule can be expressed by the following equation:

$$dF_z/dF_x = dx/dz \quad (1)$$

where:

F_z is jacking or anti-dive force induced in the suspension
F_x is longitudinal force at the contact patch
x is longitudinal displacement at the contact patch
z is vertical displacement of the suspension, at the contact patch

dF_z/dF_x can also be called the jacking coefficient for braking. On a kinematics and compliance (K&C) rig, we can measure dx/dz by cycling the suspension freely with the brakes locked and the sprung mass constrained longitudinally and laterally as it is moved vertically, and measuring longitudinal displacements of the wheel support pad. We can measure dF_z/dF_x by holding the sprung structure at fixed ride height (or a series of fixed ride heights) with the brakes locked, applying rearward force at the contact patch, and measuring change in vertical load at the wheel support pad. Measured results will not follow the equation exactly. The differences between predicted and measured values will give us some indication of the compliances, clearances, and frictional effects in the system.

When we are designing the car, or analyzing an existing car from point measurements, for outboard brakes the car has positive anti-dive if the side view instant center (SVIC) is either behind the wheel and above ground or ahead of the wheel and below ground, or if the SVIC is undefined (side view projected control arms parallel) and the side view projected control arms slope upward toward the rear.

For inboard brakes, assuming no gears in the upright, the rule is the same, with one important change: the car has positive anti-dive if the side view instant center (SVIC) is either behind the wheel and above wheel center or ahead of the wheel and below wheel center, or if the SVIC is undefined (side view projected control arms parallel) and the side view projected control arms slope upward toward the rear. Stated another way, the linkage or control arm system can only have torque anti-dive when torque reacts through it rather than directly to the sprung structure via a jointed shaft.

So anti-dive geometry in the suspension linkage must either create caster change with heave, or make the wheel move forward with heave. In the case of an inboard brake, only the latter of these will produce anti-dive.

It is worth noting that absence of caster change with heave (equal displacement at all four wheels; vertical translation of sprung mass) does not mean the caster never changes. Controlling caster change in both heave and pitch is very much like controlling camber change in both heave and roll: we can’t have zero change in both modes. The best we can do is compromise so we don’t have a lot of change in either mode. One recommendation I often make is to have the side view virtual swing arm length (SVSA) about equal to the wheelbase. That gives similar amounts of caster change per inch of wheel travel in heave (caster increase) and in pitch (caster decrease).

What limits how much anti-dive we can run? Two effects, basically. First, we tend to get wheel hop with large amounts of anti effect, of any kind, mainly at the point of wheel slip. Second, anything that makes the wheel move forward with respect to the sprung mass when the suspension compresses makes the suspension less able to absorb bumps. When the wheel hits a bump, it is best if it can move rearward as well as upward. If it has to move forward to move upward, that makes the suspension less compliant, and increases ride harshness and wheel load variation.

Briefly, that's how anti-dive geometry as we have known it works. What is this reactive ride height control thing that Lotus came up with, that has been recently in the news for getting banned by the FIA? How does it work? Would it offer advantages over conventional anti-dive? Would it be applicable outside of Formula 1?

The device uses a small hydraulic system, actuated by brake torque, to raise the front end a little bit, compensating for compression under braking. Rather than being attached to the upright in a completely rigid manner, the caliper is mounted to the upright on a bracket that can rotate with respect to the upright – somewhat like a brake floater on a beam axle. The caliper bears against a piston, or pushrod acting on a piston, in a small master cylinder attached to, or built in unit with, the upright. Under braking, the master cylinder sends fluid under pressure through a short hose to a slave cylinder built into the lower end of the suspension pushrod. With sufficient hydraulic pressure, the pushrod extends, raising the front ride height.

An anti-dive effect is thus achieved, without any wheelbase or caster change in heave. If such a system is tested on a K&C rig, Equation (1) will not apply. There can be a positive jacking coefficient without having a locked wheel contact patch moving forward when the suspension is compressed. In fact, if the suspension geometry provides zero anti-dive as conventionally analyzed, the system will show slight rearward motion at the contact patch as the hydraulics work, when rearward force is applied to the contact patch. The system will create a compliance. However, this particular compliance will be accompanied by an increase in wheel load if the sprung mass is not allowed to rise, or a ride height increase compared to behavior without the system if the sprung mass is not constrained vertically.

The FIA banned the system under the rule prohibiting movable aerodynamic devices. Some have suggested that they should have used another rule that prohibits suspension devices primarily intended to influence aerodynamics. Personally, I think either of those is a reach, and the latter rule is highly ambiguous. Claiming that a suspension system is a movable aerodynamic device in the same sense as a movable wing or a sucker fan is baldly absurd. It may be that on a very smooth track, with ground-effect-critical wings and bodywork, the suspension does affect aerodynamics as much as it affects anything. But in that case, anything at all in the suspension is a device that influences aerodynamics. Certainly a third spring in the front suspension is that. Those were introduced after the advent of wings, to limit pitch, and control how far front wings are from the ground. But they're legal.

Apparently, interpretation of this rule depends on a judgement of what the primary intent of the device at issue is. I would still say it's a big reach to say that control of aero properties is the primary function of any kind of passive anti-dive strategy, when cars having no aerodynamic downforce devices at all use various strategies to limit dive, and so do motorcycles. F1 cars had anti-dive before they had wings. This is just a slightly different way of passively harnessing the forces in braking to reduce front suspension compression.

F1 legality aside, does this idea offer functional advantages? Can it do anything that ordinary anti-dive cannot?

Lotus is saying this is all much ado about nothing, because they tested the system and didn't like the way it behaved, so they weren't going to race it anyway. That may be so, but I don't think the matter is as simple as that.

I would expect that, as with so many things, brake-reactive ride height modification can hurt or help the car, depending on the details of how it's applied and its interaction with other design and setup elements.

Although I would defer to actual experimental results on this, I am inclined to suppose that anything that produces a large jacking coefficient will cause wheel hop or chatter at the limit of adhesion, and also exact some penalty in ride quality and wheel load variation when braking, even if the jacking coefficient is obtained without caster or wheelbase change in heave. I would also expect that this propensity would depend on the overall jacking coefficient of the front suspension system – meaning the total from the brake-reactive elements plus any conventional anti-dive or pro-dive.

Most F1 cars nowadays have little or no anti-dive or pro-dive. If we simply add a reactive anti-dive system to an existing F1 car, we would either end up with a fairly small effect (if the reactive system does not create a large jacking coefficient), or end up with wheel hop or chatter in limit braking.

But since everything in a chassis interacts with everything else, we don't necessarily get a meaningful read on an idea's true potential by simply bolting it onto an existing car and seeing if the driver likes it or if the lap times come down. What if we designed the car to have reactive anti-dive, and took advantage of the system's properties by changing other things?

For example, suppose we designed the front suspension with a side view instant center near the vertical rear axle plane and well below ground. The wheel would then move rearward considerably as the suspension compresses, and the system would have considerable geometric pro-dive. Suppose we then combined this with enough reactive anti-dive to give us just a modest amount of net anti-dive. We would then have improved ride and reduced wheel load variation most of the time, with braking behavior about like other cars.

We would get this without the disadvantages of compliance struts as used in passenger cars to allow the wheel to move rearward on bumps. Compliance struts allow compliance caster change in braking and also in cornering due to the rearward force the tire generates when it is steered and running at a slip angle. This in turn results in compliance camber change. Those effects are the reason compliance struts are not used on race cars. Even for a passenger car, where the benefits of compliance struts are deemed worth the penalties, the combination of pro-dive geometry with reactive anti-dive could allow really low impact harshness, and/or allow lower control arm compliance to be reduced, resulting in a handling improvement without ride penalty.

Reactive anti-dive is just as suitable for inboard brakes as for outboard ones. Ordinarily, the only way to get anti-dive with inboard brakes is to make the wheel move forward when the suspension compresses. With reactive anti-dive, we can have a wheel that moves rearward in compression, and any amount of anti-dive we want. In many cases, we might even be able to dispense with hydraulics. Hydraulics are pretty much inescapable if the brake has to steer with respect to the suspension, but if the brake does not steer and is part of the sprung mass, in many cases we will be able to get the effects we want with purely mechanical actuation.

Compared to conventional anti-dive, reactive anti-dive is a roundabout way of getting the job done, and it does inevitably involve using more parts. However, as with other roundabout ways of doing things (e.g. pushrod and rocker suspension), the extra bits afford a convenient way of introducing intentional nonlinearities into the system's behavior. With either hydraulic or non-hydraulic actuation, we can use preload springs, limiting springs or stops, and linkages with rising or falling motion ratios to get all sorts of non-linear anti-dive. We can, for example, have anti-dive that is extremely aggressive for small amounts of brake torque, then drops away to near zero or even goes negative for the high ranges of brake force where we are likely to encounter chatter or wheel hop.

Therefore, it is my opinion that brake-reactive anti-dive has a future, perhaps most of all for street use. In outlawing it for F1, the FIA is passing up an opportunity to have racing justify its existence by improving the breed.

The Mark Ortiz Automotive

CHASSIS NEWSLETTER

PRESENTED FREE OF CHARGE
AS A SERVICE TO THE
MOTORSPORTS COMMUNITY

March 2012

WELCOME

Mark Ortiz Automotive is a chassis consulting service primarily serving oval track and road racers. This newsletter is a free service intended to benefit racers and enthusiasts by offering useful insights into chassis engineering and answers to questions. Readers may mail questions to: 155 Wankel Dr., Kannapolis, NC 28083-8200; submit questions by phone at 704-933-8876; or submit questions by e-mail to: markortizauto@windstream.net. Readers are invited to subscribe to this newsletter by e-mail. Just e-mail me and request to be added to the list.

MORE ON REACTIVE ANTI-DIVE

Last month we considered the Lotus reactive anti-dive or ride height modification system that was recently outlawed in F1. I have been giving more thought to the details of applying this idea.

I have decided that Lotus made a fundamental mistake in putting the slave cylinder in series with the ride spring, by placing it in the suspension pushrod. For a reactive system to work in a beneficial manner, it needs to act in parallel with the spring, as conventional anti-dive does.

This might be done by arranging for a slave cylinder to act on the rocker in parallel with the main coilover or torsion bar. With an inboard brake, it could be done by having mechanical linkage to the rocker from the caliper floater.

When the slave cylinder is in the pushrod, or otherwise in series with the ride spring, the slave cylinder will not extend until there is sufficient braking to overcome the load on the pushrod or whatever the slave cylinder acts on. This means the system generates no anti-dive on light brake application, and then has to generate a lot of anti-dive on heavy brake application if it is to have a useful effect on ride height. If anything, it would be better to have the system generate lots of anti-dive on light brake application, and then not so much in the range of brake force approaching wheel lockup. At a minimum, we'd like the behavior to be roughly linear.

It might be possible to add preload springing to the system so that the slave cylinder was preloaded to the point where it was just short of extending at static condition. But then the preload would not have similar effect when pushrod loading was something other than the static value. When greater load was on the pushrod, we would still get anti-dive effect only beyond some braking force threshold. When lesser force was on the pushrod, we would have an undamped spring in series with the main ride spring.

If the mechanism acts in parallel with the ride spring, we can have it affect ride height from the smallest amounts of braking force. We can also create nonlinearities with preload or limiting springs, but we will inevitably have those acting in non-braking situations as well.

So to get the best from this concept, not only should we make it act in parallel with the ride spring, but we should also be wary of unintended consequences and avoid getting carried away by the temptation to introduce clever nonlinearities.

Even within those constraints, the concept still offers the possibility of obtaining anti-dive while letting the contact patch move rearward somewhat in compression when not braking, thereby improving the suspension's ability to ride bumps without the usual penalties. It would be possible to arrange for the slave cylinder or perhaps a mechanical linkage to act as, or in series with, a third spring rather than individual wheel ride springs. This would mean that side-to-side differences in anti-dive force generated by the brake-reactive system would not affect dynamic diagonal percentage.

If the system is arranged so it acts in parallel with the ride spring, and so that it creates anti-dive for all levels of braking force, that means that the caliper bracket will always rotate whenever the suspension moves, although it may be possible to make the relationship somewhat nonlinear with respect to suspension displacement. This motion of the caliper will create some additional unsprung mass inertia, which is generally not desirable, but the effect will generally be small. The added inertia will be present whether the brake is inboard or outboard.

This also means the system will cause the contact patch to move forward in compression when the suspension is cycled with the brake locked on a K&C rig, and the rate of forward displacement with respect to vertical displacement will equate to the jacking coefficient in the usual manner. However, it still will not be possible to infer the amount of overall anti-dive purely from the suspension geometry. It will be possible to have a side view instant center that would normally create pro-dive, and still have some net anti-dive from the system as a whole.

Furthermore, when the reactive system acts in parallel with the spring, it does not have to generate forces in the same direction as the spring. It can generate a downward jacking force just as easily. That means we could have brake reactive anti-lift at the rear.

It would even be possible in some cases to have drive torque reactive anti-squat and anti-lift with independent suspension. This would entail mounting the final drive so it can rotate, and using that rotation to energize some sort of system acting in parallel with the springing.

This might not be particularly useful at the rear of the car, however, because it is possible to get anti-lift and anti-squat the conventional way by having the contact patch move rearward a bit in compression, unlike the front of the car, where it ordinarily has to move forward. Reactive anti-lift under power at the front might have some potential, when the front wheels are driven.

FORCES ON CONTROL ARMS WHEN CORNERING

In a double A arm suspension are the forces opposite for the upper A arm vs. the lower A arm? Example: in a left hand turn is the force for the right front A arms inward on the lower and outward on the upper? I hope I made sense because I really want to know. I have watched "Minding Your Anti" [my video of a lecture I gave at UNC Charlotte in 2003, still available on DVD for US$50] *a dozen times. It has helped me immensely. I just want to make sure I understand the forces correctly.*

For this discussion, we will assume that the suspension is of a type that has an upper control arm and a lower control arm, or can be approximated as such for modeling purposes. We will assume that the lower control arm plane intersects the wheel plane below the wheel axis and the upper control arm plane intersects the wheel plane above the wheel axis. We will assume that there are no drop gears in the uprights.

If we are talking about the forces on the control arms induced by ground plane forces (longitudinal and lateral forces at the contact patch, the forces that create geometric anti-roll and anti-pitch effects), for the most part the answer is yes, the forces induced in the upper and lower control arms are opposite in direction. The main exception would be the case of longitudinal forces from braking or propulsion, where the torque is applied to the wheel through a jointed shaft and only the thrust acts through the suspension linkage – that is, propulsion with a sprung final drive (independent or DeDion suspension) or braking with inboard brakes.

With regard to the lateral (y axis) forces, it is necessary to remember that there are usually tension and compression loads on the upper and lower arms in static condition, just from holding the car up and holding the wheel in position. Ordinarily, in a front suspension the ball joints are inboard of the wheel plane, and the ride spring acts on the lower control arm. That means there is a bending load and a tension load on the lower control arm and a compression load on the upper control arm, when the car is not doing anything but resisting gravity. The loads from cornering or braking are additive to (or subtractive from) the static loads.

On the outside wheel when cornering, the y axis ground plane forces will reduce the tension load on the lower arm and the compression load on the upper control arm. There will also be some increase in the normal or z axis force, which will have an opposite effect. As long as the vector sum of the z and y forces has a line of action that is outboard of the upper ball joint, the lower control arm sees tension and the upper control arm is in compression. When the vector sum line of action is inboard of the upper ball joint but outboard of the lower ball joint, both control arms are in compression. If the vector sum line of action passes inboard of the lower ball joint, the lower control arm is in compression and the upper is in tension.

When the vector sum line of action passes through the upper ball joint, the lower arm sees neither tension nor compression, and the upper arm sees compression. That is, there is no moment about the upper ball joint to generate a force at the lower ball joint, but there is a moment about the lower ball joint that can generate a force at the upper one. When the vector sum line of action passes through the lower ball joint, there is no compression or tension load on the upper arm, and there is a compression load on the lower arm.

But when we are considering jacking coefficients for x and y axis forces, for purposes of determining geometric anti-roll and anti-pitch effects, we are concerned with the changes from static conditions. For y axis forces, for an outside wheel the changes due to ground plane force are always in the compression direction for the lower control arm and in the tension direction for the upper control arm. For an inside wheel, the changes are always in the tension direction on the lower arm and compression on the upper.

For braking, if the brake is outboard there will be a rearward force at the lower side view projected control arm and a forward force on the upper side view projected control arm. The rearward force on the lower will be greater than the ground plane force, and the sum of the forces on the upper and lower (which will be subtractive from each other) will equal the ground plane force.

But if the brake is inboard, there will be rearward forces at both the upper and lower side view projected control arms. The torque of the brake will not act on the upright and the control arms. It will react directly through the caliper and rotor mounts on the sprung structure. Only the retardation force will act through the upright, and it can be thought of as acting on the upright at hub height.

The upper and lower forces will each be less than the ground plane force (and additive to each other). Their sum will still equal the ground plane force.