Commodore's Garage
Article #1 - Welcome to My Shop

Six years ago I began writing a setup guide. There are tons of these things that exist around the internet for any racing sim you can imagine. In the past, they were quite good for what they were to be used for, but when iRacing showed up in 2008 with a radically new approach to the garage, these guides started to show some issues. Most racing simulation garages are static, and changing one option doesn't directly affect anything else on the car. When I pulled up the garage for my Rookie Legend car way back in the first weeks of iRacing's public release, I realized we were dealing with a whole new animal. The car could *move* in the garage, and one change could affect every other thing on the car.

Guides made for NR2003 could be packed full of information, but could lead you down a road that would get you really messed up in the iRacing garage. In August of 2010 I was a sophomore in college, studying Mechanical Engineering here in Charlotte, North Carolina. I had also just finished my first season of real-world racing experience as a technical inspector and track official for US Legend Cars International. After sim-racing for five years, I had finally gotten my hands dirty working on real cars, and realized that something was missing from the then-current state of sim-racing. There needed to be a setup guide that explained not what each component did, but what it was, how it worked, and how it affects other parts of the car.

[GFR&D Car]

After two months, I posted it on the iRacing forums in September of 2010. Since then, iRacing has updated the garages quite a lot, to the point where my old guide doesn't really hold up anymore. I've had countless messages asking when I was going to rewrite it for the new garages, new cars, and new physics, and I've never had an answer. A couple of years ago, I spoke to Cale Gale, the "Gale" in Gale Force Racing, and talked about rewriting the guide after the Gen 6 NASCAR cars were released. The first outline would have produced no less than 67 pages, so we scrapped it. There's so much information, and so many minute details that make race cars work, that it's difficult to *not* write a novel when it comes to this stuff.

Since writing the first guide, I've left US Legends to focus on furthering my career. The five years in Legends Racing taught me more than I ever expected to learn at little short tracks with little cars, and looking back through the old guide makes me laugh at how much detail is lacking from that guide. There are things we know today about chassis setup that would have seemed alien at the time, and I'm sure there will be things in 2020 that would make us ask "...what?" That is how motorsport works, on all levels. Even in things like Legend cars, which you expect to be cheap and easy to be competitive with, can be engineering way beyond what the designers expected. I'll never forget standing on pit road at Charlotte Motor Speedway and watching a car enter turn 3 on the quarter mile and seeing the front springs bind perfectly and realizing the scale of engineering in motorsports. The same kind of evolution occurs in sim-racing, whether by software updates or someone with a wild imagination.

Mr. David Phillips gave me the green light to finally update what has been sitting dormant for six years. Through a series of articles here on iRacing's Sim Racing News, I'll go over the entire garage just like I did with the first guide, as well as stuff that isn't in the garage page, but will have a direct impact on chassis setup in the sim. While I'll focus mainly on stock cars and the tin-tops, I hope to also touch on things that are present in the more exotic racing cars.

[pic of something]

I will not sugarcoat this, and I believe it's necessary to say this at the start: I *will not* "dumb down" anything in any of these articles. We're dealing with a simulation of real-real world physics, not simply a video game. While there are tiny areas of the simulation physics that can be broken beyond the real world limit, and some that can't be pushed to the real world limit, it is an extremely accurate representation of real world physics. Like real world physics, almost nothing is cut-and-dry in the simulation. Some things work *this way* under one set of circumstances while the same thing may do something different in other circumstances.

I will rarely, if ever, say a component will *do this* if you *adjust this*. That's not how vehicle dynamics works, and that's exactly what we're dealing with. While things like setup matrices are a great help for someone who is learning the effects of an adjustment, they're becoming more and more obsolete as simulations become more and more advanced. Few adjustments in the real world are 100% consistent on what will happen, and the iRacing sim reflects that. Of course I'll explain what happens after you make certain changes to a car, but understanding what that adjustment is actually doing, and how it's doing what it does, is key to really understanding how to set up a race car.

[truck]

I'll structure all this in the best way I can think of, in the most efficient way to help everyone learn what is going on. Not every article will be about an adjustment, and instead will cover the reason behind the adjustment (bumpstops, for instance) or something that's going on that needs to be understood (downforce vs. sideforce). The first few articles will cover basic ideas about what you'll be looking to accomplish, what's happening with the race car, and the simple ideas behind everything. Before we talk about springs, we must understand what the Center of Mass is. Before shocks, we need to look at jacking forces. Before the track bar, we have to know what a roll center is and how it affects the car.

This stuff is confusing, and for many people seems like too much to figure out. "I need an engineering degree to set these things up!" is something I see a lot on the forums, and it couldn't be farther from the truth. You'd need an engineering degree to design the car, yes, but building the car and diagnosing issues with it just requires time, patience, and a little bit of effort. I'm always willing to help anyone who wants to learn, and the best learners are the ones with questions. So if you read something and don't understand it, feel free to send me a message, I'll help out as best I can.

There is *a lot* to go over, and I don't know how many articles it will actually take to get from one end of the cars to the other. The outline we wrote for 67 pages was a while ago, and the physics and garage have changed a lot from that time. I'll cover all the adjustments, but if there's a principle or specific method that needs to be addressed, we'll go over that too. Hopefully those who are looking to learn this stuff, and really learn it, will get a lot out of this series. I can't wait to get started, so check back next week for the first of many articles. Class is in session, let's get to it!

Commodore's Garage
Article #2 - Step 1: Grab a Pencil
by Matt Holden
June 28th, 2016

"We learned how to make them loose, now we need to learn how to make them fast."

Mark Martin said this during a post-race interview after an Xfinity Series race a few years ago and it's stuck with me ever since. It's a simple quote that largely went un-noticed, but carries so much information that it probably should be framed and hung on the wall of every race shop that is ever built. I don't remember whether he was driving for JR Motorsports or Turner Motorsports at the time, but I've never seen the quote in a frame at either shop so someone dropped the ball.

All joking aside, the first thing we need to address in this series is what direction to go in. We need to know what we're trying to achieve, and figure out how we're going to go about getting to that goal. Mark's statement following his race showed that their team knew what they were trying to do and they were going about it methodically. Being able to learn the cars and how they react to different things is key to working through various problems as they come along, and that means having goals and working through them one by one until you get where you're going. I wrote two articles last year to touch on this subject, which can be found here (http://www.iracing.com/behind-the-scenes-the-nascar-peak-antifreeze-series-powered-by-iracing-com-for-the-alpha-the-omega-part-1/) and here (http://www.iracing.com/behind-the-scenes-the-nascar-peak-antifreeze-series-powered-by-iracing-com-for-the-alpha-the-omega-part-2/). I recommend going over those before moving on!

[mark's nwide car]

Luckily, our path in the sim-racing cars is shorter, and less twisty, than what Martin was alluding to a few years ago. While many of their issues likely stemmed from chassis construction, body shaping, or manufacturing processes, we don't have to worry about any of that. Our cars are already built, they're essentially stuck in how they're going to be for at least three months at a time, and we can overlook some issues because everyone will have the same issue. With those out of the way, we need to look at the one thing that will affect our chassis setup the most: The driver.

One of the biggest changes in iRacing's physics over the past year or so is a shift from a one-size-fits-all setup to a more tailored setup for a given driver. When our team took part in the 2013 NASCAR PEAK Antifreeze Series, we could have had 4 or 5 cars on the exact same setup, and they'd all do fine. Today that's not the case. We put four cars on the same setup for Watkins Glen in 2015, one driver won the race and the other three did not do so well. One driver disliked the car so much that he gave up, one pushed through to finish in the back of the field, and one says (very confidently) that he'll never have a car that bad again.

Determining what you need in the car can be a difficult process. The old ways of saying, "I like a loose/tight car" is useless today, as well as a signal that a driver has never had a car that really worked well for him or her. If a driver says the car is loose or tight, the driver doesn't like the car. I helped a friend out a few months ago who was having issues getting the Super Late Model working, and when we finally got something together that worked for him, he said, "It's not tight or loose, it just drives well. It's a strange feeling." That's typically the response from most drivers when they get what they need in the car.

[Indy Class B]

When we're looking for what a driver needs, we typically look at two things: Driver inputs and the tire data. The driver inputs are fairly simple, and can be checked through telemetry. Typically, we'll take a "control" driver, one we know fairly well, and have him run 5-10 laps on a specific setup, usually the Fixed setup for whatever track we're testing at. When I'm doing something like this, I will always do the test at Charlotte Motor Speedway simply because it's not incredibly bumpy, eliminating the possibility for the track itself to dictate what a driver does. Kentucky, for instance, is incredibly bad for finding driver-related issues, since the large bumps can toss a car around and produce reaction inputs from the driver. Once the control driver has run, we'll record the 12 temperatures from his tires and find an average (wear numbers can be ignored). Then we'll get the other driver to make the same run in the same setup and record the temperatures in the same way. Both drivers will record a telemetry log file, and that will be analyzed together overlaid on top of each other.

When the data is compared, we're really just looking for major differences in how the driver is controlling the car. Both drivers are likely going to turn the wheel to the same maximum angle, both drivers will hopefully go to 100% throttle around the same point in the turn. We're looking for how they get to those points, though. Does one driver turn the wheel faster than the other, does one slam the throttle down while the other slowly applies it? What about the brakes, is one driver pushing the pedal hard for a short time while the other drags the brakes for a longer period? All of that will change how the car travels through the corner, so we need to know that ahead of time when we get into the fine-tuning adjustments.

The tire data is a good indicator of how the car is responding to the driver's inputs, and is the major indicator of what the car is doing. All we really need to look for with the tire data is a difference in where the hottest tires are. Nick Ottinger almost always has a cooler right-front tire than I do, while I have a cooler right-rear tire than he does. That fact alone means that we need a different weight distribution in our own cars.

[telemetry]

Once you have an idea of how you drive through the corners, the next step is to determine how stiff you want the car. Some drivers, like myself, prefer a car that is extremely stiff, with as little vertical travel as possible and as little roll as possible. Other drivers, like Nick, tend to respond better to softer cars that move around and shift as they go in and out of a corner. Dale Earnhardt, Jr. has mentioned multiple times this year that his car was "too rigid" on one corner of the car during a race. At Las Vegas, he said the right-rear spring was too stiff, while he said the right-front bumpstop was too hard at Charlotte. Unsurprisingly, his pace picked up once the crew was able to soften those corners to fit his driving style better.

The differences in a setup's stiffness has its own pros and cons, mainly centered around the car's aerodynamic dependence. Typically, a stiffer car will be able to produce more downforce for longer, but will suffer when the air is taken away in traffic while a softer car won't be as aerodynamically superior, but may navigate traffic better. For road courses, stiffer cars generally do not take curbs as well as a softer car, so if a driver likes to jump off of big curbs, he'll need a softer car underneath him to keep everything in control.

In terms of the inputs, a more aggressive driver will need stiffer everything, while a smoother driver can get away with softer shocks and springs. Stabbing the brakes and yanking the steering wheel can get the car moving around quickly, so that driver would need stiffer front springs and shocks to keep the car from slamming into the pavement. That said, the driver that does that will get the nose of the car down quicker, and generate a lot more downforce than a driver who slowly applies the brakes.

[pic]

This all seems tedious, I know. Most people want to get in the garage and start clicking buttons, but they rarely wind up with anything like they want. Taking the time to figure out what you're looking for will provide a huge boost to help you get what you're looking for with your car's setup. If you don't know what you want, how do you expect to know what to do? It's not a matter of "My car won't turn, we'll do this until it does" anymore, because that can lead to some problems that you'll never fix. You have to know what you're looking for and you have to know how to go about addressing issues as you go.

It should also be said that you should *never* be afraid to take what you're doing and throw it away. Sometimes ideas simply don't work and no amount of tinkering will fix it. When iRacing shifted the Class A series from the Gen-5 COT to the new Gen 6 cars, I had a series of setups I labeled as "PC*x*", basically meaning "Project Car" and then a number denoting which version it was. What those setups did was help to determine what had changed from one build to the next, and helped us to understand aerodynamic package or tire changes. Every time I threw out an old setup and rebuilt it, I increased the number for "x". I started that in late 2012 and the current test bed is PC8, so in that time I've tossed out 7 other iterations. It's simply part of the process, and it's a very frustrating part of the process. The key is to not forget what wasn't working and make the changes to correct those issues.

Finally, it's extremely important to not adjust the car for things the driver needs to fix themselves. When I was a kid, I raced R/C Cars at a local hobby store that had their own tracks. I was constantly changing out things on my cars trying to find a way to make up the time I was losing during the races. It rarely worked, and I'd always get extremely frustrated as to why things that these other guys were doing were not working for me. One day, this older man who I'd raced with for years saw me fussing with my car and after I'd explained that nothing I did made me any faster, he took me over to the board that had all our lap times on it. He pointed out that my fastest lap was almost a second and a half quicker than my average lap, which is a massive discrepancy on a lap that was only about 25 seconds long in a 5-minute race. He pointed out that my problem was not with the car, it was me. He said, "If your fastest lap is consistently close to your average lap, your car is holding you back. If your fastest lap is a lot faster than your average lap, you're holding the car back." As a 12-year old, I wanted him to be completely wrong because how a 12-year-old is, but I took the advice to heart because it was the only thing I hadn't tried yet. Over the next few weeks, I didn't adjust the car at all, I simply worked on running the same laps over and over. Eventually, the gap between my fastest lap of the race and my average lap fell to only a few tenths and I started winning races despite only having enough money in the budget for a can of Coca-Cola each week from the track's vending machine. It's a very simple thing to understand: If you're not driving the car consistently, it's not going to respond consistently, and if it's not responding consistently, you can't make good adjustments.

In the next few articles, we'll go over the main principles that cover how a car gets around the track. Aerodynamics, Roll centers, and I'll even show off what little I know of tire dynamics. Take the time until then to put all I've said here to the test. Figure out how you drive the car and see if you can find someone who drives completely different than you do. Once you understand what you need to do, we can start working on what you need in the cars. Good luck!

Commodore's Garage
Article #3 - Centers of Mass
by Matt Holden
August 4th, 2016

By now, hopefully you've got a good idea of what you're looking for in the car and how you compare to other drivers in terms of driving styles. If you're still searching for what you need that's no big deal, it can take some time to find. We've still got a couple of weeks before diving deep into the setup process, so no need to worry. I've set up a new way to keep track of all these articles, and the info for that will be at the end of this article.

One of the first things we need to understand with chassis setup is one of the most basic principles of physics: the Center of Mass. Everything that exists in the world has a Center of Mass, from a tall skyscraper to rain droplets. By definition, it is a single, intangible point in some object where mass is evenly distributed in every direction from that point. Most people know this point by a different name, the "Center of Gravity", but calling it by that name in an Engineering or Physics sense will get you at least two sets of dagger-eyes from other people. I'll go into why it is commonly called the Center of Gravity later.

The easiest (and probably one of the most well-known) demonstrations of Center of Mass is the "balancing bird" toy. It's very simple, and can be found in almost every gift shop, but by simply placing your finger underneath the bird's beak, the toy will balance perfectly on a single point, with the rest unsupported. This toy works as it does because the bird's Center of Mass is directly above the point in the bird's beak, allowing the bird to balance perfectly on your finger. The forward-swept wings place an equal amount of weight in front of the beak as well as behind, and since it's symmetrical, it balances! Theoretically, you could do this with a race car, but I do not condone such actions, and if you drop a car on your head because you were trying to balance it on your finger, it's your own fault.

[bird picture]

Since the Center of Mass has an equal amount of mass in all directions, it's simple enough (and a good learning tool) to find its location using the iRacing garage's corner weights. For practice, we'll use the corner weights from my 2012 Class B car at Indianapolis:

Left Front: 931 lbs
Right Front: 957 lbs
Left Rear: 962 lbs
Right Rear: 804 lbs

We'll use the dimensions given in the car's info page, which would have been the Nationwide Series 2011 Impala: a wheelbase of 110 inches and a width of 76.75 inches. First, we need the car's weight by adding the four corners, which equals 3,654lbs. Then, we'll find the left-side and front-end weight bias by adding the left-side weights and the front-end weights, respectively. The left side weighs 1,893lbs and the front end is 1,888lbs, and we need percentages, so divide both of those by the car's weight (1888/3654 = Front %). The left side is 51.8% and the front-end is 51.7%, which checks out in the garage, so we did well. The crude and easy way to find the CoM location is to multiply that percentage by the dimensions, which would give us the *opposite* distance to the CoM location. Yes, that is confusing, so I'll work through it. The front-to-rear location would be 51.7% (0.517) multiplied by the length (110"), giving us 56.87". The CoM would be located 56.87" forward of the rear of the car, or 53.13" behind the front of the car. We'd do the same for the side-to-side, using the 76.75" dimension, and get 39.756" to the left of the car's right-side door. So by doing just that, we've approximately located the car's Center of Mass in the car as about 1.87" ahead of the car's center and 1.38" to the left of centerline.

Vertical CoM location is nearly impossible through the sim. To find it in the real world, it involved lifting one end of the car while it's on a set of scales and measuring how much the weight on the other end increases, which can be plugged into an equation to find the height. For our purposes though, a 2-dimensional location is fine since we can only move it forward and rearward through ballast adjustments.

[Beam-pogo]

I'm sure there are some asking why that was helpful at all, and it's a very valid question to ask! The Center of Mass, in terms of the kind of analysis we're doing now, is the acting point of all forces acting internally on the car. Lateral forces, acceleration forces, and yes, even *gravity* can be considered to be acting at this single point on the car. This idea that all forces act on a single point has led the Center of Mass to be commonly called "Center of Gravity", or "CG", because drawing out the forces acting on something will always result in Gravity being drawn at the Center of Mass. Both terms are equally safe to use, and both mean the same thing. The picture to the right shows that moving the car's Center of Gravity still results in forces acting at the same point and the same angle, no matter where that point is located in the car.

[CG Move pic]

The car's CG location is important when trying to understand how the car is reacting to various adjustments and inputs when on track. As we saw earlier with my Indianapolis car, the CG was located slightly off-center and towards the left-front wheel. That was the older Class B Impala, and because the CG was located so close to the center, the forces generated from Indy's low banking angles wouldn't pull the left-side of the car down very easily. For that race, I ran a large rear spring split to make up for that, which would help the left-front to drop more easily when the car was loaded into the banking.

[Indy pic]

Conversely, if we compare the more offset chassis of the Super Late Model, the CG is much farther to the left. The higher left-side weight bias from the chassis offset means it will turn a whole lot easier to the left than my Class B car would, and banking should have a much more pronounced effect on the car's tendency to roll to the left as banking increases. Because of that, we can run a narrower rear spring split if desired to help drive off the corner, since the banking is going to do most of the work for us.

[superlate pic]

Finally, we're going to look at Inertia, defined in Isaac Newton's first law of motion. Paraphrasing, the law says that, *A body in motion will remain in motion unless acted upon by an external force*. Earlier I said that all forces can be considered as acting upon the car's Center of Gravity, and like normal G-forces, the same can be applied to the car's inertia, or its tendency to continue moving in a straight line.

When we try to turn a car from traveling in a straight line, we begin fighting the car's natural tendency to continue in that straight line. For a stock car, we're fighting over 3,500lbs of stuff that does not want to do what we want it to do. Steering, aerodynamics, and track banking all introduce forces that help convince the car to move in a different direction, but at the basic level the car is completely happy going straight ahead. This law of inertia is the reason why we see cars continue in a perfectly straight line when something happens to a car. If a tire goes down or suspension pieces break, the path the car will take becomes almost perfectly straight until it interacts with something else, usually a wall, in a direct demonstration of Newton's First Law.

The car's front-to-back CG location has a huge role in how well a car will navigate the turns on a racetrack, and we have a direct control over where the CG is in this regard: **Ballast**. The Ballast adjustment is one of the simplest adjustments we have in the garage. I'll go into more detail with it in a later article, but its purpose is directly related to what we're discussing. In most cases, a race car will be built underneath the minimum weight required by a racing series' rules. To reach the minimum weight, teams will add ballast to the car which can be moved around in the car. Being able to move this weight around is important to making a car handle well, so a lot of effort is made into making the cars as light as possible so that more moveable weight can be installed on the car.

In a stock car, ballast is placed within the chassis frame rails at the bottom of the doors. A large square tube is used for the outermost section of the chassis' center area, and teams will use either blocks of Lead or Tungsten (depending on what the series allows) as ballast. Moving the ballast forward or rearward in these rails will adjust where the CG is located in the car, and directly affecting how the car handles at high speeds. Moving the ballast forward in the car will shift the CG more forward, creating more understeer and a "tight" condition, while moving it rearward will do the opposite.

[Kart pic]

To finish this section up, I'll leave you with my own personal experiences with ballast in a racing vehicle. Before a few years ago, I always assumed that there was a speed where the ballast began doing the opposite of what it would do at high speed. At 100mph, more front ballast would make the car tighter, and at 20 or 30mph, it would make it looser since there was more weight on the front tires. When I built my go-kart in 2014, I tried putting weight up in the front of the kart, even drilling holes to mount blocks between my feet to help the front tires dig into the pavement. During my first outing on-track, the kart wouldn't turn at all, plowing straight through a very low-speed hairpin and proving to be *extremely* exciting on the high-speed sections of track. Keeping in mind that the hairpin was probably 15-20mph and the highest speed turn was around 50-60mph, that's a big swing in speed.

I made two runs to make sure it was consistent, and then went back to my pit area to fix it. I removed the blocks from the pedal area at the front of the kart and mounted them to the side of the seat, right beside my hips. In total, I shifted about 20lbs about three feet back on the kart. The result was astounding, with the kart suddenly being *too* loose in the high speed sections (about as exciting as the first runs) while being perfectly tame and predictable in the low-speed hairpin. I've since adjusted the ballast slightly forward from where it was but have not placed any more weight in the pedal area of the kart. This example was very eye-opening to me, since it took what I thought to be true and broke it completely. Today, I use the ballast adjustment constantly to fine-tune the car's handling when we get to a new track, and it's an invaluable adjustment.

Commodore's Garage
Article #4 - Aerodynamics 101
by Matt Holden
August 17th, 2016

Aerodynamics: the one word that spawns happiness in engineers and dread in race fans. In the past 50-or-so years, aerodynamics in motorsports have done nearly complete 180° turn, going from slippery, super-low-drag vehicles to high-downforce behemoths with the sole purpose of bullying air molecules into doing work that they really don't want to do. Want to experience this yourself in the iRacing sim? Go run a few laps at any track in the Lotus 49 and then switch to the Lotus 79. You may notice that the 49 will not navigate a turn much faster than your bicycle is capable of, doesn't believe in "braking", and will immediately be somewhere else when you hit the gas. The 79 will do the opposite, and that is the power of downforce. To understand how it affects racing vehicles, we need to start with downforce's opposite: Lift. For that, we need to look to the sky.

The Wing

My father is a recreational pilot, something he started doing long before I ever came along. Sailplanes (also called "Gliders") are his wheelhouse, and growing up on the grass runways of the Carolinas bit me with an aerodynamics bug very, very early in life. "Why does this thing stay in the air?" When NASCAR began seeing the benefits of aerodynamics in the late 90s and early 2000s, this question changed into "How is it stuck to the track?"

The airfoil shape is the basis of lift-based aerodynamics. If anything is producing a force to keep it aloft or press it into the ground, there's an airfoil shape on it somewhere. The basic shape is pictured below, but modern wings have evolved into more radical shapes with curves, thickness changes, and tiny "flicks" at the trailing edge. In simplest terms, wings operate on a phenomenon known as Bernoulli's Principle: *As a fluid's velocity increases, pressure decreases*. The curved surface on the top of an airplane wing causes the air to speed up, thus reducing pressure. The air on the flat side doesn't change speed much at all, and thus maintains a constant pressure. This pressure difference exerts a force on the wing in the direction of the lower pressure, and moves the wing.

[wing shape]

For racing cars, we'd prefer them to stay on the ground, so we flip the wing over. Now, the flat side is facing up, the curved side is down, and the high pressure sits above the wing and presses down. This is downforce. Road racing vehicles have gotten incredibly efficient at doing this, and can efficiently make tons and tons of downforce at relatively moderate speed. Stock Cars don't do this very efficiently at all, but still make it work.

The major paradox is that a stock car has an airfoil shape, yet still produces a net downforce to keep it on the ground. Because of this, they have a tendency to depart the racing surface if spun around at high speed, known as a "blowover". Despite no externally visual aerodynamic devices aside from a tiny splitter at the front and a near-vertical spoiler at the back, they still make a lot of downforce, almost exclusively from the car's underbody.

The Splitter and Radiator Pan

Based on aero numbers, stock cars make the entirety of their downforce from underbody effects. Essentially, the downforce is generated from air going *beneath* the car instead of over it. In most oval series, cars use what's known as a Valence or Air Dam at the front, essentially a strip of plastic at the bottom of the nose that scoops air up over the front of the car. NASCAR's top-3 series use a splitter, a flat piece of composite material that "splits" the air stream and sends some up over the car's nose and some underneath the car. It's not a new technology, it's been used for years in sports car racing and even open-wheel racing.

NASCAR mandates a flat-bottom splitter, which is bad in terms of aerodynamic shape. Teams will often angle the splitter very slightly so the lead edge is lower than the rear of the splitter, which creates a "diffuser" effect. Essentially, the space behind where the air finds the splitter is a larger volume than the air that enters, and produces a low pressure. Due to the angle of the splitter, that area gets larger as the air goes farther back, thus reducing pressure even more.

The same happens at the radiator pan, essentially a flat piece of metal sitting underneath the radiator. It has a built-in angle like the splitter, and operates in the same way. NASCAR usually messes with the size of the radiator pan to change the front downforce numbers because they're very easy to produce and swap out. Below is an image from StockCar Engineering Magazine that shows airflow underneath the splitter through a CFD simulation.

[cfd splitter]

While the color legend isn't explained for obvious reasons, it's easy enough to see what is happening. As the air approaches the splitter from the left, it's compressed underneath the splitter's leading edge and then expands quickly as it passes under the splitter and radiator pan (large rectangle behind the splitter) before hitting the sway bar tube and then going everywhere as it goes into the engine bay. Despite the initial compression, this expansion under the splitter and radiator pan produces a very low pressure, basically a vacuum, which we call downforce.

Due to how the splitter operates, it is necessary to have air going underneath it. Without air underneath it, it produces a full vacuum, which pulls the splitter into the track, which unloads the front tires, and turns the car into a sled. Losing the air underneath is known as "stalling" the splitter, and it can end your day very quickly. In addition, the lower the splitter is, the more downforce it will produce, and NASCAR's splitter is almost linear in downforce numbers as it travels down towards the racing surface. As a result, it's desirable to have the splitter as close to the track as possible *without* it touching, which prompted the return of bumpstops and extremely stiff front springs when the COT debuted in 2007.

Fender Wells

The second major point of downforce production is the fender openings. The air that went under the splitter has to go somewhere, and having it fill up the engine bay will produce a high pressure under the hood and lift the car up...which is also detrimental to performance. By working air around the fender openings properly, teams can create a low-pressure zone just outside of the wheels, pulling air out from underneath the hood area, and causing the engine bay to have a low-pressure zone underneath the hood. Just like the splitter, this produces downforce. The same can happen at the rear fenders of the car, but usually to a lesser extent. This is why we saw teams pulling out the right-rear fender and bending the right-side skirts up in 2015.

Decklid and Spoiler

While the spoiler has basically been reassigned to sideforce production (later article), the shape of the rear of the car still produces downforce. Not all of the air that goes into the front of the car will escape through the wheel wells, so some will come out of the back. Similar to how sports cars and old Formula 1 cars utilized skirts to "seal" the sides of the car, NASCAR teams have done the same. Skirts are essentially walls that will channel whatever air is left to the back of the car, where the angled trunk pan can produce a diffuser effect similar to the radiator pan. If you remember back in the 2011 and 2012 seasons, NASCAR made attempts to cut the skirt length by a few inches to reduce rear downforce and allow air under the car to "bleed" out of the side. This was almost a complete failure, since almost everyone (except NASCAR, apparently) knew that softening the rear springs would fix that issue basically overnight.

The spoiler also causes air to be trapped on the top of the trunk area, known as the decklid. This air getting stopped in place produces a high pressure, which works with the low pressure under the trunk to produce downforce.

The Greenhouse Effect

The major problem on a stock car is the roof, aka "The Greenhouse". It has a curved shape which, as we know from airplane wings, produces lift. Below is another image from StockCar Engineering magazine of the wind tunnel numbers from the 2014 Sprint Cup aerodynamic package.

[CFD Numberz]

In the engineering world, it's not proper to use two terms to describe the same force, and "Lift" is used for all aerodynamic forces acting up or down. Positive Lift acts upward, negative lift acts downward, we know that as Downforce. From these numbers, we can see that the downforce produced by the car's underbody exceeds the car's weight at a staggering 3,817 pounds of downforce. That would be enough to do that "drive on the roof" thing, but that gets killed off by the top of the car, primarily the greenhouse, which produces over 1,400 pounds of lift. All told, the two numbers resulted in just under 2,400 pounds of downforce in 2014, which has since been reduced tremendously as of 2016.

Interestingly, NASCAR's R&D released information in early 2016 following the High-Drag package tests at Indianapolis and Michigan that showed cars trailing another car produced *more* downforce due to the roof being unable to produce the lift it does in clean air. This basically confirmed that downforce was not the problem for trailing cars, and is the reason why we see a lot of effort going into reducing sideforce this season.

The Aero Package

With all that laid out, you should have a basic understanding of how stock car aerodynamics works. We essentially want the car as low to the ground as possible without anything touching the track, and we want it to stay that way for as long as possible. Since we cannot do much to alter the car's aerodynamic devices, most of the cars will have a very similar lift-to-drag ratio, and in most cases the effects of drag can be ignored in terms of chassis setup, unlike something like Formula 1 where a car with more drag needs to have that taken into consideration. In the mid-90s, NASCAR teams wanted as little drag as possible, and relied primarily on mechanical grip, which resulted in cars that rolled a lot, had a large gap between the chassis and the track, and weren't as fast as today's cars. In fact, of the 23 tracks on the NASCAR Cup schedule today, only three tracks have qualifying records that were *not* set with a Generation 6 car: Daytona, Talladega, and Atlanta. Daytona and Talladega's records obviously stand due to restrictor plates, and Atlanta's record stands because it was repaved in 1997. The record there is a lap of 197.4mph, while the closest anyone got this year was Kyle Busch at 193.3mph. It's safe to say that the 4mph would be made up easily by new pavement.

Looking through the past few years, it's easy enough to see that teams are still going for the same result in their aerodynamic platform each year. Despite rule changes to remove ride height minimums, the cars still look the same on-track. I'll end this article with a picture of Jimmie Johnson's cars from 2007, 2011, and 2016. All three cars, despite being three different generations of NASCAR's top vehicle, all have the same aerodynamic platform. As we start talking more about springs, shocks, and chassis setup, this is what we're trying to accomplish, and this is where your raw speed and handling will be found.

[Jimmie]

Commodore's Garage
Article #5 - Roll Stiffness
by Matt Holden
September 1st, 2016

When we watch a car go around a corner, we almost always take a note of how much it rolls over, or lifts the inside of the car, but how much do we really think about what's happening? In the yester-years of NASCAR, it was probably never given a second thought to how much a car was rolling over in the corners, but here in 2016, it's become a science. There's one more section we need to cover before diving into the garage, a very complex but necessary concept that is key to understanding what the car is doing while on-track: Roll Stiffness.

Roll Stiffness is something that gets covered extremely early in any motorsport-related schooling. When I was in college, I had a Formula SAE car on scales performing a roll-stiffness test within the first month of motorsports classes. On the surface, it's very simple: How much does the car's body roll for an applied force? I learned this characteristic in "Pounds per degree", or lb/°. So if the car had an overall stiffness of 250lb/°, it would take 250 pounds of lateral force to roll the car over 1 degree from horizontal.

[old jeff pic]

On that Formula SAE car, which was specifically built for 60-second autocross runs, we could get away with doing a test of the entire car's roll stiffness. For a racing vehicle that needs to cover 300 miles, we need to break it down further to "balance" the car. While the entire car will have a given roll stiffness, that value is derived from the roll stiffness values of the two suspension systems in the car. For most racing cars today, that's two independent suspension systems. However in stock cars, we have an independent, double A-arm suspension in the front, and a Solid Axle, truck-arm suspension in the rear. Two *extremely* different systems that need to be looked at separately.

Roll Centers and Moment Arms

We'll start with looking at how the suspension rolls in the first place. On every suspension system, we have two points that play into this: The Center of Gravity and the Roll Center (or Moment Center, to some). Center of Gravity was covered earlier, so go brush up on that if you need to. The Roll Center is an imaginary point somewhere in the suspension which can be thought of as the point the suspension rotates around as forces are applied. For an independent suspension system, locating this Roll Center is very complex and usually requires computer software, or an intern, to locate it without burning up a few hours in the day.

Finding the Roll Center's location involves extending lines through each of the suspension arms, locating their intersection points, then drawing more lines from those points to the center of the tire contact patch. Got it? Below is an image of how the Roll Center would be located in an independent suspension system like what is on most race vehicles. Each of the suspension arms has lines extended from them (blue for right arms, green for left), and each of those pairs intersects at a point, known as the "Instant Center". We then draw lines from the Instant Centers to the centers of the tire contact patches (red and pink lines). Where these two lines intersect is where the roll center for that suspension system is!

[RC pic, independent]

Ready for the rear suspension? This is where it gets complicated: The Roll Center in a solid axle suspension, like on NASCAR stock cars, is right around the center point in the track bar. That's it, seriously. I've got another picture below of the rear suspension assembly on an old COT (it's a show car, but still works for this explanation. I've marked the center of the track bar with a green "X", and that's literally all there is to it.

[Rear RC]

We've got all that, our brain has melted, now what? To visualize how suspension systems deal with lateral forces, take another look at the rear suspension picture above. You'll notice that there's also a white "X" up in the car, the Center of Gravity, and those lines are connected with a pink dotted-line. Those two points (Roll Center, CG) and the pink line create what's called a "Moment Arm", named as such because it generates a Moment, or torque, on the suspension. In the article about the Center of Gravity, I mentioned that all forces act on the car at the Center of Gravity, and this comes into play again here with the suspension! The best way to think of the Moment Arm is as a wrench: Place the bolt you're trying to turn at the Roll Center, and your hand is the Center of Gravity. When you push on the free end of the wrench, you rotate the bolt. The same phenomenon occurs in suspension systems.

When the car goes into a corner, it experiences a lateral force pushing outward from the inside turn, known as Centrifugal Force. For analysis purposes, this force acts on the suspension systems at the CG, and the Roll Center will basically remain fixed in the car, rolling the suspension to the outside of the turn. In race shops, while preparing for a race weekend, a lot of work goes into moving various suspension components to determine the optimum roll center location for each suspension system, and getting this right before a team loads up for a track is often the reason why a car wins or loses on race-day.

Balancing Systems

Suspension stiffness is directly adjustable by moving the suspension's Roll Center (RC). If we raise the RC, we shorten the moment arm, and the suspension becomes stiffer in roll. Similarly, if the moment arm lengthens by lowering the RC, the suspension becomes softer in roll. While we have almost no control over the front roll center location in iRacing's Stock cars (it's directly adjustable in some cars, such as the V8 Supercars), we have some small control over the moment arm length in the front. On the rear, however, we can directly adjust the RC's location.

Ideally, we want the suspension systems to work in harmony and not fight one another over a long race. In a sprint race, we can get away with a chassis imbalance since tire life isn't a big deal, but over a long race it becomes crucial. When one suspension system is much softer than the other in roll, one tire will become overloaded while it's buddy-tire will become under loaded. For instance, if we have a very soft rear suspension in roll but a very stiff suspension in the front, the right-front tire can become overloaded while the left-front can become under-loaded. Think of a dirt late model and how they used to run on three tires everywhere for a very exaggerated example. We don't see that anymore though, mainly because teams realized the left-front tire wasn't simply a balancing ballast and began using it to help the car turn. When we have the suspension systems working in harmony with all four tires doing their job equally, we have a dynamically balanced race car.

[DLM pic]

There have been numerous methods to assist engineers in designing cars with balanced suspension, the most well known being the Roll Couple Distribution method. This method was used to calculate the load that would transfer from the inside to the outside of the car in a turn, and attempted to mathematically allow engineers to match the front to the rear of the car and wind up with a balanced race car. How it works, I have no clue, because by the time I started learning all of this stuff there was a new method: equal roll angle. This method attempts to match the suspension by taking the known lateral force for a corner, and building the suspension systems so they'll roll to the same angle when that load is applied. Doing so should produce a balanced car, and it's been used at least at the lower levels of stock car racing.

Adjusting RC and Moment Arms

We'll definitely revisit these things later in the series, but there are ways you can adjust the moment arms and roll center heights in the iRacing garage that can help you get a feel for what is going on. First, and the easiest, is overall track bar height. This is extremely simple, and is typically what I do to the car in a "final practice" type session. Keeping the height split the same on both sides, just raise or lower the *entire* track bar equally. So if you raise the right-side 2", raise the left-side 2". This raises the RC height 2 inches and shortens the Moment Arm (not necessarily 2" though). A shorter moment arm means a stiffer suspension system in roll. A higher rear roll stiffness will make the car looser, or more prone to oversteer. The opposite will happen by lowering the track bar. Raising the bar will cause less weight to shift off of the left-front while turning, and can help your car maintain speed over a longer run, too.

The front suspension is less adjustable, however. If you are running any of the stock car series in iRacing, you can alter the front moment arm slightly by the front end's height. A higher front end will *usually* cause the car to be much looser due to a softer roll stiffness. Gale Force Racing actually ran into that exact issue in the 2014 NASCAR Pro Series when we got to Dover. Some of the drivers raised the cars to allow the car to drop from the high-load turns, but it wound up making the car extremely loose. Since then, I've suggested that drivers find a ride height setting they like in the front of the car and make sure to keep the heights at those settings in all cases. If you start changing springs, shocks, or whatever else, but don't keep the heights the same, you run the risk of changing the moment arm length. That's two changes, and you probably only intended to make one!

Springs

Our next article will start covering springs. Springs are directly tied to the suspension's roll stiffness in the same way the Roll Center and Moment Arm are, but to different effects. In some cases, coil springs don't handle roll very well and contribute very little to the roll stiffness. In others, such as anti-roll bars (or sway bars), they contribute *a lot* to the stiffness. And then you can add in bumpstops, or sway bar asymmetry, and you have a ton of things to think about.

We've covered all the basic principles necessary to start looking at what we can do in the garage. Next we'll look at springs, from 5-inch Big Spring cars to Coil-over Late Models and Modifieds, to the mysterious bumpstop and bump springs, and even to sway bars, also known as "torsion" springs. I'm looking forward to it, and I'm sure all of you are as well!

Commodore's Garage
Article #6 - The Spring
by Matt Holden
September 9th, 2016

In the history of sim-racing, nothing has been more associated with chassis setup than the springs. At the same time, nothing has been more misunderstood than the simple coil spring. There are countless different kinds of springs, from coil, to leaf, to torsion springs, and everything in between, and racing series regulate them heavily. When I worked at US Legend Cars as a technical inspector, there was no part of the car more strictly defined than the springs themselves. They had to be a specific length, made of a specific material, possess a constant rate, and maintain a constant diameter (No barrel springs!).

There is even an entire industry dedicated to spring design, and creating the most consistent rates and performance possible with today's technology. Luckily, we don't have to worry about things like the spring losing rate over time, or bending, or anything like that in sim-racing, but it's still important to know how these components work, what they're used for, and how they can affect your performance on the virtual race track.

Spring Rate

To understand spring rate, we must understand Hooke's Law, which governs the physics behind almost every spring in existence. Hooke's law simply states that the distance a spring is compressed (its *Deflection, x*) is directly related to the force the spring is exerting, *F.* The equation is completed by adding in a spring constant, or "rate", represented by *k*. So for all springs, $F = kx$. It's that simple.

A spring's rate is a representation of how much force it will exert for a given amount of compression, or deflection. The rate is expressed in a force-per-distance unit, most often as "pound-per-inch" or "Newton-per-(milli)meter". So if we have a 500 lb/in spring, it will exert 500 pounds for each inch it is compressed. If we compress it one inch, it exerts 500 pounds. Two inches, and we'll see 1000 pounds, 1500 pounds for three inches, and so forth. In basic terms, the rate will be an indication of how stiff a spring is, or how resistant it will be when forces are applied to it. A 50 lb/in spring will compress fairly easily under a 25 lb load, while a 1000 lb/in spring will not compress enough for us to be able to notice.

Springs with a constant rate through compression are known as "linear" rate springs. Most real-world springs do not have exact rates and may change slightly through travel, but it's not enough to cause the spring to not be classified as linear. If the rate of the spring changes through travel, this is known as a "progressive" spring. In many racing series, a progressive spring is illegal, however the vast majority of composite and rubber bump stops are progressive by nature. Below is a graph that would look similar to the output of a spring rating machine:

[rate graph]

This graph shows the force vs. travel characteristics of a given spring. I intentionally altered how much force was added per inch of travel to produce a trace that wasn't perfectly linear, but close enough to be applicable to what we're talking about. The rate of the spring is the value of the chart's rate-of-change throughout travel, and is determined by the coefficient of *x* in the trend line equation shown in the graph. So based on what we see here, the spring I made shows a rate of 248 lb/in and would be labeled as a 250 lb/in spring. The springs we have in iRacing are perfectly linear and are the rate we see in the garage. It's an impossibility in the real world, but it eliminates a few headaches in the name of simplicity.

Springs in Race Cars

If you've driven a car, you've ridden on springs. That's a fact, and we don't have the widespread technology to have anyone in existence who can't say that. Springs are a key component in car suspension systems, and are shouldered with a lot of responsibility, from dealing with forces applied from bumps, to controlling body pitch and roll, and in higher performance cars, even handling aerodynamic loads at high speeds.

When I was a kid back in the 90s, I remember seeing NASCAR Cup cars heaving up and down and swaying over to the side as the drivers slung them around the big speedways. I specifically remember watching Jimmie Johnson's #48 and Bobby Labonte's #18 jump up and down over the backstretch bumps at Charlotte when they were fighting for the win in the 2005 Coca-Cola 600. Today, that sight is a little more rare, and usually only found on local short tracks, but why is that?

When aerodynamics started to become a problem to consider in the mid 2000s, team engineers realized how they could get a car around the track faster by harnessing the air going around the car and generating grip instead of relying solely on the suspension and tires to do the job. The result was faster cars, and a dramatic change in how the cars looked on track. Even in world-class series such as Formula 1 and WEC's LMP1 class, the cars are extremely rigid, with little to no suspension movement. I always enjoy pointing out how a Formula 1 car's suspension will seem to be completely solid at Monza, while the suspension arms move a relatively large amount at Monaco. It's simply to use aerodynamics as efficiently as possible, and the simple fact that if we set a car up like Johnson's 2005 Coke 600 winner, it may not even make the show, let alone win the race.

I had a professor in college tell me a very useful thing: "Set the springs for speed, the roll stiffness for cornering." We discussed roll stiffness last week, but we need to cover the main reason springs are used today: *SPEED.* As I hinted at with the Monza/Monaco example, springs in today's race cars are used to control aerodynamics and the car's attitude at speed, not primarily for mechanical grip. If we have a car that will be traveling at high speeds and we need to control the aerodynamics, we need stiffer springs. If we can sacrifice aero for mechanical grip, then we can get away with softer springs. We'll go more in-depth with this in another article.

Spring Types

While I mentioned a few types of springs earlier, we really only have to worry about three types of springs in the sim-racing world: Coil, Leaf, and Torsion.

Coil springs are the most recognized, and are exactly as their name sounds: a coil of wire. In a coil spring, we have a central axis around which the wire is wound, usually at a constant diameter. If the diameter changes on the length of the spring, it will be named either a “Barrel” spring (narrower at either end) or a “Conical” spring (cone-shaped), but they all work in the same way. Applying pressure to the ends of the spring causes the wire to flex up and down, which it would rather not do, and thus exerts a force to counteract the applied force. Pretty simple!

Torsion springs are found on most cars in the form of a “Sway Bar”, or Anti-Roll Bar. Torsion springs are simply bars of metal that are twisted through a lever, and thus exert a twisting force to counteract that force. Some racing vehicles, such as Formula 1 and LMP1 cars, may use torsion springs on the corners as the wheel springs, twisting the spring through a complex system of pushrods and rocker arms. In those systems, one end is fixed to the chassis and the free end is twisted.

Finally, Leaf springs are a rather antiquated type of spring, but we still see it in some racing vehicles such as the iRacing Street Stock. Leaf springs are made by stacking strips of metal and then bending the stacked set, creating a curved shape. Applying a force to the top of the curve results in the stack acting as a spring. These are seen very often on modern pickup trucks and jeeps, but they seem to be getting replaced slowly by coil springs.

Spring Motion Ratios

Something we don’t have to consider with iRacing is what’s known as a spring’s installation ratio, or “Motion Ratio”. Whenever a spring, shock, or sway bar arm is not mounted directly on the centerline of its respective wheel, it doesn’t move at the same rate as the wheel. For a simple example, if we have a 12” suspension arm, and the spring’s center axis is located 6” inboard of the wheel’s centerline, it will probably move at ½ the rate of the wheel. If the wheel moves up 1”, the spring will move 0.5”, and we have a 2:1 motion ratio. This can obviously get dramatically more complex by introducing spring angles, but for a very simple explanation, we'll keep everything vertical and parallel.

While it's not something we deal with directly in sim-racing, it's a common question: "Why does setup A work in this car but not in this one?" For instance, why can't we take a setup out of the Class B Xfinity car and drop it straight into the Gen 6 Cup car? More likely than not, the two cars have slightly different suspension systems, and the springs are mounted with different motion ratios, resulting in a different behavior on-track. There's also the case of different aerodynamics and weight distribution, but it's something to consider. Also, it's the main reason why you can't always take the setup from your real-world Super Late Model and put it in the iRacing Super Late Model. If you don't know how the motion ratios differ, it's going to be hard, or nearly impossible, to make it work properly.

Multiple Springs per Suspension System

NASCAR uses an extremely old suspension system in its race vehicles, opting for a solid rear axle suspension system instead of the more common independent systems we see on the road today. The solid axle suspension system is interesting because the two corner springs both affect how the other corner works, specifically in the overall vertical stiffness of the rear end. In an independent rear suspension system, the left-rear spring won't have much to say in how the right-rear suspension travels, but that's not the case here. Matt Kenseth's flip at Talladega this year is going to be extremely helpful in a lot of things, this time we'll look at where the springs are located in the rear suspension (orange arrows):

[kenseth flip]

The springs are located on the far end of the truck arms (painted white). Despite being located on two different components, they both influence the same suspension system. The two rear springs in a stock car are mounted in parallel, and thus combine rates to get an overall rear stiffness. To find this, we simply add the two rates together, so if we have a 250 lb/in left-rear spring and a 1000 lb/in right-rear spring, the rear suspension has a 1250 lb/in compression stiffness.

The front suspension acts in a similar fashion, however much more complex. If we're using bump stops, we can wind up with three different springs acting on the suspension in the simplest example: Main spring, bump stop, and sway bar. These are also in parallel, so they can add to find the wheel's spring rate. This is something to consider as well, since the suspension's spring rate changes the moment you've contacted the bumpstop. Here's an image of a modern stock suspension system with the various spring highlighted:

[suspension]

Real systems get even further complicated by suspension arm flex (which we have in iRacing, but it's not really a concern), tire spring rate, and the spring rate of anything else that flexes while the suspension is in motion.

The alternative is when springs are mounted in series, or in-line with one another, usually on the same shock. This is common in dirt racing, especially Rally and dirt oval racing where a softer spring will be used next to a stiffer spring, often called a "tender" or "helper" spring. For this situation, the combined rate is much more complex, found by multiplying the two spring rates, then dividing that by the sum of the rates. For these arrangements, the softer spring is usually expected to bind quickly, meaning the suspension is now solely on the stiffer spring rate.

Commodore's Garage
Article #7 - Selecting Springs

by Matt Holden
September 9th, 2016

Last week we looked at how springs work, how they're rated, and the different types of springs that may be available on our virtual race cars. The big question, however, is always "What rates do I use?" Older sims had flaws in the physics, so spring choice was simply based on what you could exploit to get as much speed as possible. Recent builds for iRacing have brought spring choice closer to reality, and choosing springs based on track characteristics is becoming more and more important. Things such as aerodynamics, banking, and grip level play their part in laying out what a good setup will be, as well as what can get you into trouble.

The Car

Obviously, the biggest factor in choosing springs is going to be the car itself. Lighter cars can use softer springs, while heavy cars need beefier springs to keep them from slamming into the race track. In five years of working with Legend cars at Charlotte Motor Speedway and Concord Speedway, I never saw a spring that was labeled over 400 lb/in. The car itself weighs only about 1300 pounds, so that was more than enough! On the other hand, I've seen NASCAR trucks in recent years that had all four springs above 1000 lb/in!

Every car will have a spring range where it's happy. Too soft and it may flop around as you try to navigate even the slightest bend, but too stiff and the tires may never be effectively in contact with the racing surface. Finding this "happy" range can be a drawn-out task, but narrowing down what you're going to work with can really simplify the process down the road. Some cars, like the Mercedes GT3, are limited heavily on spring choices, so it's not that difficult to decide what you want, and it's much harder to work yourself into a massive problem.

The Driver

Drivers like to complain, and when something isn't how they like it, they complain *a lot*. Alan Jones famously left Formula One in the early 80s because the suspension systems were becoming too stiff and making it very hard on the driver following a ban on ground-effect skirts. Drivers today are no different, and will not hesitate to let you know about it. One of my teammates, Alex Scribner hates the feel of high-rate rear springs in his NASCAR vehicles, and will jump at the opportunity to swap them out for a softer spring. I, on the other hand, prefer a stiffer suspension.

In my second article of the *Commodore's Garage*, I mentioned how important it is to know what you're looking for in a car so you don't waste countless hours using things that just don't work. Now's the time to put all that to good use! If you like a more rigid suspension under your car, go for some higher rates. Spring rates alone do not make a setup fast, so don't be afraid to go against what the "norm" is if you simply do not like it.

Aerodynamics

Remember last week when I mentioned the difference between suspension movement in a Formula 1 car at different tracks? We're going to revisit that now, and I'll again use the extremes to look at it: Monza and Monaco. Teams regularly bring their minimum-drag package to the Italian Grand Prix in Monza, but want a maximum downforce package for Monaco. These changes in aerodynamic packages are not limited solely to the aerodynamic bits and add-ons to the car, they incorporate the suspension settings as well. With low-drag comes low-downforce at Monza, and teams not only want to keep some amount of downforce in the car, but also want to keep the car sitting with an attitude that will produce the lowest drag possible. Since mechanical grip is almost a non-issue, this calls for very stiff suspension, and almost no vertical suspension movement can be seen during a lap of the track.

Monaco, on the other hand, is relatively slow, run on public streets, and thus isn't a great platform to produce downforce with. The low speeds (Lewis Hamilton went through the Casino Hairpin at 45kph, or 28mph in his 2015 pole lap) of the track require the car to generate a lot of mechanical grip when the downforce falls away, so the cars run softer suspension than at higher speed tracks.

This is a big thing to consider when choosing your springs in any car. You're obviously not going to stuff some 1500lb/in springs in the Street Stock because it's just going to bounce away, and at the same time you're not likely to mount 50lb/in springs in the Williams because the downforce would crush them.

Furthermore, think about how much aerodynamics is going to be a factor at a track. Aerodynamics are important at *any* track, but some tracks, like the example above, will benefit more from a bit of aerodynamic sacrifice and a little focus on mechanical grip. Martinsville is commonly thought to be a track where aero does not matter, but teams still show up with the front of the car low in the corners trying to get whatever advantage they can. If a car is moving through air at any speed, it's being influenced by the air.

Track Loading

The next thing to consider is how much loading the track is going to produce. For road racing, this is not a massive deal save for a few places such as Zandvoort, but it's *huge* for oval racing. For example, Dover and Atlanta, despite their differences in everything, produce very similar vertical loading. For a general rule, higher banking will require stiffer springs regardless of the track's length. Bristol is only 0.5-mile in length, but its high banking angle produces very high vertical loads. New Hampshire is one mile long, just like Dover, but vertical load is almost non-existent because of the low 7° banking in the corners.

It's also important to consider *how quickly* the car will "load up" in the banking of the corner. I mentioned Dover and Atlanta being similar, but Dover tends to produce the load much quicker. Dover's corner entry is almost like a drop, while Atlanta is smoother. You can control a lot of this with shocks, but Dover may require slightly stiffer springs to keep the chassis from slamming into the pavement.

The only effective way to determine a track's loading is to record telemetry for a track and look at the vertical G-loading for the car in the corner. It can also be done with some simple math, but it's necessary to know the exact banking of the track, the car's speed in the corner, and the turn radius, which is difficult to find sometimes. The vertical G value can be multiplied by the car's weight, producing a load on the car. This can be compared to other track figures, and help you get a starting point for your next race.

Banking Characteristics

A big thing to consider is almost always an oval-track problem: banking. There are some banked corners in road racing, but suspension asymmetry is not usually considered based on those corners like it is in oval racing. How the banking forms in relation to the straight is a big factor in how the car itself loads up in the corner. If we think about SMI's Atlanta, Charlotte, and Texas Motor Speedway, they're all the same length, they all have about the same banking (~24°) in the corners, but they're all unique in how their banking is built.

Banking is formed in two ways from the most basic standpoint: The top stays at the same height and the bottom drops away, or the bottom stays and the top rises. Rarely does a track follow either of those exactly, but each track can be related in that sense. Texas and Atlanta both have banking where the top rises more than the bottom falls, while Charlotte's Turn 1 entry has the bottom drop away in an extremely similar fashion to Dover's turns. Tracks like Atlanta or Texas, where the banking forms more on the outside of the corner, will load the right-front corner more heavily on entry, and will typically call for a stiffer right-front spring in relation to the left-front spring. Similarly, these types of turns tend to have a very flat corner exit (think of Texas' Turn 2), and may require a little extra left-rear rate to keep the back end of the car planted leaving the corner.

A general rule of thumb is to increase the right-front spring rate as banking increases. Whether you do that through the spring itself, or by other means such as sway bar asymmetry, is up to your own preferences, but this can help the tires last longer into a run. Conversely, flatter tracks, such as Martinsville or New Hampshire, may allow you to run more even front spring rates, or even a stiffer left-front in some cases. This is also hugely based on driver preference, since some drivers simply do not like having a stiffer left-front spring.

Bumpiness, Grip

Finally, we're going to look at how bumpy and/or worn out the track is. A bumpy track will usually require softer suspension and higher ride heights, while a smoother track can call for stiffer springs and much lower ride heights. Similarly, older, more worn-out tracks will want softer springs, high-grip tracks will allow the use of stiffer springs.

We'd all like to have a lot of downforce at every track, but bumpy or worn-out tracks will often benefit less from aerodynamics than they will gain from mechanical grip. If we install high-rate springs into a car at a bumpy track, there's a good chance that the car can bounce around too much, the driver loses confidence in the car's grip, and lap times suffer. However, if we sacrifice some of the aerodynamics and install some softer springs, the car can glide over the bumpy sections of track and the driver has a more predictable platform underneath the car. The same can go for a worn-out surface. Chicagoland isn't extremely bumpy, but the pavement is relatively old and has lost some of its grip. We may want to install some softer springs for a little extra mechanical grip to keep the car from skidding when the tires can't grip the track as well.

New pavement is always fun, and it's not uncommon to see the biggest springs show up on newly-paved tracks. These situations have almost no bumps and a ton of grip, meaning we can sacrifice mechanical grip for the sake of aerodynamics. Extremely high-rate springs, aggressive shocks, and low ride heights can be used effectively, combining the best of aerodynamics and mechanical grip to produce some extremely high speeds.

Go forth!

I've intentionally kept spring rate numbers vague because I want everyone to go try to find what is comfortable for them. If I were to say, "I used a 900 lb/in right-rear spring at Kentucky", you may not be inclined to experiment with what would work best for you, and that's not what this series is for. I've given some guidelines as to what to think about, and when those are combined they can narrow down your choices considerably. Also, keep in mind that a set of springs doesn't make a car fast by itself, there are many ways to go about this. Never be afraid to try something new, and if it works, stick with it! Try not to do the "setup dance" at a new track, but instead find something that is working well for you and adjust it to the new track. Over time, with minor adjustments each week, you'll consistently run well each and every week while you build up a notebook on what the car wants in certain conditions and what things don't work at all.

Commodore's Garage
Article #8 - Coil Binding

by Matt Holden
September 23rd, 2016

Last week I went over the various thing that can influence your spring choices for a given track. How much banking, amount of bumps, and even the weight of the car can drastically alter what you need in your race car, but can those rules be broken? Of course they can, and it's not uncommon for them to be broken for the sake of a stopwatch. NASCAR in the 90s and early 2000s was purely a fight for mechanical grip, but that all changed sometime between 2003 and 2005 (the title of "First" is debated) when a rather ingenious crew chief found a way around the minimum ride height rules and was able to harness the air like nobody before him. This started a trend that still exists today, we know it as "coil binding".

The Basics

In a coil spring, we have a limited amount of available travel, or the total deflection the spring can undergo before the coils themselves will start touching each other. Once the coils touch, the spring is "binding" (the height of the spring in this condition is known as "block height"). Regardless of the spring's rate prior to binding, whether it be a 300 lb/in spring or a 3000 lb/in spring, the effective spring rate once bound can be thought of as *Infinite* for our purposes. In reality, further compression would depend on the spring material itself because any more travel would come from literally squishing the metal together, kind of like you might smash Play-Doh in your hands. That considered, racing springs are usually made of steel, and very tough steel at that, so the chances of the spring compressing any further past its block height are almost nonexistent, so it's safe to say that a bound spring has an Infinite spring rate and can be treated as completely solid.

WHY?

In most engineering circles, if you design something with a sprung mass, but the spring isn't strong enough to support it, you get made fun of. During a design process, a lot of work goes into simply isolating and dampening vibrations, usually by choosing springs that will prevent harmful oscillations from destroying the product. This goes for automotive suspension as well, with car manufacturers putting a lot of time into matching spring rates and damping rates to produce good handling and good ride quality. So with all this considered, the question becomes: "Why on earth would you bind a spring?"

The genesis of coil binding is both secretive and openly debated. Some say that Michael Waltrip's car in 2005 at Charlotte was the first to use it properly, others say it was going on a few years earlier when the common-template bodies showed up in 2003. I've also never gotten a straight answer as to whether or not Jeff Gordon's 1997 T-Rex car was coil-binding or not, which would explain why it was so blisteringly fast compared to the rest of the field. Regardless of when it started, it stems from a need to get around ride height rules, and nothing really beyond that.

Ride height rules were circumvented as early as the 1980s in Formula 1's Ground Effect war, when teams would run soft springs to keep the car at legal heights while at a standstill, but utilized bumpstops to keep the car off the ground at high speeds. NASCAR teams wanted something similar after bumpstops were banned in the early 2000s. The mandated frame height rules were *way* too high to be of any value to aerodynamicists, who want the cars as low as possible. Instead of stiff springs to keep the car up off the track, they opted instead for softer springs (usually just strong enough to maintain legal ride heights for tech) but allowed the springs to bind just before the car hit the track. Thus, they clear ride height checks in tech inspection as well as have a really solid...no pun intended...aerodynamic platform on the track.

Developing the Coil-Bind Setup

With the introduction of the COT, NASCAR allowed bumpstops back into the picture but kept the minimum ride height rules. In an attempt to save money by offering teams an alternative to coil-binding, it inevitably wound up costing a bunch of money in the long run. Early setups in the COT looked similar to stuff from the 90s, and that first race at Bristol was known for high body roll angles and giant splitter gaps. Eventually, the soft spring idea returned, and the bumpstops were used as the ride springs at speed. While not an infinite spring rate, bumpstops were still very stiff because they want aerodynamics and don't particularly care about the driver's teeth while on track. Stiffer bumpstops wound up producing sharp changes in wheel rate for small amounts of travel, and then we wind up with the "pigtail" spring.

The last major evolution in coil-binding still exists in the Truck series (and possibly Xfinity) series today, and is known as the "pigtail" spring. NASCAR mandates that front springs have a "closed, ground coil end" and one "open coil end" (pictured to the right). In a standard configuration, the open end (top of the picture) would sit in the lower suspension arm and the closed end meets the perch. Teams began running angled spring buckets where the only part touching the lower mount was the open end of the spring. This caused the spring to bend under aerodynamic load with little effort and effectively bound only half of the spring. This allowed much stiffer rates to be run (upwards of 2000lb/in) but still got around the ride height rules. The high rate spring allowed better control of the bumpstops, and produced a more consistent handling. A picture of a spring binding in this configuration is below, with a link to a demonstration video in the description.

[Pigtail]

With the removal of ride heights in 2014, the practice of coil-binding has vanished from the Sprint Cup garage since it's now possible to run whatever rate springs you want, as well as get the aerodynamic platform for the best downforce.

Coil-binding in the iRacing Sim

iRacing has a very accurate coil-binding system built into the physics, however a lot of the major issues have been stripped away for the sake of the user. The major thing to remember is that the iRacing springs will bind across the entire spring equally at the same time, so there's no pigtailing going on. This means that once you've bound a spring, you have no suspension left, so timing and rate choice are extremely important. As of writing, all of the fendered stock cars can run a coil-binding front end, but the Sprint Cup, Xfinity, and Super Late Model offer bumpstop options instead. There's no right or wrong way to set the cars up, but some drivers prefer the bumpstops over coil-binding, so play around and see what you like better.

If you do decide to go the coil-binding route, it's important to know the springs you're working with. Across the stable of big-spring coil-binding cars (*not* coil-over like the two Late Models), from the National car to the Sprint Cup, it's safe to assume that any spring under 430 lb/in will bind under all circumstances. Springs from 440-470 will bind in specific cases (like if you raise the nose and have a high vertical load), but anything above that will likely not bind and stay down for the straights. It's very simple to run tests to find out which springs can provide the travel you need, simply by starting with a baseline and changing the springs out one by one to see how early or how late each one binds. Run a few laps on a pair of springs, change one (and reset the ride heights!!!), and go back out and see how much difference there was in ride height. A general rule of thumb is that a softer spring will raise the nose of the car if you are binding them, while a stiffer spring will lower it. For instance, if you have a 350 lb/in left-front spring and that side is hitting the track, swapping to a 330 lb/in should bind earlier, preventing the chassis from hitting the track. This is backwards from the traditional school of thought, not a typo!

Below is a picture from a 2016 test where I charted the spring heights for the Class B car. Both runs pictured were taken using the same starting ride height, but changing the left-front spring. Both springs are less than 425 lb/in, but there is a tremendous difference in splitter height.

[xfinity pic]

Adjusting with Coil-binding

Once you've chosen your front springs to give you the proper ride heights and splitter clearance, you're undoubtedly going to have to adjust things to get the car handling properly. There are a few things to remember though:

- *Avoid Sway Bar changes if possible:* While balance changes with the sway bar are effective when you're using freely-moving springs, sway bar changes while binding can mess up your aerodynamic platform by allowing too much body roll with a small change. If the car does happen to roll up off of the left-front spring with the right-front bound, it can get extremely tight. A natural thought would be to decrease the sway bar diameter, which would make that issue much worse.
- *Remember that your spring rates no longer work:* Like I mentioned, a bound spring is effectively at an "infinite" spring rate, and is no longer the rate represented in the garage. Prior to binding, the rate would be in effect, but once bound it's thrown out. Because of this, front spring rate changes won't actually result in an expected result because both the old spring and the new spring will still be "infinite" lb/in.
- *Keep the same front springs week-to-week:* Changing the front springs requires you to start completely over with setting the heights, sway bar, and alignment. iRacing's chassis don't change much at all for a given track type, so it's wasteful to change front springs when going to a different track. What worked at Charlotte will work at Homestead, Dover, Texas, and Kansas, while your Richmond package should work at Martinsville or Iowa as well. All of your major tuning adjustments can be done with the rear end without hurting the aero platform at the front of the car.

Finally, crossweight adjustments gain a new level of complexity when coil-binding as well. If crossweight was adjusted properly (same ride heights, no preload changes in front ARB and rear housing), it will change how much travel is available in the front springs, which is where close attention to spring deflection comes into play. It's important to keep track of how much travel is available in a spring before and after a crossweight adjustment and to adjust accordingly. If you aren't sure how to find spring travel, take a look at the spring deflection numbers. The first number is how much it's compressed in the garage, the second number is how far it can travel before binding. The available travel is the difference in these numbers, so if you have a spring deflection "4.55 of 6.10", you have only 1.55" of available travel in that spring! Changes in this value will result in the car's height when the springs are bound, so pay attention closely.

Coil-binding is complex and the main reason why there was such a discrepancy between higher-funded teams and low-budget teams in the higher levels of NASCAR. Luckily, we don't have to go buy special springs or perches for our virtual racing cars because they've done all the hard work for us already. Still, it's a very detail-oriented aspect of racing even in the virtual world, and will reward those who pay attention to the most minute details. Doing so will save you a lot of headaches, keep you from chasing handling issues that you may never find, and keep you running up front.

Commodore's Garage
Article #9 - Bump Springs

by Matt Holden
September 23rd, 2016

So far we've gone over how the main springs work on the car, but we still need to cover secondary springs, or "bump stops". iRacing has recently overhauled the Gen 6 Cup cars and the Xfinity cars to use bump *springs*, but those still fall under the category of bump *stops*. To understand why the Gen 6 car was changed over to bump springs, we need to know the reason behind bump stops in the first place. As you can imagine, there's a lot to go over.

Bump Stops

Bump stops have traditionally been used in motorsport to literally stop "bump" travel in suspension systems. Bump travel specifically refers to compression of the suspension...usually over a bump...or from aerodynamic loading, braking, and even acceleration. One of the earliest uses I know of is Williams F1 using bump stops to get around Formula 1's 6 cm ride height rules, which were intended to reduce downforce levels from the evolving ground effect phenomenon. The cars ran softer main springs but eventually compressed into bump stops at full compression to keep the car off the ground under the high aerodynamic loading.

NASCAR teams did something similar when the COT was introduced in 2007. The bump stops that were allowed with the new car were used to prevent the car from hitting the track, while the springs used in the cars were huge by the day's standard, a far cry from the "baby buggy" soft springs used in previous years. Larry McReynolds said during the spring Bristol practice in 2007 that the springs in the cars that were on track were of rates he hadn't seen since the 90s. The cars "sucked" to drive, probably attributed to the horrible aerodynamic platform they had.

Bump Stop Rate

At some point (my guess would be lap 2 of that Bristol race) someone realized the rubber bump stops had their own spring rate, and they began using *that* as the at-speed spring rate for the front end of the car. If we fast-forward to 2012, bump stops had been engineered to be much more consistent, specifically in rate progression through travel. When coupled with the pig-tail spring assemblies, this could be used to keep the splitters sealed and the car handling relatively consistently through a race.

The issue lies in the behavior of spring rate in a rubber bump stop. While coil springs are engineered to produce a consistent rate through as much of the full travel range as possible, rubber material basically cannot do that if used in a solid shape. Instead of a constant rate, a rubber bumpstop will gain rate through compression, known as a progressive-rate spring. While a coil spring may produce forces of 250 lb at 1", 500 lb at 2", and 750 lb at 3" (a 250 lb/in spring), a bumpstop could produce 100 lb of force at 0.1" of compression, 250 lb at 0.2", but 700 lb of force at 0.3" of compression. An example graph is shown below. When compared to the conventional coil spring, a rubber bump stop could potentially increase the suspension's overall rate by a *lot* in a very short amount of travel. In simplified terms when the above example is considered, we're working with 1000 lb/in of rate at 0.1" of compression, but we're going to be over 2300 lb/in at only 0.3" of compression! This sharp increase in rate would be immediately apparent to the driver, and could produce unpredictable handling swings if not controlled properly.

[bumpstop rate charts]

In the real world, rubber bump stops can be further complicated by the shape of the bump stop itself, the spacers in between the bump stops, and even how many make up the entire stack itself. Remember that springs in series wind up producing a lower rate than when used individually, and combine that with the exponential gain in force these can produce, and you can run into some major headaches! This led to the need for stiffer springs (for rate control) without super-stiff shocks, but there was still a need to circumvent the ride height rules, so we wind up with incredibly expensive pig-tail springs and bump stop-specific shocks. Below is a composite picture (credit: RE Suspension) of three types of bump stops that were common in oval-track racing. These are all shown with the shock compressed to the point of contact, but take note of the different style of washers in the middle, "puck"-style, bumpstop assembly. Each has a different angle on the cone-like face, which would change the way the bump stop generated force under compression!

[RESusp bumpstops]

The Bump Spring

[eibach spring] Enter: The Bump Spring. These likely started out on the weekly late-model scene, but the rubber bump stop eventually got unseated by what's now known as a "bump spring". The bump spring is very simple to understand, and works exactly like a typical main spring would, however it's located on the base of the shock shaft where the bump stop would be located. Using a coil spring instead of a rubber bump stop effectively eliminates the exponential force increase present in the rubber bump stops, meaning the handling can be more consistent across unexpected conditions like bumps, aerodynamic issues (trailing behind another car) and even banking changes at a track like Homestead or Iowa. As you can see from the Eibach spring pictured to the left, it looks just like a normal spring, so there's no witchcraft or confusing magic going on with how it works.

The mechanics of its operation are also quite simple, and mirror the bump stop in its operation. Whenever the bump spring contacts the shock body, this point is set specifically with "packers" (more below), the bump spring rate is then included in the overall spring rate of the suspension. If you think back a few weeks, I described how multiple springs in a given system will combine for a new rate. In this case again, we have a parallel spring system once the bump spring is active, so the bump spring rate adds to the main spring rate to provide an "active" spring rate. For instance, if I have a 500 lb/in main spring, the suspension is acting on 500 lb/in until the bumpstop is in contact with the shock, which I set at 2500 lb/in for this example. Once the bumpstop is in contact, I now have a suspension system that is behaving with 3000 lb/in through travel. Of course, the *actual* rate would be lower than this due to motion ratios between the main spring and the shock, as well as ratios between the wheel and spring locations. This can be disregarded though, because all of these ratios are constant across all of our cars.

It's use, in iRacing's garage, differs based on the car it is being used on. Currently, only the Gen 6 "Cup" car and the "Xfinity" Class B cars have bump springs and while the Xfinity car basically presents only one type of use, the Cup car can utilize the bump spring in many different ways. The Cup car is basically infinite in its possibilities. You can choose soft main springs and stiffer bump springs, you can run stiffer main springs and soft bump springs. You can run one, you can run both, or you can run no bump springs, there's literally *any* configuration that you can make work. The main decision is based on the driver, and what they are looking for. I've worked with drivers who liked stiffer main springs with softer bump springs, but there were other drivers who *hated* the feel of the big main springs and refused to race that way. In 2015, when Gale Force Racing had almost seven cars running in the DWC-level PEAK Antifreeze Series, it was not a surprise to see seven different front end styles on track at once, and they all wound up doing the same thing.

The Xfinity Class B car, however, is fairly simple since it still runs with the minimum garage heights. For most tracks, you can run a soft main spring (careful to avoid binding if you go this route!) and then choose bump springs to ride on when you're up to speed. Since the bump spring will activate at a lower ride height (and cannot be preloaded), it's not uncommon to need a much stiffer bump spring than you'd normally choose as a main spring when starting at a high ride height. For instance, if we start the left-front corner at 5.5" (minimum height) and we'll hit the splitter on the track at 2.5" of height, we have 3" of height to take up. For the sake of examples, we'll say the motion ratio in the suspension is 1:1, so we have 3" of suspension travel to gather up to get the splitter down. We'll install a 1000 lb/in spring on that corner and compress it three inches to get the splitter down, so it's exerting 3000 lbs of force at this point. Now, let's install a bump spring and a softer main spring to drop the nose quicker, and this bumpstop engages with 1" left in the total travel available. I'll use a 400 lb/in main spring at this point, so at 3" of travel it's pushing 1200 lb of force. We're 1800 lbs short of our goal here, but we know that the bump spring will compress 1", so it's very simple: We'd need an 1800 lb/in bump spring to achieve this. We've gone from a spring rate of 1000 lb/in with a single spring to 2200 lb/in with the bump spring, but wind up with the same force at the end. It's not necessary to go through all these steps with bump springs, but it's intended as a way to explain why a stiffer set of springs is often seen when the bump spring is added!

[indy pictures]

Bumpstop Gap, Packers/Shims

Finally, we're going to look at what's used to tune the bumpstop contact point. The point where bump stops engage is crucial to how the car handles. In terms of driver feel, contacting a bump stop is basically the same as multiplying a given corner's spring rate in an instant. It's also possible to run into big swings in cross weight if the bump stops don't engage properly, which can lead to major issues at tracks where the car "lands" heavily in the banking, such as Dover or Charlotte. Our tuning tool for this is the *packer*, basically a small shim that sits on top of the bumpstop stack to manipulate when the bumpstop engages in the suspension. Below is another picture from RE Suspension of their metal-dish bellows bump stop (very common in late models and possibly even Cup) that has a similar spring rate behavior as the coil bump spring:

[bellows pic]

What we're looking at is the actual bumpstop stack at the bottom of the shock. Here we have a lot of metal washers shaped like cones that are stacked in opposing directions, re-creating a bellows-style spring. The large part on the outside (with a cutaway) is the bumpstop cup, used to align all the components and keep them all compressing in a consistent manner time after time. Inside the cup are the metal washers, all stacked at the bottom of the cup. On top of the washers is a red washer, used to ensure the bump stop engages consistently across the entire surface of the bump stop. The very left of the image has the bottom of the shock body. The distance between the shock body and the red washer is known as the "bump stop gap". The bumpstop will not be in play until the shock body contacts the red washer. We can add shims (which sit on top of the red washer) to reduce the bump stop gap, or remove them to increase the gap. In the image of the three bump stop styles, there are two packer shims installed at the top of the upper-most (white "Christmas tree" bumpstop) assembly. We'll look at how to set the packers in a later article, but for right now, take the time to consider why having the bump stops engage at different points would produce an ill-handling or unpredictable car.

The bump stop trend is confusing, I'll admit. There are a lot of things at play when they're installed on a car, and the days of slapping them onto a shock simply to keep the car from bottoming out are gone completely. The learning curve is very steep, mainly because everything about them goes against traditional thinking, but once they're understood it becomes fairly easy to manage what's going on. It's okay to get frustrated, that's a given, but when that happens take a step back and try to envision what all is going on and what you can do to fix it. Working with two springs per corner is not easy, but everyone can do it, it just takes a bit of patience!

Commodore's Garage
Article #10 - Track Bar

by Matt Holden
September 23rd, 2016

Before moving into initial setup options on our cars and how to go about getting what you need to race, we still need to cover the "bars" in the car. Stock cars typically have a front sway bar and a rear track bar (or Panhard Rod), while road racing cars have a front and rear sway bar. In this article we'll cover the track bar and the purpose it serves in the car, the sway bar will be covered in the next article. While I try to include a lot of images to show off how the components operate in the car, but the track bar is an incredibly difficult part to see on a car because of its location. Plus, teams often paint the bar black, making it even more difficult to see! Still, I have tried to provide images when possible to point out the bar in real-world cars.

Tracking Straight

In the style of suspension used in a stock car (truck arm, solid rear axle), the rear end assembly is designed to naturally move around a little bit in operation. The entire rear end housing (the part that contains the rear axles and differential) can move left to right in the car as well as up and down under normal operating conditions simply because of how it's built. It's simple enough to understand how having the rear tires wander left and right at 180mph would be a bit of a problem, so we have what's known as the "Track Bar". One end of the track bar is installed solidly to the left end of the rear end housing (just inside of the left-rear wheel) and the other end of the bar is attached to the frame behind the right-rear wheel, usually on an easily-adjustable slider mechanism. This installation method links the floating rear-end housing to the frame and body of the car, reducing or eliminating side-to-side movement of the housing in the car, and keeping the car driving in a straight line.

[t-bar image]

To look at the track bar's mounting system, we'll look back at the image I used to point out where the roll center is for the rear suspension. While I want to note that this is the track bar in a show car, it still shows off the general configuration that is common today in stock cars. We're looking at the back of the car, with the bar oriented in the way it would be oriented on a setup sheet in iRacing's garage. The left end of the track bar (height shown with blue bar) is bolted onto a slotted mounting plate attached to the rear end housing. Each slot on this mounting plate would represent one "adjustment" for this end of the track bar. The right end of the bar (shown with a red bar) is mounted to a sliding adjustment mechanism attached to the bottom of the car's frame. It's difficult to see in this image, but the bar itself is attached to a sleeve around a large tube on the right of the image. This sleeve has a screw running through it, and turning this screw (whether by the window bolt or an electronic driver adjuster) will raise or lower this end of the track bar.

Adjusting Overall Height

The track bar can be used primarily in one of two ways: overall height and height difference (also known as "rake"). There is quite a lot of general stuff that goes on with the track bar in terms of how the rear end behaves in the car, how the springs and tires are loaded, and how the car rolls left and right in the corners. All of those are covered by these two simple options though, making a very simple device into a very complex and important part of the car.

The overall height of the track bar directly influences the rear suspension's roll stiffness (covered in article #5, "Roll Stiffness") by locating the roll center of the rear suspension. For the type of suspension we see in a stock car, the roll center is located around the center of the track bar, regardless of its configuration in the car. So for instance, if we have the track bar set at 8" high on the left side and 10" high on the right side, the roll center is located about 9" off the ground (RC location: (*Left-side height + Right-side height)/2*). In terms of overall balance of the car, this is the easiest adjustment we can make to the car to cure long-run issues. As we raise the rear roll center, the rear suspension will become stiffer in roll and the left-front/right-rear will be loaded heavier than at a lower setting. Raising the bar in this manner will reduce a long-run swing towards a tight car by evening out loads around the car. Similarly, if we lower the track bar, we'll reduce a long-run swing towards a loose car.

Where this becomes a more complex problem is in driver feel. If the rear of the car is significantly softer in roll than the front of the car, raising the track bar will typically not be something the driver feels at all initially, but may comment on it later into a run. My own experience with this adjustment has shown that the real effects won't show up until lap 15-20 at the earliest when the driver may comment that raising the bar freed the car up through the center. On the other hand, lowering the track bar is almost always apparent to the driver immediately, with the driver typically saying that the rear is more secure, or that it has more "side bite".

Furthermore, it's important that this adjustment includes *zero* change in the track bar's rake. Moving either end of the bar will change where the roll center is located, but adjusting one end without the other will change how the rear end is loaded under lateral forces. If your goal is to change the car's handling over a long run, and you believe that a roll center change will alleviate some of your handling issues, make sure you move the track bar equally on both ends. If you don't, and you change the bar rake, you'll wind up making *two* adjustments.

Adjusting Rake Angle

Track bar angle adjustment is the most common adjustment we hear about in NASCAR. The left end of the bar is fixed in place once the car is on track, but the right end of the bar can be moved up or down while the car is on track or in the pits through the screw adjustment mentioned earlier. While this does adjust the roll center location (and thus long-run balance) it primarily changes how the bar is loading the rear end of the car.

When the car goes into the corner and has a lateral force applied to it from cornering, the track bar has to exert forces to keep the body and frame from moving over to the right. If the bar is perfectly flat and parallel with the track, it will exert a force in a purely leftward direction to keep the rear end housing from moving. If we run the bar with an angle then the bar will exert a force at an angle, relative to the chassis. This produces both a leftward force and a downward force on the right-rear of the car.

[bar angle image]

In this image I've drawn a simple track bar (black bar) mounted at a large angle for demonstration purposes. In the corners, the car will experience a large lateral force, represented by the purple arrow. This force will try to push the car and its parts to the right, and since the rear end housing is directly connected to the tires, which are stuck to the track, the body of the car will want to rotate to the right, away from the rear end housing. The track bar will take over and pull the body back in line, exerting a force down the bar from the right-side mount to the left-side mount, represented by the red arrow. Since this force acts at an angle, it can be split into a horizontal component and a vertical component. The horizontal component of this force, shown as a blue arrow, is what keeps the housing and body aligned under the lateral G-force generated by the turn. The angle also produces a vertical force, represented by a green arrow in the image above. This vertical force pulls down on the frame, and due to the right-side mount's location loads the right-rear spring and wheel. If we were to run the right-side lower, or "negative rake", this would pull *up* on the right-side of the car and unload the right-rear. Forces such as this vertical component are what's known as *jacking forces*, or forces that are placed on the tires as a result of suspension component movement, which we'll cover later.

Other Considerations

In real-world racing, the track bar's location can tell you a lot about the way things are set up in the rear of the car (and sometimes the front, indirectly). In the days of the Generation 5 "COT" car, we frequently saw the bar very low, an indication that the suspension had a naturally high roll stiffness. For the Gen 6 cars, we see it a little bit higher than back in the COT days. We can link this to the fact that the COT was allowed to use a rear sway bar, thus providing a simple and effective tuning option for the rear roll stiffness. Despite the COT having traditionally soft rear springs (usually less than 1000 lb/in in most cases), the track bar was set low on the car because the sway bar was able to handle the roll stiffness dilemma.

When teams started running the rear bar in a manner that would produce extra skew in the rear end, NASCAR removed the rear bar for the Gen 6 car and we now see higher rear spring rates (in some cases) as well as a higher track bar. Teams have also started painting the bars black, making it harder to see them in the shadows underneath the car. What works on a real car may not work for the sim-racing car simply because you don't know what else is under the real car. Below is an image of the #48's car up on jack stands in the garage. Even though the ends of the track bar are visible in the image (the bar itself is black), trying to build a setup based simply on seeing this for your sim-racing vehicle would be difficult, if not impossible. Instead of seeing something like this on TV and trying to replicate it in the sim, think of what may be going on with the springs to result in this track bar configuration, and then start there. Furthermore, those of you who pointed out that the car being on jack stands produced a higher angle than it would on the ground get a high-five.

Commodore's Garage
Article #11 - Sway Bar

by Matt Holden

October 20th, 2016

I had a professor in college say, "Springs for speed, bars for handling". It's a short, but interesting statement, and it's a concept that's easy to forget and go about things in the other way when tuning a race car, be it real or virtual. Modern race cars, oval cars especially, have seen huge advances in the technology that goes into something as simple as a sway bar, or anti roll bar. Just like the coil springs, sway bars can be used for both mechanical grip and aerodynamic grip and come in many different forms. However, they all perform exactly the same role regardless of what car (or type of car) they're installed in.

What is it?

A sway bar, in simplest terms, is a torsion bar spring linking two corners of a car's suspension to prevent opposing movements. If the wheel on one side of a car goes up and the other doesn't, or goes down, the bar exerts forces to bring the two back in line. The bar is typically round, frequently steel (special case for the Indy car at the end of the article), and part of the bar is fixed to the chassis in some way. Sometimes the bar is solid, sometimes it's hollow, but it all depends simply on how the car's engineers want the bar to behave.

[bar image]

This image, from *Speedway Engineering*, shows the components typically seen in a modern racing sway bar. While the sway bar in your passenger vehicle is likely one continuous tube bent around various components in the chassis, racing cars usually employ a "three-piece" sway bar, where the bar (left of picture) can be changed without changing the arms (right of picture). The sway bar arms can feature their own adjustments, such as offset or length, and can also be changed independently of the bar itself.

[formula bar]

This image shows the suspension arrangement typical of an open-wheel car, or any car utilizing push- or pull-rod suspension. In this specific arrangement, the anti-roll bar is situated in front of the suspension components with one end fixed to the chassis (not shown) and the other end clamped into a t-shaped arm. When the suspension moves and rotates the rockers, the linkages to the arm twist the bar, creating the anti-roll effect.

NASCAR-style Sway Bar

Following the great war over rear sway bars being "improperly" mounted during the 2012 season in an effort to generate extra rear end skew, NASCAR put a stop to rear sway bars in the Cup series when they introduced the Gen 6 cars. The Xfinity "COT" also arrived on the scene without rear sway bars following a vote among the owners in that series. Similarly, many oval series in the US do not allow the use of rear sway in stock car racing series, and some disallow sway bars altogether. This places a heavy emphasis on the front sway bar to do a lot of the work to "flatten" the front of the car for aerodynamic purposes, while placing the responsibility of rear roll on the rear springs themselves.

The front sway bar in a NASCAR-style stock car is mounted in front of the front suspension, just behind the radiator, in a tube welded to the bottom of the front clip (known as the "sway bar tube"). The bar is placed within this tube where it is allowed to rotate freely and steel arms are clamped onto the ends of the bar. Linkages are placed on the end of the sway bar arm which link it to the suspension's lower control arm.

[Hendrick]

I took this picture on a visit to Hendrick Motorsports a few years ago. This particular car is one of Jeff Gordon's test cars from the 2013 season, so there are some extra bits in the picture we wouldn't normally see, such as the laser sensors on the splitter and the blue suspension travel sensor behind the shock. Still, the sway bar is mounted in a fairly typical fashion, and can be seen just to the right of the brake assembly. The large, dark grey triangular-shaped piece is the sway bar arm, and the silver circle at the front end of the arm is the end of the sway bar.

[Chevyotodge]

Here is a picture I took in college of the Red Bull Racing car my school acquired following the team's closure. This is the right-front suspension with the main spring and bumpstops removed, as well as the sway bar disconnected. The arm can be seen pointing downward on the right, and we also have a clear view of the sway bar tube at the front of the chassis.

Sway Bar Diameter

The bars themselves are available in many different diameters, with this being the major factor in a sway bar's stiffness. Bars are "rated", similar to how a spring is given a rate, however sway bars (and torsion springs as well) are usually rated by giving them a force at a specific angle of twist. For oval track racing in the US, sway bars are typically rated by twisting the bar 5° from rest, and then given the force they exert at that angle of twist. If we have a bar rated at 1200lbs, then that specific bar would exert 1200 pounds of force at 5° of twist. Bar length is typically mandated by the sanctioning body, and is typically a non-factor in sway bar selection.

Sway Bar Arm Lengths

Whenever a torque (twisting force) is applied, the resulting force is largely dependent on the length of the arm applying the force. Here is the equation for torque:

$T = r*F$

T = Torque (lb-ft, N-m)
r = Radius, or arm length (inches, meters)
F = Force applied at the end of the arm (pounds, Newtons)

We can see from this that, for a given torque, if the arm length is increased, the force at the end of the arm must decrease to keep from changing the torque. This is what happens when you go and get a bigger wrench, or get a big pipe to put on the end of the wrench, for a bolt that won't budge. The longer arm allows you to place a weaker force on the end of the arm to get a large torque at the bolt. This same effect occurs in a sway bar, but typically on a smaller scale.

If we know that the bar will exert a given torque at a specific angle of twist, changing the arm length changes the forces acting on the suspension at the end of the arms. If we increase the length of the arms, we reduce the force acting at the suspension, effectively "softening" the sway bar. Conversely, if we shorten the arms, we will increase the force at the suspension, "stiffening" the bar.

[zoomed Hendrick pic]

Zooming in on the picture of the #24 test car from earlier, we can see that the sway bar arm installed on the car has three length options, with the linkage installed in the middle of the three available settings. The picture at the top of the article has similar holes. These holes allow for an easy way to fine-tune the sway bar's characteristics without having to remove components of the sway bar.

Sway Bar Asymmetry

Whenever we run the bar with different-length arms, we are using what's known as "arm asymmetry" for our car. This is typically where, on a left-turning oval car, the right side arm is shorter than the left side arm. This shorter arm applies a heavier force to the right-front suspension relative to the left-front suspension, which causes the body of the car to roll to the left. The simple result of this is that the right-front wheel rate (spring rate felt at the wheel) is increased while the left-front wheel rate remains the same, effectively the same as increasing the rate of the right-front spring. This, in turn, creates a "heave" characteristic in the sway bar itself, meaning it will begin to resist vertical movement in the chassis instead of purely countering body roll.

This can be used to tune steady-state cornering (car is settled in the corner, no more pitch or heave is occurring) characteristics. Increasing asymmetry will help to flatten out the front of the car without adding preload, which can be extremely helpful for an aerodynamically-dependent car. The increased right-front rate will keep the car from rolling left and right in the corner, and provide a more stable aerodynamic platform.

Keep in mind, however, that increasing asymmetry also increases the vertical stiffness of the front suspension, so an increase in arm asymmetry will also cause the front of the car to run higher, especially on the straights, which increases drag. While this can be ignored at a track with short straights, it can spell trouble for a track like Indianapolis where the straights are fairly long.

[kansas picture]

Sway Bar Preload

The final adjustment on the sway bar is the bar's preload. By shortening or lengthening one of the linkages on the sway bar, we can place a load on the sway bar artificially, without the car experiencing any lateral forces. Starting the bar with some amount of preload will roll the car to the left so that the lateral forces wind up rolling the car to a "flat" attitude.

A by-product of preload is a change in dynamic crossweight, or crossweight in the car while cornering. An increase in preload (negative values) will increase the crossweight of the car as lateral forces change, specifically on corner entry and exit. While some drivers dislike preload in any situation, other drivers have been known to like the extra mechanical grip offered on corner exit. Keep in mind, however, that larger sway bars will apply the crossweight change faster than a smaller bar, which can lead to very strange results.

Sway Bar Adjustment

Sway bar tuning has become complex, and is no longer a case of just swapping out diameters to get the handling characteristics you're looking for. The sway bar is often chosen based on the springs used, track characteristics, and aerodynamic needs, so in the same way that a few articles will be dedicated to spring tuning, the same will be done for sway bars.

<u>Special Cases</u>

Indy Car Arm Material

The Dallara Indy Cars in iRacing offer two choices of sway bar arm material: Steel and Titanium (Ti). Since Titanium is a weaker material than the steel used in sway bar parts, and will result in a "softer" acting sway bar than with steel components. This adjustment could be likened to the arm length adjustment found in stock cars, however this specifically concerns the *bending* of the components instead of the dimensions of the components.

In-car Adjustable Sway Bars

[Turner sway bar arms]

It was remarkably difficult to find a good picture of these, so I want to give credit to Turner Motorsport's website for the image above.

Many road-racing vehicles feature a driver-adjustable sway bar in the front and rear, operated by levers in the driver cockpit. These adjustments change the acting angle of the sway bar arm itself by rotating it. When the arm is vertical (Full Stiff Setting), the arm itself is very resistant to bending, and will transmit a large amount of the sway bar's torque forces directly to the suspension. As the arm is rotated to horizontal (Full Soft Setting), the arm is likely to flex as forces are applied to it. This flex will cause some of the sway bar's forces to be "lost", resulting in a softer-feeling sway bar.

Commodore's Garage
Article #12 - Tire Data

by Matt Holden
November 3rd, 2016

Whenever we put a race car on track, be it real or virtual, we have two things that can give feedback on how the car is performing: The driver, and the tires. Drivers are usually very literal when they communicate what the car is doing. If it's starting to spin, they'll say it's loose or oversteering, and if it won't turn, tight or understeering. Going off of only what a driver says can sometimes provide quick-fixes to chassis setup, but can also make a problem worse when something else is causing what the driver reports. Tire data, such as hot pressures, temperature, and even tread wear, can be used together with driver feedback to diagnose issues more accurately and in less time. In this article, I'll use a few examples taken directly from iRacing events. A link to download the spreadsheet I use to record tire data during the NASCAR PEAK Antifreeze Series can be found at the end of the article.

Individual Tire Loads

With load comes friction, with friction comes heat. As we run the car on the track and expose the tires to high loads in the corners, the interaction between the tire and racing surface produces heat. The heat put into the tire can be from sliding, but will also come from the work done by the tire to maintain traction with the racing surface. As we load a tire, it will become hotter, and as a tire becomes unloaded, it will cool off. In many cases, this is the only thing that will be constant in all situations, so it's important to understand how to see which tires are loaded more heavily than others when you finish a run. To walk through a basic tire analysis, I'll look at the first set of tires that came off of Nick Ottinger's #05 car in the 2016 NASCAR PEAK Antifreeze Series race at Kentucky:

[kentucky pic?]

To analyze the load on each tire, we need to get an average temperature for each tire. This is fairly simple since we only have three temperatures to work from, and it will often give a much clearer picture of what the car is doing when compared to the 12 temperatures from the garage tire sheet or the black-box tire sheet. The tire averages were as follows:

LF: 207° RF: 262°
LR: 241° RR: 248°

Nick said during the first run that the car was loose off of turn 4, but he was managing it well. Looking at the tire averages to get an idea of the tire loads, we see that the rear tires are fairly close in temperature, meaning they're seeing a similar load. The front tires, however, have a very large difference in loading, with the right-front tire being *much* hotter than the left-front, indicating that it is experiencing a very heavy load. While he is turning left only, and it's a given that the right-front will take more of a beating than the other tires because of this, a large difference between the front tires such as this is typically not desirable.

This type of tire readout is also an indication of a suspension imbalance. In this specific case, we have the rear end too soft (in roll stiffness) when compared to the front. This results in the rear of the car rolling over and unloading the left-front tire, placing the load removed from the left-front on the right-front. Over a run, this will result in the car progressively becoming tighter, or producing more understeer, as the right-front tire wears and can no longer support heavy loads. In a race, this is basically an un-solvable issue, but at a track like Kentucky where there is very little grip and a lot of bumps, it's often more comfortable for a driver to have this kind of car due to the extra security from the large amounts of rear grip.

Tire Wear

When in practice or testing, it's almost useless to tune the car off of the tire wear, especially oval cars. Since oval cars only turn in one direction, the right-front tire will almost always be the most worn tire on the car, and one of the left-side tires will be the least worn tire on the car. Not to mention that a tire will wear when it's loaded and working hard as well as when it's sliding across the track with no load on it. That said, wear does have its uses, especially when combined with tire temperature analysis. Here are the wear averages from the same set of tires above:

LF: 93.7% RF: 78.7%
LR: 90.3% RR: 78.3%

Looking at these tires, the one thing that's immediately apparent is the large difference in the left-side wear. Nick was reporting a loose-off condition, but the temperatures were showing a tight condition. The wear numbers, in this case, finish the full story and give the message that he was more than likely sliding the left-rear tire more than the others which would result in a loose condition. Since his tires showed that the left-rear was very hot, the adjustment during the pit stop was to remove weight from the left-rear corner via the spring perch, reducing the preload on that spring and allowing the left-rear tire to work a bit more than before. This proved to be the right adjustment and the second run was faster than the first run.

Alignment and Temperatures

One of the major things to look for during practice and testing is any issues with camber, which will only show up on the tire data in the garage. For ovals, we typically want the left-side of the tires to be hotter and slightly more worn than the outsides, so negative camber on the right-side tires and positive camber on the left-side tires. For road courses, we'll typically want negative camber at all four corners, but the amount of symmetry between the sides will depend on the track's characteristics.

A good starting point for camber values is around 10-15° difference between the inside and outside edge temperatures with the wear numbers showing close to the same for all three zones (a few percent is okay). At larger tracks with long straights, such as Michigan, Indianapolis, or Pocono, the optimum camber can provide a larger temperature split than normal, since the tires will ride on the left-side edge for the straights. Small tracks like Bristol or South Boston may require a smaller split, since the distance between turns is short and the outside edges won't cool off between turns. Optimum camber is simple enough to find though, and is usually the setting that provides the most grip with the lowest lap times. If you find what works best for your car, write down the temperature and wear split in the Notes section of the garage so you can fall back on it later or at a similar track.

Starting Pressures

While the minimum pressures are usually dictated by the tire manufacturer (and reflected in the iRacing sim), we still have a bit of room to work with starting pressures. In the broadest terms, we would like to have the pressures set so that the center temperature on the tire sheet is between the inner and outer temperatures. As of writing this article, the current tire model averages the center temperature and will show the center temperature higher than what would be measured in the real world (Read Mr. Kaemmer's explanation here on the forums: http://members.iracing.com/jforum/posts/list/1808602.page). With that issue known, the best setting seems to place the center temperature around the temperature of the inside edge of the tire. With that set, further pressure adjustments can be made based on the wear numbers: If the center is wearing too much, reduce the pressure. If the center is not wearing, increase the pressure.

Starting pressures can also be changed for the situation as well. Our fastest laps will occur when the tire is at optimum working temperature and pressure, so getting those to match is key to a fast car. For short runs, especially qualifying, it's common to see pressures raised quite considerably in an effort to get the car up to speed quicker. Longer runs will need lower pressures to allow for pressure buildup over many laps, with the aim of hitting the optimum pressure at maximum buildup.

Pressure Build-up

Tire pressures can make or break a race, and among the many things we have to deal with through a race, pressure buildup is one of the most frustrating. As the tires heat up, the heat in the tire will heat the air within the tire and cause the pressure to rise. Over a full run on a set of tires in oval racing, it's possible to see upwards of 20 psi buildup on the right-side tires, so it's important to make sure the tires are building up to the proper pressure through a run. The garage shows us Hot Pressures next to our Cold Pressures, and the hot pressures can be just as important to diagnosing a car as the temperature numbers themselves.

The next article will take an in-depth look at tire pressures, but tire pressure changes the "spring rate" of the tire. This spring rate is very similar to the suspension spring rate, and it has a big influence on how the car drives. The biggest thing to look at is the difference between hot pressures on a given side of the car. For instance, if we set the cold pressures like this...

LF: 25psi RF: 45psi
LR: 25psi RR: 45psi

And hot pressures end up like this...

LF: 32psi RF: 65psi
LR: 34psi RR: 62psi

We have a problem! Knowing that the tires on each side of the car are the same for oval racing, equal pressures would produce equal spring rates. Thus, when we start, we have equal spring rates on the left tires and equal spring rates on the right sides. However, once we've run the car a while, we start getting higher spring rates on the right-front tire and left-rear tires. Plus, the left-front and right-rear tires will be a little flatter on the track than their counterparts and will produce a little more grip because of that. These two issues will undoubtedly lead to a tighter car late into a run.

For this situation, we need to adjust the cold pressures to compensate for the difference in pressure buildup. For the right-sides, we have a 20psi build on the RF and a 17psi buildup on the RR, so we can either add 3psi to the RR, take 3psi from the RF, or go halfway with each and take 1.5psi from the RF and add 1.5psi to the RR. Any of these will help the issue, and we can do the same type of change to the left-side tires. Getting the pressures to build up evenly can help address issues with the chassis, since a difference in pressures would add another set of spring rates to the car's behavior.

Telemetry Temperatures

[telemetry trace]

The temperatures shown in the garage are the temperatures that would be taken using a pyrometer in the tire carcass. Essentially, these are the temperatures of the tire's tread based on how much work it's done *during the run*. On the other hand, telemetry exports display a more instantaneous temperature reading, and can differ from the garage temperatures quite considerably. Basically, the telemetry temperatures can be thought of as being taken from infrared sensors around the tire, and can give a good indication of how much of the tire is in contact with the pavement at a given moment. If the telemetry trace is showing a steep drop in temperature, there's a good chance that section of the tire isn't contacting the pavement and is doing no work at all. These temperatures, and especially the rate at which they are changing, can be extremely effective in tuning corner performance, while the garage temperatures are better for tuning long-run handling and overall balance.

While each set of data for tires has its own advantages and disadvantages, it's important to consider *all* of it when setting up your car. This is all given to you with no work needed and is the only way your car can tell you accurately what's going on, so don't ignore it!

Commodore's Garage
Article #13 - Tire Pressures

by Matt Holden
November 18th, 2016

Tires are a racing vehicle's only form of contact with a race track and, as a result, can be the reason why a car is a race-winner or a mid-fielder. Many factors can determine how well a tire is gripping the track surface, from tire compound, to load, and especially its own behavior when exposed to different loads. Tire pressures can directly affect how a tire responds to various loading conditions, and we'll take a look at what goes on with the limitless tire pressure adjustments.

A Lot of Hot Air

Pressure is a value of force for a given unit of area. For tire pressure, the most common units are Pounds per Square Inch (psi) and the Pascal (Newton per Square Meter). Under static conditions, the tire's internal pressure is set by simply adding or removing air inside the tire. This is fairly simple, and most everyone knows about this from driving on the highway. In motorsport, however, we're concerned not only with the cold pressure, but the *hot* pressure as well. Cold pressures (set in the garage) are the pressure when the tire is "cold" or at ambient temperature. Once the car has gone on track and the tire is worked, we now have what is known as the "hot" pressure.

Hot pressure can be influenced by a few things, mainly heat from the tire carcass (tire loading and friction) or heat from around the tire, usually the engine or brakes. The difference between the cold and hot pressures are called "build-up", essentially how much the tire pressure increases as the tire is worked on the race track. Different tires will transfer heat to the air or gas inside the tire differently, so while you may know how quickly the Goodyear Eagle builds up in NASCAR's vehicles, it may not be quite the same as the short-track tires on the Late Model or Modifieds, and *definitely* will not be the same as the Pirellis on the Formula 1 cars.

Pressure and Contact Patch

The most immediate result from pressure adjustments is a change in the tire's contact patch, or the part of the tire that is in contact with the pavement itself. In relation to the size of the tires themselves, the contact patch can be remarkably small, and pressure adjustments can have a dramatic effect on how large this area is. As you might expect, a larger contact patch will provide the opportunity for more grip, while a smaller contact patch can result in less grip. At lower pressures, the tire is allowed to "squish" a bit more into the track, the same way play-dough can be smashed into a surface and spread out. Higher pressures produce a more rigid tire, and with that a less-compliant tread on the tire.

In addition, pressures can also have an effect on how stiff the tire is laterally, with a direct effect on how much the tire will "roll over", or fall over onto the sidewall. At very low pressures, drivers often claim that the car feels as if it is moving around on top of the tire, not a very confident feeling at all. At higher pressures, the tire sidewalls will become stiffer, and keep the tread of the tire in place under cornering loads. While this can provide a more positive feel to the driver, it can be a problem when pressure build-up is considered (more on that below).

Pressure and Tire Spring Rate

Where the tires really start to affect handling is the tire's spring rate. Similar to the springs installed in the suspension, tires will compress and exert an opposing force to any vertical loads the car feels while going around the track. This spring rate, like the suspension spring rate, has a direct impact on the car's handling, especially in oval racing. In stock car racing, the tire spring rate can vary from track-to-track, especially in NASCAR where the tire supplier brings a unique tire to suit a track's conditions. The rate can also vary from left-to-right, with the left-side tires typically having a lower spring rate than the right-side tires for a left-turn oval. For example, the Goodyear tire used in 2015 at Charlotte indicated (on the data sticker) a right-side spring rate around 1410 lb/in and around 1350 lb/in on the left-side tires. (*NOTE: These figures are determined under controlled conditions and intended to help "match" tire sets, used here as an example of general tire behavior.)*

As you might expect, we're always looking to tailor the pressures (and spring rates) to a given track. At shorter tracks with less vertical load, we can get away with much lower pressures. Tracks such as Martinsville, Langley, or Stafford Motor Speedway allow for extremely low pressures to be run without major issue. This produces a lot of mechanical grip due to lowering the tire's spring rate, as well as increasing the size of the tire's contact patch, which adds up to a lot of grip! On the other hand, big tracks such as Michigan or Charlotte, and even the faster short tracks such as Bristol and Iowa, need higher tire pressures to deal with the vertical loads and higher speeds. For example, it's not uncommon to see left-side tire pressures lower than 10psi at Martinsville in the Cup series.

Pressure Tuning

Now that the basics are out of the way, we'll start looking at how to set pressures in the garage. Cold Pressure starting points will be primarily influenced by the driver's style. For instance, I tend to build up the RF tire very heavily and the left-side tires evenly, while Nick Ottinger tends to build up the right-side tires evenly and the LR tire heavily. I've seen other drivers who build up the RR tire heavily, and know of real-world drivers who were notorious for causing both front tires to build up tremendously. This has to be considered when setting the initial pressures, and is simple enough to find in a few test sessions. In addition, setting the left-to-right difference in pressure can be important as well, depending on what the driver is looking for. For short races where speed is key, a narrow split between each side can be useful in giving the driver confidence early in the race. Longer runs may call for larger differences in pressures, and the driver's ability to deal with low-pressure tires is important in this situation. I know of situations in the NASCAR Class B series where drivers have won races on as much as 25psi from left-to-right, and even as low as 7psi between each side. It's simply a matter of what is necessary, and how much the driver can tolerate.

While it's impossible to get all four tires to build up *exactly* the same, you can adjust the cold pressures to prevent major differences in what the pressures build up to over a tire run. If you have your car set up properly, but the right-front tire builds up 5psi more than any other tire on the car, it will inevitably have a significant understeer condition, and there's almost nothing you can do to the chassis that will fix that. A simple adjustment to the cold pressures can bring the build-up closer together, resulting in a more consistent feel for the driver throughout a run.

Setting the pressures based on tire temperature and wear is also extremely important. If the temperatures on a tire show a hot or worn center, it's a sign that the pressure is too high, and center of the tire is effectively "lifting" the outer edges of the tire. Similarly, a cool or un-worn center is a sign of an under-inflated tire. Both conditions can result in a loss of available grip and should be adjusted accordingly.

In-race Adjustments

In a race, it's important to know how your car builds pressure, and which tires are likely to over-build in a long run. In the event that your car starts to swing one way or the other on handling, a look at the notes will tell you which tire may be causing the issue, and simply your decision on what to adjust when it comes time to pit and make adjustments. For instance, if all of your practice runs showed that the left-rear and right-front tires were building up more than the right-rear and left-front, a swing towards a tight condition during the race is very possible. In that event, reducing the pressures in the right-front or left-rear tires would be a very safe adjustment, and will likely alleviate whatever is happening on-track.

The most important influence on how pressures build up is simply the driver. How the driver approaches the first few laps on a set of tires can lay the groundwork for how the car will change through a run. If a driver pushes very hard and makes aggressive moves early, it can build the pressures up differently than if the first few laps were more conservative. Drivers with a good eye for how a race is progressing can change how they "break in" a set of tires in the first few laps. If a race is becoming a series of short runs, they may push harder after restarts, since longer runs aren't expected. However, if a race is producing long green-flag runs, they may run more conservative laps following a pit stop, expecting to be on that set of tires for a long time. If done properly, the difference in lap time at the end of the run can be substantial, making late-run passes very quick and easy.

Another adjustment available to the driver is brake-induced pressure buildup. At a track with any form of braking, heat from the brake systems will soak into the tire, increasing the pressure. It's not uncommon for some drivers to adjust the brake bias in an attempt to build up pressure in a given tire to cure handling issues. For instance, if a car is becoming loose, or oversteering, the driver may shift the brake balance forward in an attempt to heat the front brakes up and increase front tire pressure slightly. While the brake adjustment can fix corner entry issues, the resulting pressure change can cure issues in the center and exit of the corner. This is not a significant adjustment, however, and does *not* replace simple pit-road pressure adjustments!

Notes are Important!

The notes section of the garage should *never* be ignored, and is a great place to record what pressures were for a given practice session. Whenever you go to a different track, it can give you a massive head start to look back at a similar track and see what pressures you ran at the previous track. Or, if you run a track for a second time in a season, this can also be a big help! It's a good practice to include in your notes section what the pressures were building up to over a run, and what you may have adjusted to during the race, in order to make changes to compensate at the next track. Tracks with similar characteristics, such as Iowa and Richmond or Texas and Atlanta, can often use the same pressure settings for the tires and it's important to know what happened in the last race so you aren't doomed to repeat the same fate. On a similar note, if you are in the process of developing a new setup for a car, keeping pressures consistent is key to understanding what is changing from track-to-track. Eliminating unnecessary problems with the tire pressures can help you focus more on the chassis itself, and save you from a lot of headaches in the long run.

Commodore's Garage
Article #14 - Ride Heights, Perches, and Deflections
by Matt Holden
December 2nd, 2016

The last major piece of the puzzle we need to look at before diving into building a setup are the three simplest things in the garage: ride heights. Older (and even some newer racing games) have a "Ride Height" adjustment in their setup garages, completely independent of the rest of the car. Anyone who's ever worked on a car at all knows that simply breathing on part of the car in the wrong way can move the ride heights around, and iRacing included that extra bit of realism (some call it a "frustration"!) to their garage system. Where ride heights have been a set-and-forget setting in the past, completely ignoring them can run you into some major issues that may seem un-solvable. In addition, ride heights (and the requirements for those ride heights) are where every setup begins, so now that you [hopefully] know the rest, let's look at the final variable!

Ride Heights

Ride heights are simply a measurement of how far a car is from the ground. In many cases, this is *not* the lowest point on the car, and while the ride heights may read as 3 inches (or 7.6cm...), there could be part of the car that is much lower, such as suspension components or body panels. In addition, most teams have their own way of measuring the car's ride height, and it can vary from car to car. For road racing, ride heights can be measured in multiple places across multiple teams, even if they run the same car. It's important to learn where these ride height points are measured so you know what you're dealing with at each track. For oval racing, it's a bit simpler, since most oval cars measure their ride heights at the frame rails at the bottom of the "door" panels. Below is a picture of the NASCAR Xfinity Camaro with the approximate front and rear ride height points highlighted. Keep in mind that the ride height is measured at the *frame*, not the skirt, in the garage!

[Alex xfinity pic]

Ride heights can be determined in multiple ways, primarily dictated by the series rules. Almost every racing series has a minimum ride height rule in place, so setting the car at those minimums is often the way to go. Beyond the rules, track characteristics play into the necessary ride heights as well. If the track is bumpy, like Kentucky or Lime Rock, or has a lot of vertical loading changes, such as Zandvoort, higher ride heights may be necessary to keep the car from contacting the racing surface. On smoother tracks like Kansas or COTA, lower ride heights and a more aggressive aerodynamic attitude is possible. Setting the car too low can be an aerodynamic advantage, but on a rough track can cause an excessive loss of mechanical grip, and a loss in speed.

In some cases, minimum ride heights at all four corners is *not* the optimum situation for a given car. Testing will help you find what you like best, but some cars perform better with extra *rake* set in the car, or higher rear ride heights. This can provide an aerodynamic benefit, resulting in more front downforce and increased oversteer, while other cars may behave better with a flatter set of ride heights. For oval racing, it's common to only have one corner set at the minimum ride height, typically the left-front corner, with the highest usually being the right-rear.

[Gen 6 ride height rules]

Your starting ride heights will often dictate the spring package you choose in your car. In every situation with every car on every track, we'd like the car to settle down and drop below the minimum required heights and "seal" off the bottom of the car, producing maximum downforce. Where we start the car in the garage tells us what we need to do to achieve the lower ride heights on track. If we have to start the car up higher for either track characteristics or rules, we'll need a softer set of springs to allow it to drop under aerodynamic loads. Conversely, if we start the car lower, we'll need stiffer springs to keep the car off the track surface.

Coil-binding and bump stops can go against those general rules, however, since we can start the car with a softer riding spring rate and wind up on a much stiffer rate after compression. This is most often used to circumvent and essentially ignore ride height rules by installing springs that will keep the car at legal ride heights when stopped, but utilize a stiffer spring package for better aerodynamics while on track. How you choose to set up your car is up to you, and I suggest trying multiple packages to see what fits your driving style best. When we start looking at putting together a spring package for a track in the coming weeks, we'll consider all of these options and go into each in more detail.

Spring Perch Offset/Shock Collar Offset

The way to adjust ride heights on a stock car in both real life and iRacing is through spring mount offsets. On a "big spring" car such as the NASCAR Truck, Xfinity, Cup, and National cars, this is the *Spring Perch Offset* adjustment. On a coil-over car like the Late Models, Modifieds, and Legend cars, this is the *Shock Collar Offset* adjustment. In both adjustments, the upper mount is being moved up or down (in relation to the chassis). This moves the car up or down, adjusting the ride height in turn. Due to the tight packaging for modern race cars, it's very difficult to get a good picture of how these are mounted in a car, however I did find a way to show it. Below are two images, the first is a NASCAR-style Big Spring front suspension assembly, the second is an AFCO racing coil-over shock. The first image is actually a display at Hendrick Motorsports' Museum, so if you're ever in the Charlotte area, you can head over there and take a look at it in more detail if you like!

[spring images]

The Spring Perch/Shock Collar Offset value is typically a measurement from a "zero" value for that component. For both of these, the zero setting would be where the perch or collar is as far up as it can possibly go in the adjustment range. Values are typically negative, since they are a distance *downward* from the zero setting. Any positive value would be an adjustment of the perch upward, any negative would be an adjustment downward. So for instance, if we have a Spring Perch Offset of -2.500", we've set the upper perch 2.5 inches below its highest possible adjustment. This does *not* mean the spring is compressed 2.5 inches, and in fact has no meaning for the spring at all, it's simply a way to locate the upper perch or shock collar.

[bjones]

This adjustment has a further use in allowing various length springs to be used in the same car. Springs are rarely manufactured at a specific height across a range of rates, and auto racing is no exception unless you're willing to shell out a lot of money. In fact, the only exception to this I know of is Legends racing, where springs are required to either be 8" or 10" long, with the shocks having an adjustment to accommodate either length. Still, those springs will vary in length by a few tenths of an inch, especially if they're showing signs of fatigue. Whenever you change spring rates, it's more than likely going to introduce a new spring length in the process. This will always change the ride height on that corner, and to correct that you *must* reset the ride height using whatever spring perch adjustment is available. Failing to return ride heights to what they were prior to a spring change can introduce a lot of "adjustments" that weren't intended, from sway bar preload changes, alignment, aerodynamic changes, and spring preload changes all around the car, so while you intended to change just a spring, you wind up changing the entire car if you don't reset the height on that corner. It's very simple, just use the perch or collar to change the height on that one corner (not the other three!!), and only takes a few seconds but will save you a lot of time and headache.

Spring Deflections

Spring deflection doesn't have a major influence on your chassis setup, but it does need to be covered. I figured that the perch/collar offset section would be the best place to put it, since the deflection is in the same realm of thought as the perch offset. I know some guys who build their setups specifically off of the spring deflection values, and I know others who do not. Both approaches have worked well, even in the same cars, and both approaches won championship *in the same car* in 2016. Personally, I've found that a certain arrangement of deflections have worked best for me, which I'll explain after going over deflection in general.

Deflection, in an engineering and physics situation, is the distance anything has moved from a point of equilibrium. Equilibrium being where this object is when no external forces are applied to it, deflection would be how deformed that object is once a force is applied to it. For example, if we have a beam that is perfectly straight with zero forces applied to it, it would be said to be in equilibrium, since all forces acting on it are "equal". When we go and push on one end of the beam, it will "deflect" from that state of rest. In a spring, its length without an external force is its length in equilibrium. If we compress it by one inch through an external force, we now have one inch of *spring deflection*.

Knowing what we learned about springs in earlier articles, these spring deflection values are a way to look at the static force being exerted by the spring in the garage. The deflection in the garage is represented by two values, typically in this format:

4.896 in of 7.850 in

The first number is how much our spring is deflected, the second number is how much travel is available in the spring. So, in this case, we have about 2.954" of travel left in the spring before binding. The first number can be further used to determine the static force from the spring itself by simply multiplying it by the spring's rate. So if the spring in this situation is a 500lb/in spring, *500 x 4.896 = 2448 pounds*.

In my own experience, I've found that a certain arrangement of static forces wind up producing a setup that I like, while setups that don't fit this arrangement don't do what I would like for them to do. At the end of the 2016 Class B season, I gathered all the setup files I'd used in the full 33-race season and calculated the static forces from the springs on each car. I knew which setups I liked and which ones I didn't like, and for each set of static forces I calculated a "crossweight" force in the same way you'd calculate crossweight on the corner weights themselves. For the setups I liked, this value showed as between 48% and 49%, while the setups I disliked the most were above 52%. This isn't a hard-and-fast rule, and the figure is hardly "real", but the trend was surprising, and a prime example of why testing is important in finding what you like in the car. I can now take this information and apply it to 2017 with the intention to fix the tracks I struggled at this season.

Commodore's Garage
Article #15 - Setting Spring Package

by Matt Holden
December 2nd, 2016

We now know the basic ins-and-outs of the spring components in our race car, but we need to apply that to the car and get started with our race setup. "Where do I start?" is probably the most common question asked, and in almost every case, the answer is the *spring package*. We have to look at many different things that will affect the spring rates and consider all of them, eventually winding up with a spring package we are comfortable with. If the springs aren't set properly, it will lead to handling issues during the race such as handling swings from imbalance, over- or under-worked tires, and possibly even a crash! This article and the next one will look at getting your springs arranged leading up to the holiday break.

Aerodynamics and Minimum Ride Heights

Most cars will have a minimum ride height defined by the series rules. Whether that's 7" or 7 centimeters, there is no getting around it while you're in the garage. For the most part, ride height rules could be deemed more traditional than necessary, since most teams have found ways around the ride height rules. Overcoming the limitations presented by ride height rules is the first thing we need to consider when choosing our springs, and how we go about overcoming those rules can often be the difference between a leading car and a mid-field car.

First, we'll look at an aerodynamically-dependent car, or one that relies heavily on downforce (or sideforce) to get it around the track. Formula cars, NASCAR's Cup/Xfinity/Truck cars, Prototypes, and, to an extent, some GT and short-track oval cars. Aside from the NASCAR Cup car, we're going to have ride height rules that keep the car elevated above what we want. For best aerodynamic performance, we want the car low, so we need to get around the ride height rules quickly once we're on track and out of the inspection bay. We have three options in this situation: 1) Deal with it, 2) Bumpstops, or 3) Coil-binding.

The first option, "Deal with it" is the common choice for cars without bumpstops (or limited bumpstop options). If you don't have the option to effectively drop the car and keep it from slamming into the track, the best option is probably going to be going with springs that will keep the car off the ground, but are soft enough to give the necessary mechanical grip to get around the track. A common example of this in the iRacing service would be the V8 Supercars. The Supercars generate a bit of downforce, but lowering the car for aerodynamic purposes can often hurt laptimes more than raising it up and sacrificing a bit of aerodynamics for mechanical grip.

The second option involves dropping onto bumpstops from aerodynamic load. This is very common in high-downforce racing cars, may or may not be adjustable in the garage screen. In these situations, softer springs can be used with the intent that they'll collapse under aerodynamic loads and the car will ride on the bumpstops in high-speed corners, producing a lot of downforce. As the car slows down for a slower corner, it may rise off the bumpstops and use the soft springs for mechanical grip. Keep in mind that bumpstops are often very high rate springs, usually progressive, so a coil spring of 1000 lb/in (175 N/mm) will often be much softer than whatever the car was riding on as a bumpstop.

Finally, we have the dreaded coil-bind. I did an entire article on the basic mechanics of coil-binding here: http://www.iracing.com/commodores-garage-7-coil-binding/. For the crash course on coil-binding, we simply want the car to have an extremely soft spring rate so that, once it's up to speed on the track, the springs collapse quickly from downforce loads. The result is that the entire spring is compressed to where all of the coils are touching each other in some way, preventing any further downward travel from the springs. This is extremely effective, however head-scratchingly complex to figure out and make it work properly. While coil-binding went out the door in Sprint Cup following the 2014 ride height rule elimination, it is still present in Xfinity and Truck, whether by full-coil bind or half-coil "pig-tail" binding. We don't have the means for half-coil binding in iRacing, but full-coil is still a viable option in these cars.

Finally, we have the case where the car is not an aerodynamically-dependent workhorse at all. Cars like the Legend Car, the Modifieds, or the Lotus 49 all have very little aerodynamic dependence (or none at all). In these situations, the thing to consider is how far we need to drop the Center of Gravity to produce a good-handling car, if we need to at all. Again, it's going to be a balance of advantages and disadvantages for these cars, so the only thing that will produce a good base to work from is old-fashioned testing. That said, none of these cars will behave well with excessive spring rates. Spring rates that would be considered extremely soft for the cars in the other three categories will be perfect for this range of vehicles. All we're looking for here is mechanical grip, and that will come from softer springs.

Once we know what suspension system we're working with, we should know how we could attempt to get around the ride height rules, if we can at all. With that known, we can look through our range of spring options to choose what springs to pair up with our aerodynamic needs. On some cars, such as the Mercedes GT3 car, this range can be extremely small, where it can be a relatively wide range on other cars, such as the BMW Z4 GT3 car. Driver preference can also be a factor, especially in whether or not the driver wants a stiff or soft car. If the driver wants a softer spring package, the car will need to be raised up higher to start, while a stiffer spring package will be just fine with lower ride heights.

Track Characteristics

The other major factor in spring choice is the track itself. Is it bumpy? Smooth? Are there hills? Is it flat as a board? While the aerodynamics and ride height rules for a car will determine the general range of springs to run, the track itself will narrow this down even further. Again, as is with every other thing, testing is the only way to determine what springs you like and what will work best for you.

A major track characteristic that will affect your spring choice is how much of the track is high-speed and how much is low-speed. If the track has a lot of high-speed sections, such as Monza or La Sarthe, you're going to want to go with a relatively stiff spring package and set the car to generate a decent amount of downforce while having the car trimmed out for low-drag. Soft springs at these types of tracks will produce a large amount of vertical movement in the car, which can result in a lot of drag if the car starts moving excessively. Onboard videos from Formula 1 cars at Monza are filled with the sounds of the car skipping along the pavement, as well as virtually no movement from the suspension arms at all. The opposite would be somewhere like Zandvoort or Summit point, where slow corners dominate a lap. At these tracks it may be difficult to generate a lot of downforce, and you're going to want to lean more on mechanical grip and soft springs to run a quick lap time. In general, you can think of the spring rates going up as downforce goes down as a general rule of thumb.

On ovals we would consider the same thing: Higher-speed tracks will need stiffer springs. Big 1.5-mile speedways will need a stiffer overall spring package than a short half-mile oval, and the same rule about downforce-vs-spring rate will apply. However, we have to further consider the amount of banking on the oval as well. Higher-banked tracks such as Charlotte or Bristol will produce very high vertical loads, and to keep the car from slamming into the race track we need high-rate springs. Small, low-banked tracks will see the lowest spring rates, and real-world rear springs under 200 lb/in are not uncommon in the upper levels of NASCAR at a track like Martinsville.

One of the rare things that applies to both road and oval racing is the bumpiness and grip level of the track surface itself. Smooth, high-grip tracks can allow much stiffer spring rates and a more "aggressive" approach to generating downforce. Old, worn-out tracks like Kentucky Speedway or Lime Rock Park will produce a need to deal with bumps to keep the car from jumping around and eventually finding its way to a retaining wall. Softer springs are necessary for these tracks to keep all four tires on the pavement, as well as reducing complaints from the driver. A track such as Kansas Speedway or Circuit of the Americas, where the pavement is very smooth and void of large bumps, doesn't produce the need to deal with those bumps and a stiffer spring package can be used.

Putting it all together, forming a baseline

Once all of these options are considered, a general idea of what springs to use should be very apparent. For instance, if we're setting up a NASCAR Cup car at Charlotte, we know that we don't need to get around ride heights (soft springs aren't necessary), we would like as much downforce as possible (stiff springs are desirable), the track isn't extremely bumpy (stiff springs), and the banking is quite high (stiff springs). Looking at all that, we're going to probably run a stiff spring package at Charlotte and attempt to generate as much downforce as possible while all but ignoring mechanical grip. However, Atlanta Motor Speedway has the same banking as Charlotte, is the same length as Charlotte, but is tremendously bumpy in comparison. We could take our Charlotte spring package to Atlanta, but it would likely bounce around and we'd have no grip. Instead, we need to adapt with softer springs for the bumps, which means we'll need to raise the car a little bit to keep it from hitting the track, and we'll sacrifice a bit of aerodynamic grip in lieu of some mechanical grip.

I tend to beat the testing horse beyond death, but I cannot stress how important it is to actually *test* various configurations. A lot of sim-racers have the idea that "testing" is simply where you go build a setup for a track, which it isn't. To get what is necessary to build a setup week-by-week and know how to adapt a car for different conditions requires true *testing*. If you're going to race a new car (or if you feel you still struggle with a car you want to drive), find a track you like, and try all kinds of different things to see what works, and what responds best to *you* as the driver. Try stiff springs, try soft springs, try a combination of both! Move the car up and down on ride heights, see if something clicks and you hit on something that you're comfortable with. Once you have the car to your liking at that track, that setup becomes the "baseline". Those springs are what you're going to base every other setup on until there's a major change to the car (such as aerodynamics, tires, or weight). Starting over with a new setup at a new track (or doing the "setup dance") simply puts you back at Step 1 every single time, and you have to re-learn what works with that spring package. In many cases, one spring will apply to many tracks, and in some cases, it will apply for an entire season. So don't over-think it, and never throw springs out just because you've gone somewhere new!

Commodore's Garage
Article #16 - Adjusting Spring Package

by Matt Holden
December 24th, 2016

As we approach the holidays and the close of 2016, instead of putting more information out there, I'd like to send everyone on a mission. It's pretty simple, and the few weeks we'll have should be plenty of time: Get your spring package right! In previous racing sims, that was pretty simple, right? iRacing's garage is a little more complex, and throws in new challenges, but it's the same challenges that real-world teams see on a weekly basis. In the last article, we looked at what will influence your spring package's characteristics, and this time, we're going to look at how (and when) to make spring changes to get the car behaving more like we want it to behave on-track.

When Changes are Needed

I can remember when I started sim-racing 10 years ago that spring changes were how I got the car to go around the track. Rotating springs in and out of the car like they were candy was commonplace in race practices, but that really isn't a very realistic situation. Watch a high-level racing series practice on television and you'll quickly notice that spring changes are rare, and the commentators will likely make a big deal out of it. In today's aerodynamic and engineering-based world of NASCAR, a spring change in the track garage usually means somebody messed up on the in-shop simulator. It's very important to understand that the vast majority of your handling will come from weight placement, shocks, alignment, and roll center locations, while the springs will have a greater effect on the car's long-run handling trends. Yes, a bad spring package will cause the car to handle quite badly, but changing out springs constantly can lead to headaches that you may never figure out, and late-race handling issues that can take you out of contention.

Knowing that, the question becomes "When is a spring change necessary?" Three major factors can be the warning flags that will identify spring changes:

1) Is the car changing heavily throughout a run? Does it start loose and abruptly become tight? Does it start "okay" and then refuse to turn?

2) Is one tire consistently *much* hotter than the other three? Further, are two tires on one end of the car the same temperature?

3) Is the car hitting the track? (Walls do not count here....see #1 for walls)

If any of these produce a "yes" from the driver, it's possible that a spring change is necessary. If the driver is just saying, "It's understeering here", or "It's a little free here", that doesn't necessarily mean a spring change is the fix for the problem.

Spring Change "Physics"

Whenever a spring is changed we're going to possibly see a few different changes in the spring itself. Obviously we're going to see a change in rate since we don't have the ability to move to a longer/shorter spring with the same rate. Many of the cars in iRacing now have different spring lengths for each rate, which we need to adjust for as well. In special cases, we can also see a change in the number of coils, which directly affects coil-binding cars. For this article, I'm intentionally avoiding bump-stop suspension for now in an effort to start simple and add complexity as we go.

The major thing to understand about springs is that we're *not* going to change the load on a corner through a spring change by any amount that would be significant. For constant conditions, a wheel will see a given load. This is based on the track, speed, downforce, roll center and CG locations, and any forces produced by the chassis itself, such as from the track bar. Whenever we change the spring, we *will not* significantly change the load the wheel sees. The spring is meant to support that load, not alter it, and as a result will not change the force acting on the tire if everything else in the car is kept the same.

To understand this, we have to look at the math behind spring behavior (it's simple, stick with me here). I mentioned *Hooke's Law* in a previous article about springs, and that applies here as well. Let's say that we see a 2000 lb load pushing up on our wheel, and we begin with a 500 lb/in spring. So from the law, we have this:

$F = k*x \;\text{-->}\; 2000\ lb = 500\ lb/in*x \;\text{-->}\; x = 4\ inches$

The 500 lb/in spring will travel about 4 inches with this given load. So if we change that to a 1000 lb/in spring, we'll have only *two* inches of travel from the new spring, but the load is still 2000 lb/in (1000 lb/in * 2 inches = 2000 lbs). Our load at the wheel didn't change, but we will inevitably see a much sharper increase in resistance from the spring and the load at the wheel will be countered much faster. Furthermore, we're also going to see a change in dynamic ride height, which alters the aerodynamics if the car is capable of that sort of thing.

The softer spring will allow the wheel to travel more easily, and keep the tire in contact with the track much longer over rough surfaces. The stiffer spring will reduce the amount of movement, meaning the wheel won't be as forgiving over rough surfaces, but will (hopefully) keep the car at a consistent height and prevent up-and-down movement, which can produce some wacky aerodynamic shifts in balance.

Balance Issues

The most prominent sign that a spring change is necessary is long-run balance issues. This has been covered so far in the articles on Roll Stiffness (http://www.iracing.com/commodores-garage-5-roll-stiffness/) and the Tire Data article (http://www.iracing.com/commodores-garage-12-tire-data/) if you want to check those out and brush up on the knowledge.

Basically, we're looking for any under- or over-loaded tires via the tire temperature data screen. You'll get the best result from extended runs in practice, usually 10-20 laps is enough to uncover issues.

[Temp picture]

We'll use this data set for an example. Once the tires are off the car, we need to average the temperatures (we're ignoring wear) and consider what was reported, in this case the car got a lot tighter through a run. We can see that the right-front tire is extremely hot, the left-front tire is cold, and the rear tires are nearly the same. This is a sign that the rear suspension is rolling over, and is too soft in roll stiffness. We need to either stiffen the rear or soften the front, and the most effective adjustments for this would be either roll bars or roll center (track bar height, on oval cars). Softening the front end by reducing the right-front spring rate would also work, however on a bumpstop front end it could play havoc with the handling and actually do the opposite (covered in a future article). These tires came off of the Gen 6 NASCAR Cup Car, and the following adjustments would be a step towards fixing the issues:

-Stiffer right-rear spring
-Softer left-rear spring
-Higher overall track bar
-Smaller front sway bar
-Softer right-front spring rate

The right-rear spring would be a challenge because of the aerodynamic issues it could bring about. Depending on how stiff it was to begin with, we could actually *increase* rear downforce with a stiffer right-rear spring by holding the rear of the car up longer, which would actually cause more understeer, or a tighter car. Like I mentioned, the right-front spring would pose many issues with a bumpstop, so we're probably going to avoid that. The best option would be either a soft left-rear spring or a higher track bar, with the spring change being a much larger adjustment than the track bar height. Either one would work, however there's another thing to consider with this data set....

Changing for Individual Tire Temps

I used this specific set of data (which came off of Nick Ottinger's car in the 2016 NASCAR PEAK Antifreeze Series race at Kentucky) because of the difference in rear temperatures. Whenever we see a tire that is slightly cooler than its counterpart (in this case, the right-rear tire), softening the spring would increase the temperature on that tire a little bit. This is counter to what most people have learned through sim-racing setup guides, but it stems from having the tire in contact with the pavement longer, and the softer spring allowing the suspension to deal with bumps and undulations in the track surface. The longer the tire is in contact with the pavement under load, the longer it will be under friction forces, and thus generating more heat.

So, with all of these things considered, we know that the car is imbalanced with the rear too soft (rolls to the right too easily), and the left-rear is cooler than the right-rear. Considering all that we know about the car (the rates themselves are irrelevant here), we can reliably assume that the left-rear spring is too stiff for this situation. Plus, since we know that Kentucky is a very bumpy, low-grip track, this adjustment makes sense for the track characteristics and should produce a freer car with less swing towards understeer during a run.

Making the Change

Before a spring is changed, we need to pay attention to a few things and make sure we return these values to where we started:

-Corner Ride Height
-Sway bar preload(s)
-Bumpstop Gap/Shock Deflection
-Alignment for the corner

Everyone knows the rule to "change only one thing at a time", and this is where that rule is most important. If we don't make sure all of these things go back to where they were, we've changed more than just the spring. Ride height changes produce aerodynamic changes, bar preload changes produce sway bar changes, bumpstop gap changes result in contact timing changes, and alignment changes result in contact patch changes. All of these can produce unexpected results, and piling them on with multiple un-adjusted spring changes is like jumping down a big hole with no ladder.

[correction image]

In the above image, I've changed a spring on the NASCAR Class B Xfinity car's left-front corner and pointed out the steps necessary to make the change. Here, we want to make sure that everything possible returns to where it started, and the major step here would be resetting the ride height following the spring change. We start with a 500 lb/in spring with the ride height at 5.54". The first step is to change the spring and let the car settle (the "Apply" button does this for us). Following the spring change, the ride height went to 6.78"...which is illegal...and the corner weight increased by 64 pounds, which likely comes from the change in sway bar preload from the new ride height. It's a simple fix though, and we just need to lower the car via right-clicks on the Spring Perch Offset until the ride height is as close as we can get it to where we started. In this example I returned the car to 5.56" ride height by raising the perch about an inch.

Another thing to consider here is that I got the shock deflection back to where it started, with only a 0.02" change in deflection. This is incredibly important for a car like this, where we would use bumpstops! I've changed the main spring, but I haven't changed the bumpstop gap, which means the bumpstop will engage at the same point in suspension travel as before. Before the ride height was reset, the shock was at 4.63" of deflection, meaning it had expanded by a full inch. This would add an inch to the bumpstop gap, meaning it would take that much more travel until the bumpstop engaged. Our smallest shim option is 0.063", so in the event that the spring adjustment changes the bumpstop gap by more than that, we need to re-shim the bumpstop to compensate.

Race-testing Spring Changes

[indy pic]

I'd like to end the year on an example of how we went about a major spring package change in 2016 on the Class B series cars, specifically a situation where I tested a new spring package for Alex Scribner around the season's midpoint. Prior to Indianapolis, we'd struggled with bump spring timing, and major changes with the car's behavior on and off the bump springs. Alex wanted to go one direction with his car's front end and I wanted to go a different direction, however he was still curious to see how a third option would behave during a race. Specifically, we wanted to run a stiffer right-front spring with a smaller sway bar, preloaded higher than we typically would have done with that car. This would have, in theory, caused the right-front bump spring to engage later in the corner, producing a very responsive turn-in (from the high-rate left-front bump spring), with the bump spring shimmed so it engaged approaching full load in the corner, which would stabilize it and produce extra bite when the throttle was applied. The lower initial right-front rate should also have reduced the long-run balance swing from neutral to tight, and reduced tire falloff over the course of the race.

We decided on the Indianapolis race to try it, since it would be a very smooth track with high speeds, and the bump spring engagement should have been very apparent to the driver. I volunteered as "guinea pig" for the test, and I raced the setup in both starts at Indianapolis this year. The setup did produce the results we expected, with a long-run handling consistency that we hadn't seen all year with excellent turn-in and very good drive off. However, the stiffer front springs (and necessary sway bar changes) caused the front end to raise up too high on the long straights, meaning I lost ground to cars ahead in a few passing attempts from the extra drag. The direction Alex went in, however, seemed to work out because he won the race that week. This marked the last speedway chassis spring change we made in the entire year (including rear end adjustments), running every race beyond Indianapolis on the same spring package that Alex used, tailored to each driver that used it.

An article to cover the process more in-depth is on the way in the form of a *Behind the Scenes* article, but it was a prime example of not making spring changes to cure handling, but instead cure chassis balance issues. My car handled exactly the same as my previous setup, it just didn't have the long-run problems I'd faced earlier in the year. Springs are not the "rough adjustment" that we've been told in the past, but instead the foundation on which to build the setup for your car. If you pay attention to what the car is doing, and look at the data it can provide from the tires, you should be able to determine whether or not a spring change is *really* necessary. If not, adjusting something else can save you a world of trouble, not to mention result in a more consistent and responsive car in the future.

Bringing the year to a close, I want to say "Thank you" to everyone who has read these articles and messaged me with their success stories and questions. When I proposed the idea to David Phillips earlier this year, I made it clear that I wanted to help the sim-racing community as best as possible. I hope I've been able to do that so far, and hope to do so in 2017 as well. I wish all of you the best for the holidays, and can't wait to get started on next year!

Commodore's Garage
Article #17 - Adjusting Spring Package

by Matt Holden
January 20th, 2017

Happy New Year! I hope everyone had a wonderful holiday season and, hopefully, everyone has been able to take what we looked at in 2016 and apply it to their sim-racing experience. Before we dive into the more complex components of the cars such as bumpstops, shocks, and alignments, let's take a final look at springs and consider a few aspects we need to consider going forward.

Motion Ratios

We're extremely lucky in sim-racing with the fact that we do not have to bother with suspension motion ratios, however I don't doubt that I may have to eat those words someday. My first experience with suspension motion ratios was when I began R/C car racing in the early 2000s. I raced primarily on a concrete road course, and I couldn't afford a large assortment of springs. However, the chassis I had featured many different ways to position the shock and spring on the suspension. By simply moving the shock in or out on the suspension, I would wind up with a softer or stiffer suspension because of where the shock was located. This allowed me to adjust the individual suspension stiffness without changing the springs and correct the handling for changes such as cooler temperatures or any oils that may have been left on the track from the nitro-powered cars running on the same track.

Having a component of the suspension mounted inboard of the wheel creates what's known as a "Motion Ratio" between the spring and the wheel. This motion ratio produces what is known as the "wheel rate", and will come into consideration when we look at bump stops. For a big-spring stock car such as the NASCAR Cup or Xfinity cars, when the wheel travels one inch, the spring doesn't travel a full inch. As a result, we wind up with a wheel rate that is much softer than the installed spring rate. We could have a 500 lb/in spring in the suspension, but the wheel only needs 300 pounds to move an inch, and thus has a 300 lb/in *wheel rate*.

Currently, there are no cars on the iRacing service that allow for the shocks, springs, or rocker arms to be moved. However, *all* of the cars consider the various motion ratios with the suspension components. For instance, we can have a very simple suspension assembly as in the Mercedes GT3 where springs are pretty much assigned and the relationship between the springs and sway bars can be effectively ignored. Conversely, we can look at the more complex assemblies in the higher-class cars. The NASCAR Sprint Cup car has three acting spring rates for both front corners: Sway bar, Main Spring, Bump Spring. Each of these components is mounted to the suspension arm at a different point, and thus has a different motion ratio with the wheel as well as the other components (The bump spring and the main spring do not compress equally). This is one of the reasons why a 5000lb/in main spring with no bump spring doesn't behave in the same was as a 2000 lb/in main spring and a 3000 lb/in bump spring. On the road side, we can go even more complex with something like the Pro Mazda, with its adjustable rocker arms, or the McLaren MP4-30 and Dallara Indy Cars, that add third springs into the mix at high speeds to control body attitude.

For simpler cars, the motion ratios can be ignored completely since we have them locked at a certain value for all of us. More complex cars need to have this taken into consideration to make sure issues aren't being brought in without welcome.

Weight Transfer

There's a common misconception in the sim-racing community (and the real-world racing community, in some aspects) that changing springs on a corner will change the load applied to the tire as a result of the spring change. This is *not* entirely true, and often leads to some confusing results on the racetrack, real or virtual.

To understand what happens, we need to break the suspension down and look at the parts. For this example, we're going to ignore motion ratios completely and assume that the spring is mounted right at the wheel. Let's say we are seeing 1000 pounds of load on our right-front wheel and we have a 500 lb/in spring. Our spring will be governed by Hooke's Law in this situation:

$F = k*x$
  $F = 1000$ lbs
  $k = 500$ lb/in
  $x = ?$
$x = F/k$
$x = 1000/500$
$x = 2$ inches

Now, let's change that to a 600 lb/in spring on the same car at the same track:

$F = k*x$
  $F = 1000$ lbs

k = 600 lb/in
x = ?
x = F/k
x = 1000/600
x = 1.67 inches

Note that the spring is governed only by three variables: Spring Rate, how much it compresses, and the external force acting upon it. When placed under load, the spring is simply going to compress until it can exert a force equal (and opposite) to the force compressing it. For our example, the external force is 1000 pounds, the springs will compress until they exert 1000 pounds and then the wheel will not travel any further. The spring cannot exert more than is applied to it, and it won't stop travel until it has equaled the external force. It's a very simple principle, but has been misinterpreted in the past. Changing the spring will *not* directly affect the load placed on the tire. The force is an external force governed by the car's speed, track characteristics, roll center and CoG locations, and aerodynamics acting on the car.

What *does* change is how quickly the spring reaches the point of equilibrium. A very stiff spring will reach equilibrium quickly and with a short amount of travel, while a soft spring will allow movement before equaling the external forces. If our spring is stiff and the suspension too rigid, it may not deal with changes in load very well and the tire could lose contact with the track. If the spring is too soft, the driver may feel the car is too unresponsive, since movement in the suspension usually means forces aren't increasing quickly at the contact patch. This is why some drivers prefer a stiffer car and others prefer a softer car, but if we leave everything else the same, the reality is that the tires are loading the same amount. It's just a matter of how you get to that point.

Aerodynamics

Where the springs *will* change the load on a tire is when aerodynamics are considered. For instance, if we look at the MX-5 or the Legend car, there isn't much to be found in the aerodynamics department, and the majority of our gains will be found in weight distribution and alignment. For cars like this, we can leave a set of springs in the car forever and tune the car simply by weight placement. On the other hand, if we look at a car that is heavily dependent on aerodynamics, springs will need to be adjusted to cope with differing levels of downforce.

At a higher speed track, it's much easier to maintain the necessary levels of downforce to keep the car attached to the track. For this, mechanical grip isn't necessary, but maintaining the car's attitude is crucial to high speeds. Thus, stiffer springs will be needed. At a slower track, high downforce levels will be difficult to maintain, so mechanical grip is needed to fill in the space when the driver cannot rely on the air. Whenever we do something that increases the downforce level on a car, we will inevitably need to change the spring to respond. An increase in downforce will increase the external force acting on the wheels, and will in turn increase the spring's deflection. Similarly, if we install a spring with less travel (stiffer rate, for instance), it could raise the car at-speed, reduce downforce, and that tire may lose grip.

Deciding when Spring Changes are Necessary

So with all this considered, when do we look at changing the springs? Most handling issues can be solved without changing the springs, but there are a few cases where it is necessary:

1) **Driver feels the car is inconsistent or unresponsive** – Driver comfort is huge, and is a major factor in fast lap times. If the driver is not confident in the car, he or she is less likely to drive into a corner like their hair is on fire. A car that feels "floppy" to the driver needs to be stiffened up, and a car that feels too rigid will need to be softened. A car that is "loose" or "tight" doesn't necessarily call for a spring adjustment. If the driver feels the car is consistent, it's best to leave the springs and try a weight or shock adjustment, or aero adjustment if you have that at your disposal.

2) **Long run handling swings** – Ever had a car that was awesome for five laps but "flipped a switch" and became terrible? That's typically attributed to a suspension imbalance, where one end of the car wants to go off and do its own thing. This is almost always attributed to a roll stiffness problem, so it's definitely a time where springs need to be changed. A simple check of the tire temperatures should show one tire excessively hot, while the two on the opposite end are rather close. The two tires with similar temperature belong to the softer end of the car. Either stiffen that end or soften the other end, and the balance swing should go away.

3) **Getting outrun between the turns** – Drag will ruin your day pretty quickly, especially at a track where high speeds are a necessity. If another car is running the same speed as you are through the corners, but walks away from you on the straights, it could be that your car is simply producing too much drag. Whether it's too much rake (rear higher than the front) or the front of your car is too high in the air, it's unnecessarily robbing you of a few tenths, so take the time to trim the car out. If there's too much rake, try softening the rear springs to let the rear drop at speed. If the nose is too high, soften the front springs!

4) **Digging up the track** – The only thing worse than aerodynamic drag is drag produced when your car is in contact with the race track. Not only are you going to unload at least one of the tires, but it's going to be like throwing an anchor out of the car. I can't stress how important it is to fix issues with bottoming-out before ever fixing anything else on the car. If you're hitting the track, stiffen up the springs before ever doing anything else to the car.

Conclusion & Bump Stop Introduction

We're going to start looking at implementing bump stops and bump springs in the next few articles. These are going to bring out a level of complexity that can make your head spin as well as breaking a few of the sim-racing setup rules of the past. Springs are not the "rough tuning" option that they've been billed as in the past, and it's become a rarity to see springs being thrown around the garage at race tracks around the world. Understanding that springs are meant to absorb bumps and control body height, not manage weight distribution, is going to be key in grasping bump stop setups. It's not a very difficult thing to understand, but it could cause you to look at things in a whole new way.

Commodore's Garage
Article #18 – Implementing Bump Stops

by Matt Holden
January 27th, 2017

Of all the things available to us in the sim-racing garage, the most confusing thing is typically the simplest component on the car: The bump stop. Ironically, the bump stop is just a tiny version of the one thing that almost everyone understands: springs. The confusion doesn't come from its characteristics, but instead from its operation and how it affects various parameters in the car's behavior. Many cars in the iRacing service use bump stops in some way, and each car's method is unique in some way. Despite that, like everything else, they are all governed by the same rules of physics. The rules for the Mercedes GT3 car will work on the NASCAR Super Late Model.

What is a bump stop?

For a quick refresher, check out the first Commodore's Garage article on bump stops, linked here: http://www.iracing.com/commodores-garage-9-bump-springs-2/ That article is a generalized view of the bump stop/spring component as a whole, and now that we know more about how springs work, it's time to go into a little bit more detail to fully understand them.

To begin, we need to know where they are on the car. In almost every case you'll come across, they're located on the shock absorber regardless of where the main spring is mounted. Some passenger cars have a bump stop molded into the lower suspension arms, but for racing purposes they're always going to be mounted on the shock to make it easier to adjust and/or replace the components to the bump stop.

The first component is obviously the bump stop itself. The most widely used type of bump stop is a composite, foam, or rubber shape that will come into contact with the shock body at some point in the shock's travel around the race track. These can come in a variety of shapes, from the square-sided "puck", the "donut", or even the "Christmas tree" common on some oval short track cars. While the shape is irrelevant for us, it actually changes how the bump stop exerts forces as it is compressed. While a steel coil spring produces a linear force increase through travel, composite and rubber bump stops do not, and instead produce a *rate* increase through travel.

[Graph]

This is an example of the difference in how a bump stop will differ from a spring in its operation. The red line represents a simple 100 lb/in spring. The increase in force is linear, with each step in deflection increasing the force by 100 pounds. The blue line represents how a bump stop would behave, but on a tremendously large scale (There's no bump stop on this planet that will compress 5 inches). Instead of a constant increase in force, we have an exponential increase in force due to a gain in rate through travel. Note that the change from "4" to "5" is the same increase as from "1" to "4". Not all bump stops will follow this exact curve, but it's not very different from what is common in racing. This is extremely important to understand when choosing how to use the bump stop in your car.

The other component to consider is a packer, or a shim. These are used to change where the shock body will contact the bump stop in travel, and thus use the bump stop as part of the suspension. When the shock is not in contact with the bump stop, the bump stop is just extra weight on the car, but having the contact set properly can make a car handle perfectly and having it wrong can make it handle like a dump truck. Packers are very simple, usually just plastic disks that sit on top of (or below) the bump stop.

The combination of both the bump stop and the packers can be referred to as the "stack". Here's an image of a typical bump stop arrangement on a modern shock. I've highlighted the bump stop (consisting of three "donut"-style bump stops), the packers, and the bump stop gap.

[Bump stop image]

What about Bump Springs?

Bump springs are a relatively new invention and have seen a tremendous wave of adoption in oval racing. Since rubber bump stops have the issue of rate change, this can lead to handling inconsistencies around the track. Bump stops at the NASCAR Cup Series level could be extremely stiff, in the tens of thousands of pounds-per-inch, and tiny changes in compression (even 0.05" in some cases) could result in a tremendous change in acting spring rate on the wheel. Imagine going from a 500 lb/in right-front spring to a 6500 lb/in spring without warning in the center of the corner and you have the issue. Not to mention that building the shocks for these very specialized bump stops was getting expensive, and lower-budget teams would struggle to get everything arranged properly in the car. Even at the local short-track level, having a bump stop gap off by a few tenths of an inch could be the difference between a front-runner and a back-marker.
Enter: The bump spring. In real-world applications, these come in two forms, Coil spring and "Belleville" dish-springs, but that doesn't need to be considered for sim-racing. Since these are simple springs, they produce a linear force increase through travel instead of changing rate randomly through travel. This simplifies the physics and mathematics involved and produces an extremely consistent spring rate through the entire range of travel. Right now, only the NASCAR Xfinity and Cup cars use rate-adjustable bump springs, while other cars either use composite/rubber bump stops or the rate is not adjustable.

Here is the same shock as mentioned previously, however with a bump spring substituted for a bump stop. As you can see, it's exactly the same! The only difference is how rate varies through travel. We still have the same packers and the same bump stop gap, and for simplicity's sake, the same bump stop height. The physics and ideas involved with both stay the same: Both add an extra spring to the suspension, both produce an increase in wheel rate, and both can be used to stop downward travel.

What does it do?

[some image here] Initially, bump stops were used to do what their name suggested: stop bump travel. Early uses were often to overcome ride height rules, most prominently during Formula 1's ground-effect period by using soft coil springs to drop the car under aerodynamic load but stop it from dropping further once the car had reached the perfect height for aerodynamics. They've evolved since then, and engineers have discovered they can also be used for acting spring rates as well, most often in cars lacking a "third spring", such as GT cars, sports cars, and oval-track racing.

The need for these arises from a trend toward aerodynamically-based chassis setup. For good aerodynamics, we need stiff springs to keep the car low to the track but not in contact with the track. Sanctioning bodies impose minimum ride heights for whatever reason, and these heights are typically far beyond what is desirable for aerodynamics. Cars need to drop down below these heights at speed, but that requires a soft spring. The bump stop will eventually contact the shock, and now we have the stiff springs necessary for aerodynamics, as well as the low ride height to go along with it.

Mechanically, bump stops behave in the same way a spring does: External forces will compress the bump stop once it has made contact with the shock, the bump stop (or spring) will compress until the force it exerts equals the external force, and the suspension will travel no further. The unique aspect of bump stops is that the wheel rate can change dynamically through travel as the shock comes into contact with the bump stop and releases the bump stop.

For example, let's consider a coil-over vehicle where the spring, shock, and bump spring (for simplicity) all move on the same axis, and thus a 1:1 motion ratio exists between the main spring and the bump spring. Our main spring for this will be 250 lb/in and the bump spring is 1250 lb/in. We've set our bump stop gap to be 1.5" from static ride height, and we'll just keep the 1:1 motion ratio going between the wheel and the spring. For the first 1.5" of travel in the suspension, the wheel rate will be 250 lb/in. Once we come into contact with the bump stop, we now have to add the two springs together, and the wheel rate jumps to 1500 lb/in for any travel beyond that point. This is a major jump in wheel rate, and definitely something that a driver would notice in the car. Compression is no problem, but an improper arrangement of sway bar and shock could cause the shock to contact and "release" from the bump stop multiple times through a single corner, especially on a bumpy surface. Coming on and off the bump stop many times in a corner results in multiple changes in spring rate *through the corner*, and there's never a good ending for that situation.

Packer Settings

The packers are a tremendously important part of the entire bump stop stack and I could do an entire article on them alone. As the car compresses and gets to where the shocks contact the bump stop, we have a lot of things that are happening. Sway bar preload and spring forces at this point, just prior to bump stop contact, will result in a "dynamic" crossweight that is likely different from the static crossweight. Once the bump stops are in contact, the crossweight can change yet again, and this is a problem. If the packers aren't set properly, then bump stop contact can send the dynamic crossweight in one direction very quickly, which is often perceived as an issue stemming from spring rates or sway bar settings. For instance, if the right-front bump stop contacts earlier than the left-front, there will be a sudden increase in dynamic crossweight, or vice versa if the left-front bump stop contacts first. Due to the high rates used in bump stops, it's unlikely that the second bump stop will be able to "catch up" to the first one.

Timing is everything when using bump stops. If the shims and bump stop gap aren't set properly, it's very easy to get bogged down looking for a problem that you're never going to solve. Most of the cars on iRacing have symmetrical suspension systems, so if the packers are set equally on both sides, there's a good chance the bump stops will contact at the same time. The NASCAR Xfinity and Cup cars are asymmetrical, however, with the right-front shock usually slightly longer than the left-front shock. As a result, it's good practice to keep the shims at the same height for rough tuning and for clearing any contact with the race track, and adjusting one shim by a small amount at a time to tune the car's handling once on the bump stops. Usually a difference of 0.125" between each bump stop stack is safe and will produce a consistent behavior without major headaches and issues coming your way.

Special Cases: V8 Supercar & the NASCAR Cup Car

The V8 Supercars got an overhaul in the last build that introduced Damper Length and Bump Stop Gap to the setup screen. Damper Length itself isn't a new concept to the iRacing setup garage as a few cars have had this option for a while. The V8 Supercar adds in the complexity of both a shock collar offset and the bump stop gap, and the three have to be adjusted at the same time to change the car's handling. If the V8 Supercar is your wheel house (If it isn't, try it, it's fantastic), the only way to understand this is going to be fiddling with it yourself. The proper way would be to adjust the damper length, then adjust the shock collar to reset to the previous ride height. If you did it right, you should have the same ride heights with a new bump stop gap!

The other special case is the NASCAR Cup Car. Since the car doesn't feature ride height minimums, the concept of the bump stop becomes more complex. There's no need to "drop" the car at speed, so the reasoning for a bump stop/spring changes. In the future, we'll go over what this adds to the mix and the various ways you can go about setting that specific car up.

Commodore's Garage
Article #19 – Crossweight

by Matt Holden
January 27th, 2017

Last week we looked at bump stops, but this week we're going to cover weight distribution, specifically *crossweight*. Why would we move straight from springs to crossweight? When we dive deeper into bump stops (specifically shims, gap, and contact timing), preloads, and shocks, crossweight is going to come into play *a lot*. And I mean *a LOT*. Crossweight is a very simple concept but, along with many other things, past racing sims have over-simplified it to the point where you can run into problems in the iRacing garage. Racing is about details, and making sure all the details are correct before you ever hit the track. One of the most important details is the weight distribution because it plays so heavily into handling. An error in how the crossweight was set can make or break your race, even if everything else on the car is perfect.

What is Crossweight?

Crossweight is the percentage of weight in the car situated on the right-front and left-rear tires. The percentage is calculated by adding the weight on these two tires and dividing that value by the total weight of the car (See diagram). While this is commonly kept at or around 50% for road racing cars, oval racing cars can vary wildly from car to car within the same class of racing. Values above 50% can be thought of as "positive" crossweight, values below 50% can be thought of as "negative" crossweight.

[crossweight image]

The way crossweight is represented for a given series or type of car can be wildly different. Indy Cars and Dirt Oval cars are the most dramatic departure from the common crossweight percentage value. Both of these cars typically represent crossweight as a value of weight for a specific tire on the car. For Indy Cars, this value is the amount of weight on the left-front relative to the right-front. iRacing's garage for the DW12 represents crossweight as "*xx* lbs to the left front". For instance, if that value read as "-100 lbs to the left front", then the left-front tire has 100 pounds less on it than the right-front. Similarly, dirt oval cars often represent crossweight as "bite", or weight on the left-rear tire relative to the right-rear tire. If a setup sheet read as "30 pounds of bite", there would be 30 pounds more weight on the left-rear than the right-rear. These are both ways to represent the same thing, and a crossweight percentage can still be derived from the individual corner weights if desired.

Effects of Crossweight

While the normal "more crossweight = tighter" is a guideline, it is by no means a law. Crossweight, by itself, is not an indication of whether or not a car will be tight or loose and should not be treated as such. Crossweight, instead, is used in conjunction with spring rates, nose weight, left-side weight, and track characteristics. Typically, a higher nose weight will need a lower crossweight to keep the car going around the track, while a lower nose weight will need a higher crossweight. Similarly, cars with higher left-side weights can get away with a higher crossweight.

Crossweight is often a key factor in forward traction, however, and is usually tailored to a track's characteristics. Specifically, bumpiness and grip level offered by the surface itself. Since crossweight adds weight to the left-rear tire, this provides more traction while on throttle (which is why Dirt folks call it "bite"). Bumpy, low grip, tracks such as Five Flags Speedway or Atlanta Motor Speedway don't offer large amounts of grip, so crossweight is set higher at these tracks to make the car a little more stable over bumps and while the car is sliding on the worn out surfaces. Newer tracks, or very smooth tracks, such as Kansas or Southern National, allow for lower crossweight to be used since the track has enough grip to provide traction without help from the car. Diving further into track characteristics, higher banking will allow for lower crossweight due to the extra grip from vertical loading, while flat tracks such as Langley will have higher crossweight to provide extra bite off the flat corners.

One of the biggest factors to determining crossweight is spring "split", or difference in spring rates left-to-right on the car, which is covered later in this article.

Adjusting Crossweight

Before setting the initial crossweight in the car, it's necessary to know how to adjust the spring perches consistently around the car, in close to equal values, to keep everything constant. If one perch is adjusted, then a second one *must* be adjusted as well. Proper crossweight settings are made by adjusting all four, and doing so will keep ride heights consistent while changing *only* crossweight. To decrease crossweight, we need to use the following adjustment buttons in the iRacing garage:

[Decrease CW]

When using these four buttons, make sure the ride heights stay about the same for each trip around the car. Different spring values and heights will result in different amounts of weight for a single click (a 500lb/in spring will have a different result per click than a 1000 lb/in spring). Whenever you're close to the desired crossweight value, adjust only one end of the car instead of the *entire* car. Typically, the rear of the car is easier to work with than the front, and will prevent any issues with aerodynamics coming into play.

To increase crossweight, we simply use the other set of buttons:

[increase CW]

Similarly, we need to keep the ride heights constant, and switch to the back of the car for finer adjustment.

Setting Crossweight

If our car is equipped with sway bars of any kind, it will have two crossweight values throughout the setup process: Bar attached and Bar detached. Since preloading a sway bar will directly affect crossweight as well as spring preloads, it's very important to know (and set) initial crossweight with the bar(s) detached while knowing how much crossweight is desired once the bar is attached. Doing so will eliminate issues that can be masked by preloads introduced during the setup process. For instance in a NASCAR-style stock car: if we set the ride heights, springs, and crossweight with the bar attached, it will preload the sway bar and the rear axle housing. Once we have what we're looking for with that, we'll then have to reduce preloads on the bar and axle housing to where we want them, which will change the crossweight again. This will result in a see-saw of crossweight and preload adjustments until we have the numbers we're looking for. However, more likely than not, the spring deflections, bump stop gaps, and alignments will all be completely wrong, causing the car to drive very poorly and abuse the tires through a lengthy run.

[Detached image]

Above is an image of the iRacing setup garage with only the necessary items for setting crossweight. Here we can see the crossweight set initially at 49.5% with the bar detached. When the bar is attached, however, we get this:

[Attached image]

We now have an increase in crossweight of 4% and large changes in corner weight values, but no adjustments made to the spring perches themselves. This is *not* bad, and is going to happen in every single car, but it's an example of what needs to be considered during the setup process. Whenever you start to build a setup, you should already have an idea of how much preload you want in the bar as well as a general idea of what you want the static crossweight to be. In this case pictured, I wanted 200 ft-lbs of preload and a crossweight of 53.5%. I managed to get both remarkably close, and I have everything set where I want it to be set.

Doing this is fairly simple, but involves *paying close attention* to everything in the garage screen while you do it:

-Start by detaching the sway bar and allowing the car to settle. At this point you should already have the springs you want in the car and have set the ride heights where you want them.
-Using the Spring Perch Offset values, adjust the crossweight to a standard value. Something like 50% will make this very simple. Adjusting the crossweight using the four perches to keep ride height constant is very important! (See above images) Record this starting crossweight value.
-Next, attach the bar and set the preload value you want. Allow the car to settle, and note the new crossweight value as well as ride height differences (How much did the left-side raise with preload?).
-Find the difference in the preloaded crossweight value and the initial crossweight value, which will give you the crossweight added from the bar's preload.
-Now that you know how much crossweight changes, you can subtract that from your desired crossweight to get your "detached crossweight". Detach the bar and adjust the crossweight to this detached value, attach the bar, preload it, and you should be at your desired crossweight.

So…I know that's confusing, so here are the same steps with my example values to give a visual:

-Set the car at 50% crossweight with the bar detached (our desired CW is 53.5% when complete).
-Attach the sway bar and preload it to 200 ft-lbs, allow the car to settle.
-The preload increased the CW by 4%, so now we'll detach the bar and allow the car to settle again.
-Since we want 53.5% and the bar added 4%, we'll need to start the car at 49.5%. Detach the bar, let the car settle, and adjust the crossweight to remove 0.5% of cross from the car. It should settle at 49.5%.
-Re-attach the bar, preload it to ~200 ft-lbs, and the crossweight should now be at or around 53.5%

Head still spinning? That's okay! I've had requests for videos to tie in to these articles, and this is going to be the first video to help explain how to do this.

Dynamic Crossweight

It's important to understand that crossweight, unlike nose and left-side bias, is (typically) not adjusted by moving physical weights, or ballast, around in the car. Instead, a combination of spring preload (static compression force) and twisting preload on the sway bar and axle housing cause forces in the suspension to be redistributed, resulting in one corner of the car's suspension to support more weight than another corner. Since there is no physical object to associate the crossweight with in the car, its value is purely dynamic and will change as the car moves through travel, through steering inputs, and as the suspension encounters bumps and track features.

This dynamic crossweight will be different based on not only springs and shocks, but also driving styles. I mentioned earlier that crossweight is not indicative of a car's handling characteristics by itself, and that I'd give an example:

In 2016, my team began racing in the NASCAR iRacing Class B Series full-time. We quickly discovered that, with 3-5 drivers per week, a single setup was only going to favor one of the drivers while hurting the others. We began tailoring individual setups per driver with different springs, shocks, and weight distributions. By mid-season, all of the drivers were running up front despite all the cars being different. There was a phrase we had: "If we have five cars on track, we have five different setups on track."

[alex, me, homestead]

Of all the cars we built for that season, none were more different than the cars Alex Scribner and I drove. Alex's driving style differed from mine so much that we could not race each other's cars save for Iowa, where we ran identical cars. All other tracks were totally different. The major issue was corner entry, where Alex's steering input and line often resulted in more load on the right-rear, while my line produced a heavier load on the right-front. Since Nick Ottinger and I share a very similar issue, we already knew how to handle it: larger rear spring split and higher crossweight for my cars. Where Alex ran anywhere from 400-600 lb/in of rear spring split, I ran between 800 and 1000 lb/in of spring split in the rear. This difference alone meant that my crossweight was often 2-4% higher than Alex's cars due to the softer left-rear spring.

This led to something very counter-intuitive: Despite Alex's cars having a lower crossweight, they were much tighter than my cars when I drove them. Similarly, Alex always said my cars were extremely loose, despite a higher crossweight. In fact, *everyone* said my cars were very loose even though I ran the highest crossweight of anybody in that season. This is a prime example of how crossweight does not directly indicate a car's handling characteristics and must be treated as part of the entire setup package. Follow the steps I've shown here and you'll wind up with a much more consistent setup process as well as eliminating nagging issues that would otherwise cause headaches for weeks.

Commodore's Garage
Article #20 – Bumpstops & Crossweight
by Matt Holden
January 27th, 2017

Following our look at the details pertaining to crossweight last week, we'll now turn our focus to a situation unique to cars with bump stop suspension systems. This is the first concept we'll look at that will not apply to every car in some way, but while your interests may not fall into a category with one of the cars in question, that doesn't mean you won't be driving one in the future! We're going to look at what happens when bump stops *engage*, or come into contact with the suspension components and become an acting part of the suspension system. We're also going to look at bump stop gap, as well as packer settings, and how this affects what happens at the point of engagement. Warning: picture overload, but this is something that may be necessary.

Acting Spring Rate

The major reason for using bump stops (or bump springs) at all is to modify the suspension's acting spring rate through travel. In most cases this is to control heavy aerodynamic loads and prevent the car from coming into contact with the racing surface, but it also has the potential to drastically change the car's handling characteristics if bump stop contact is not tuned properly. For simplicity's sake we're going to eliminate rear bump stops in our example, but the principles that apply to the front suspension can also be applied to rear bump stops.

The term "acting spring rate" refers to the total spring rate acting on the suspension. In a single-spring suspension, the acting spring rate it the rate of the installed spring. Since we have two springs (ignoring the sway bar) that will act on each corner of the suspension, that act in parallel, the rates of the two springs will combine to create the acting spring rate on a bump stop suspension system. For these examples we'll assume the following: we have a 500 lb/in "main" coil spring installed and a bump spring with a linear 1000 lb/in rate. (NOTE: For a car with composite bump rubbers, such as the Super Late Model or the GT3 road cars, the same ideas apply, but they will not have a linear spring rate through travel.) Our car will be arranged like this:

[initial]

Here, we have a symmetrical spring setup and for initial travel our bump springs will *not* be engaged (bump stops are crossed out). This will remove the bump springs from consideration and our *acting* spring rate is 500 lb/in. At low speeds with little aerodynamic load this will be our operating conditions and the bump springs can be ignored completely. The front end will have a very low roll stiffness relative to full-load conditions because of the low acting spring rates. Let's now apply full loading from aero and the track characteristics and engage the bump springs:

[doublestop]

Here, we've compressed the suspension to the point of bump spring contact and we now have new acting spring rates. We'll add the spring rates together to get our new acting spring rate of 1500 lb/in on each corner. This new spring rate will keep the car off of the track at high speed, as well as produce a very responsive front end instead of the more dull response likely with the lower spring rates. This also increases the roll stiffness of the front end, dynamically altering the balance of the car.

It's important to note that this acting spring rate change happens *instantly*. The spring rate for a suspension system changes at the moment of contact, and will similarly change again at the moment of release, where the suspension is no longer in contact with the bump spring or stop. If the bump stop is not in contact consistently through a turn (the suspension is on and off the bump stop repeatedly), the acting spring rate can similarly change many times during a corner and produce a very unpredictable car.

For those of you interested in the math portion of the bump spring engagement, it's a common mistake to only consider the acting spring and relate that with suspension travel. Here's another of my handy charts to go along with our example:

[chart]

Here, we start with a 500 lb/in spring rate until around 1.25" of total suspension travel. We're eliminating the motion ratio for both springs, so at 1.25" of travel the main spring exerts 625 pounds of force. At any point beyond 1.25" of suspension travel, we will be engaged with the 1000 lb/in and have an acting spring rate of 1500 lb/in. However, if you notice from the chart, 1.5" of travel results in only 1000 pounds of force, *not* 2250 pounds as would be the case for a 1500 lb/in spring! Were we to run a 1500 lb/in main spring and no bump stop, we'd actually see less travel and need to start the car lower due to the higher force with shorter suspension travel.

The missing 1250 pounds is found when the springs are looked at separately at this point in the suspension travel. At 1.5" of suspension travel we have compressed the 500 lb/in spring 1.5" and it is exerting 750 pounds of force. At this same point, we've only compressed the bump spring 0.25" since it engaged at 1.25" of suspension travel, and it is exerting only 250 pounds of force. This is why having a main spring of 500 lb/in and a bump spring of 1000 lb/in *does not* produce the same results as a main spring of 1500 lb/in by itself. At the point of contact, the force acting on the suspension will increase at the rate of the *acting* spring rate but will not be the same force as a single spring with the same rate.

Uneven Bumpstop Contact

Something that can prove to be both an advantage and a headache is whenever the bump stop gaps are set unevenly and contact occurs sooner on one corner than the other. In most cases, the car will be set up to contact one bump stop before the other anyway, especially in oval racing, but if it's set improperly this can cause unpredictable handling and it can very likely lead to a loss of control. Let's look at our example again, but mis-match the bump stop gaps.

We're going to start again with the bump springs disengaged and an acting spring rate of 500 lb/in on both corners. First, we'll set the right-front bump stop gap very small relative to the left-front gap and see what happens:

[RF stop]

Here, we've engaged the right-front bump spring but the car is still riding on only the left-front main spring, so (instantaneously, remember) we've moved to a 1500 lb/in acting spring rate on the right-front while the left-front didn't change. If the bump stop rate isn't that high, we can get away with this for short amounts of travel, however if there is a large amount of suspension travel this can be a problem. Let's break out the chart again!

[RF Only – chart]

With this chart, I've added a trace (in red) to represent the left-front corner. We can see at the point of bumpstop engagement (1.25" of travel) the force on the right-front (blue trace) begins to gain force much quicker than the left-front due to the higher acting spring rate.

[Late LF engagement]

If we eventually engaged the left-front bump spring, and the bump springs were the same rate, we'd wind up with something like this where both corners are gaining force at the same rate. However the right-front will *always* have more force acting on the suspension at all values of travel in this configuration.

Crossweight Implications

If the suspension travel were 100% in the vertical direction, we'd be okay in terms of weight distribution. However, we do have to consider chassis roll when the acting spring rates change and what that will do to the sway bar and rear end housing (if the car has one). Whenever one bump stop contacts before the other, the suspension will go through a change in roll stiffness and the chassis will inevitably begin a shift in dynamic crossweight until the other bump stop is engaged. Following our same example, but looking at the front of the car, we'll wind up with this:

[front end RF bumpstop]

Here, we've engaged the right-front bump stop earlier than the left-front bump stop again, which will start to cause the car to want to roll to the left (from the driver's perspective). This situation will likely result in an increase in crossweight, plus the force gain from the higher acting right-front spring rate. Once the left-front bump stop engages, the dynamic crossweight will probably decrease, but the sway bar and rear housing could still be fighting the initial roll change. This is often a very difficult problem to diagnose from the driver's seat, mainly because the effect is very small initially. Since bump stop gap differences can be fractions of an inch before motion ratios are considered, the shift in crossweight can be relatively small. Still, it can manifest into a larger issue later in a run, resulting in poor tire wear and/or uneven tire pressure buildup.

The real-world method to avoiding this problem is simply to remove the main springs and allow the car to sit on the bump stops at the heights it would be under full compression. Shims are then inserted or removed on the bump stop stack until the crossweight at full compression is equal to the crossweight at static heights. As of now, we don't have that ability in iRacing, but we can never rule that out as a future possibility.

Where do I start, then?

The best way to start setting shim height and bump stop gap is to keep them equal on both front corners, as well as equal front spring rates for both main and bump stop rates. Testing runs should show whether or not one of the front tires get excessively hot over a few laps, and you can either add a shim to the other suspension (RF is hot, shim LF) or remove one from the hot corner. In most cases, the ideal difference in shim height/bump stop gap will be small but it is something else that is driver-specific. Once the difference between the bump stop gaps has been determined, adjusting them equally up and down can be used to fine-tune splitter or valence height at speed. Should you decide to change spring rates at any point, you will need to change the bump stop gap/shim height to compensate for the change in suspension travel. Even further, any change in ride heights will change the shock deflections, and thus the bump stop gap. *Always* keep an eye on the shock deflections in a bump stop car to make sure your bump stop gap is staying constant through all of the adjustments. If you change something on the car that changes the bump stop gap, but you do not account for it, you will have changed more than one thing on the car!

Bump stops are not simple at all, and it's often the main reason why some racing teams just never have a competitive car. Adding a new spring in the middle of the suspension's travel is a mathematical nightmare, not to mention all of the things that carry a spring rate with them in the suspension. If we looked at a real example, we'd need the motion ratios between the main spring, bump spring, and sway bar, plus the ride spring rate of the tire itself, as well as any extra springiness that could come from the suspension arms! There is a lot to consider with these things, and they're not something that can be understood overnight, but taking the information in a little bit at a time and testing each little bit as you go will work wonders in the long run.

Since this is article #20, there will be no new Commodore's Garage next Friday. We'll be back on March 3rd where we'll dive into some telemetry analysis and start applying what we've learned so far. I want to wish the best of luck to everyone who is running in the three NASCAR iRacing Series that start up next week, and I hope to see you on the track!

Commodore's Garage
Article #21 – Telemetry

by Matt Holden
March 3rd, 2017

The most valuable thing to tuning a car is data. Whether it's tire temperatures, pressures, wind tunnel, or telemetry data, anything that can tell you what the car is doing at any given moment can be worth its weight in gold. Despite a visual simplicity, telemetry data is a rabbit-hole of complexity with almost limitless possibilities that can tell you what you need to know, and even things you don't really need to know (but are fun to talk about). Telemetry and data analysis is what I get asked about more than anything else, and many are surprised to find out how little we actually use it at Gale Force Racing across every series we may run. For some reason, there's an idea that telemetry will tell you how to adjust the car which is far from the truth, but it's extremely valuable for diagnosis purposes, not to mention its value when comparing drivers to each other.

For this article I will be using screenshots from MoTeC i2 Pro version 1.0.21. This is an older version, but it's where my current workbook lives until I can get a new one built for the newer versions of MoTeC. Should you be a fan of the Atlas telemetry software, the same ideas and math will still apply, it might just look a little different.

The 2014 NASCAR PEAK Antifreeze Series Season

Before we dive off into the meat 'n potatoes of telemetry, we need to take a look at where (and how) telemetry can get you in trouble. Telemetry is a method of data analysis and is simply a way to present a large variety of data from a race car at any given point or situation on the track. Therefore, telemetry is, by definition, a way for the car to communicate what it is doing on track to you or whoever is doing the setup work for your car. It is *not*, in any way, a way to show off issues with the car. During the 2014 season in the NASCAR PEAK Antifreeze Series, we fell into that trap with a few of our cars at Gale Force Racing. With the ride height rules removed from the garage, we focused heavily on telemetry and tire data to tune the cars, giving the drivers virtually no say in how the cars were built.

[indy 2015]

As you can probably imagine, this became a serious problem during the year. Instead of listening to the drivers (who are experiencing what the car is doing) we focused completely on telemetry to adjust the cars. This produced some extremely ill-handling cars and issues during the races where adjustments would produce no results or would produce the wrong result. Yes, Nick Ottinger won two races that year, but would you be surprised to learn that those races were the two races where telemetry was tossed and Nick built his own car off of feel? For 2015, we scrapped our "Telemetry-first" methods and while Nick still only won two races, the cars were much more consistent since we were adjusting based on what he was saying about the cars. In fact, everyone did better in 2015 in terms of on-track performance.

If we move a year further on from that to 2016, we have Alex Scribner's Class B championship. Believe it or not, it was developed on a relatively low amount of telemetry analysis, with the vast majority of it coming in April during the spring break. For a given week, we would have two or three 1-hour practice sessions, and telemetry may have been running for 15 minutes total during the week. I don't mean to say that telemetry should never be used, but it's very easy to work yourself into a mountain of problems if you aren't careful!

Know what you're looking for

The most important thing about data analysis is to know what you're looking for before you ever open up a log file. If you go into your data analysis program without any idea of what you need, it's like trying to cook a meal without knowing what you're going to cook. Whenever you're building a brand new setup, you need to know ahead of time where you want the ride heights to be at speed, what your temperature spreads across the tires should roughly look like, and whether or not you will be seeing a lot of suspension oscillations (bumpy track) or if the suspension needs to be locked down (smooth track). Basically, whenever you build a brand-new setup, you're going to be using telemetry to check and make sure you've gotten everything correct before you ever start adjusting on things.

For example, Alex used a brand new setup in the second half of the 2016 Class B season. The setup works so well for him that he still uses it today, three races into the 2017 season. His setup was built primarily off of the April test I mentioned earlier, during which Alex and I adjusted every single component on the car from a control setup and observed what happened to the car as a result. Once we knew what we needed about the car, we got the aero data he liked from our "older" setup and then built the new one with that in mind.

[averages]

Here's an image from one of my most heavily-used telemetry channels: Height averages. Testing showed that each of our drivers liked a specific attitude for the car in terms of rake. Some like a very high rake, some like a very low rake, but it was always constant across all tracks. This box overlays both an average of the front ride heights and rear ride heights, giving a visual of the car's rake angle which is crucial in terms of tuning the car's drag-vs-downforce characteristics. You could further generate a rake angle by creating a difference channel (Rear AVG – Front Avg = Rake, inches) if you wanted to, but I use these two channels to make sure the aerodynamics are correct for me at a new track. If I make a spring change, it's a simple case of double-checking these traces to make sure nothing has changed aerodynamically in the turns, which can mask the spring change. Once that's checked, we'll basically turn telemetry off and go from the garage temperatures alone.

Temperatures & Dampers

Tire temperatures can also be helpful in diagnosing handling issues. However, it's important to note that the telemetry temperatures are recorded differently than the garage temperature readouts. Essentially, the garage temperature readouts are taken from the tire carcass in the same way you'd take the temperatures using a pyrometer in the real world. This device is simply a probe that sticks into the tire tread and records the temperature, and is very useful for getting an idea of how the tire has been loaded throughout a run. The telemetry readout, on the other hand, can be thought of as an infrared temperature sensor, getting the instantaneous temperature reading at any given moment on the tire. These figures will usually show a much more drastic change in temperature throughout a lap when compared to the garage temperatures, and can be used to identify a tire that is getting loaded too quickly which wouldn't show up in the garage tire sheet.

[temp 4-view]

Another channel to consider and pay attention to is the Damper Position (or Suspension Position) channels, as these can be directly tied to wheel movement in the car. If they aren't moving, the wheels aren't moving, usually a sign of either coil-binding or the suspension being very stiff at that given moment. Typically, a coil-bind trace will be flat while a stiff suspension trace will have movement around a given position. If this is what you're going for, then it's fantastic, but if not….well that's a problem! Here are pictures of a coil-bind trace and a high-rate bumpstop damper trace:

[cb trace]

[bumper trace]

Math Channels

Possibly the most powerful telemetry tool that is the most unknown feature is the Math channel. In many cases, engineers can only put so many sensors on a car before it becomes a problem, but a lot of information can be gathered from small amounts of data input. Enter: The math channel. In most telemetry programs, there is a feature that allows you to create your own mathematical formulas that generate a new channel that isn't directly provided from the sensors given. The Ride Height Average channels mentioned before are math channels based on the following equations:

- *Front Height AVG = (LF Ride Height + RF Ride Height)/2*

- *Rear Height AVG = (LR Ride Height + RR Ride Height)/2*

Neither of these channels could be developed straight from any sensors, but with a simple mathematical entry we can get a new piece of data that is extremely helpful in tuning. My personal telemetry workbook has 47 unique math channels (Some are duplicates like the Averages), here are a few of the ones we use the most:

- *Wheel Lockup: Wheel Speed RR + 20 < Wheel Speed RL*
  - This channel exists for all four wheels and compares a wheel with its opposite to determine a lockup. If the wheel speed of one wheel is 20mph lower than the other, data will output this as a "locked" wheel.

- *Acceleration: derivative(Ground speed)*
  - A derivative is a calculus expression to determine the rate of change of a graph. The derivative of position is Speed (rate of change of position), and the derivative of Speed is Acceleration (rate of change of Speed). This simple channel can be used to find when an engine has "run out" of power on a long straight, or to help identify the best shift points on road courses. Not perfect, but it has proven its worth in the past.

- *Single-side Pressure Differences: LF Pressure - LR pressure (Also for right-sides)*
  - When tuning, especially with springs, we want the pressures to stay fairly close front-to-rear to prevent the tire's spring rate from influencing the car's behavior. With this channel we can see when the pressures start to deviate and whether or not a driver's behavior will play a role in producing these changes.

Driver Comparisions

The final thing we'll look at is driver comparisons. If you have a new driver, there is nothing more valuable than knowing how that driver differs from the drivers you're familiar with. We'll look primarily and the driver's inputs, usually the throttle and brake, overlaid with the car's speed, to see any major differences.

[inputs]

Here we can see two very different displays of drivers on the same track. For reference, the Color traces were slightly slower than the White traces. Most notable is the throttle behavior and how that influenced the car's speed on track. The faster driver drove further into turn one and was back on the throttle very quickly, plus he didn't hesitate on the throttle as the slower driver did. There's a similar behavior at the other end of the track as well, with the faster driver being much more confident on throttle application. It should be worth noting that these two drivers were on the exact same setup.

Small things like the differences in throttle application and how far the faster driver went into the corner can lay an excellent foundation for setup changes for a newer driver. In this case, the color traces belonged to a newer driver, and his lower corner entry speed and much calmer throttle application allow him to run softer front springs due to the lower vertical loads on entry, as well as a much more forgiving car on corner exit due to smoother throttle application. The driver with the white traces needs a much stiffer car to cope with a more aggressive throttle application and corner entry behavior. With this information known, both drivers have cars that better suit their own driving styles, and speeds are similar between the two.

Final Thoughts

When we finish going over the necessary components to build a setup and we've begun the building process, telemetry will come into play. However, its use beyond initial build won't be as extensive as some may expect. While it is invaluable through a testing process, it should never be used as a guide during the tuning stages, but instead as a way to confirm what you or the driver is saying about the car.

Commodore's Garage
Article #22 – Shocks

by Matt Holden
March 24th, 2017

If you've watched enough racing, there's a good chance you've heard at least one team mention that they've got a new "shock package". There's also a good chance that team was either significantly faster or far more terrible than they were on the old shock package. Shocks have that power in bunches, and a good understanding of how the shocks are affecting a race car can often make the difference between a good day and a "let's just go home" kind of day. While they appear very simplistic from the outside, a shock's internal operation can be tremendously complex, leading most teams to have a dedicated shock specialist for each car with the knowledge of how to adjust and rebuild shocks on the fly. Cale Gale, founder of Gale Force Racing, did just that for years in NASCAR's Xfinity and Camping World Truck Series before turning his focus to the highly-competitive world of short track oval racing. I can attest to the sheer complexity of racing shocks thanks to my years working for US Legend Cars International. When I started working there, I knew almost nothing about shocks, but my boss was nice enough to let me sit in the Bilstein trailer at Charlotte Motor Speedway every now and then and watch the shock guys do their thing.

Shock Operation

Before ever diving into how shock adjustments work, we need to understand what's going on inside the shock body. To the right is a cutaway image of a very simple shock with all of the major components labeled. For most racing shocks, we have the shock body separated into two chambers: The fluid chamber and the gas chamber. The fluid chamber is filled with some kind of oil that the piston passes through while the suspension is in travel. The second chamber holds a pressurized gas that can help prevent bubbles forming in the fluid, which is really bad, mmkay?

Our main focus is going to be on the piston itself. Pistons are typically just metal disks with a specific pattern of holes drilled through to allow the fluid to pass through. One side of the piston will be known as the compression side, the other the rebound side, and the holes can vary in design for each side of a given piston to get the desired results. Finally, on either side of the piston is a stack of extremely thin metal shims. When I say "thin", I mean *thin.* Penske Racing Shocks' website sells a shim kit with thicknesses going from 0.020" to 0.004" (0.5mm to 0.1mm) with varying diameters. The arrangement of these shims controls the "high-speed" characteristics of the shock.

For adjustment purposes, most teams will consider only two situations for the shock: High-Speed and Low-Speed. Each speed setting will be considered in both compression (shock is getting shorter) and rebound (shock is getting longer), leading to a maximum of four adjustment options per shock: High-speed compression, Low-speed compression, High-speed rebound, and Low-speed rebound. Low-speed adjustments can be made in a variety of ways, but the most common is usually a needle in the shock shaft itself that opens or closes a "bleed" hole in the shaft. Closing the needle will result in a smaller passage for the shock oil, stiffening the shock. High-speed adjustments are usually made through the shim stack (and its preload) or the piston itself. At a given shock speed, the shims will buckle and open the larger holes in the piston. This allows more fluid through the piston and instantly changes the shock's stiffness without external adjustments necessary.

Adjustment types

A shock's complexity has led to many different types of adjustments available to teams, and the iRacing garage has followed that trend. Most of the shocks in the garage will have a clicker-type adjuster. This is simply a knob on the shock that adjusts the needle mentioned earlier, and literally clicks as it is turned. "Sweeps" is another type of adjustment, and is made in a similar fashion to the clicker adjuster, just without the clicks. All adjustment types will do the same thing, and in all cases a higher number is a higher stiffness for that given adjustment.

Low-Speed Adjustments

The majority of your handling-related shock adjustments will be on the low-speed side of the shock. For cars with only one shock adjustment (oval cars, some road cars), the garage adjustment is the low-speed adjustment. Even though we don't have the vast adjustment range that a real-world team may have at their disposal, the shocks can still be used to effectively adjust for weather conditions or track state changes. While I don't like to use the old "setup matrix" principle because of how quickly you can build yourself into a hole with them, the expression "Springs for rough tuning, shocks for fine tuning" *does* apply in this sense. If we think of track characteristics first, such as how bumpy the track is, how much banking the turns may have, or how important aerodynamics are, the springs will be the rough-tuning factor of our setup. Shocks will come into play once our springs are correct, and when we're trying to get rid of issues related to suspension oscillations or weather conditions. (NOTE: *NEVER,* and I mean *never*, change a spring to adjust for weather conditions! There are other options that are more effective and less problematic to choose from.)

Starting with compression, the low-speed adjustment is in play whenever a wheel is traveling upwards into the car, with the shock getting shorter. This can be under braking, initial turn-in, banking increase, or when throttle is applied. When compression is increased on a given corner, the shock will exert more force on that specific corner in compression, placing more force on the tire while the shock is moving. Aerodynamically, this can be used to hold a corner of the car up for a longer amount of time during handling transition periods, such as during braking or throttle application. An example of how this can be used is at a high-banked, high-speed oval such as Bristol where the right-front suspension will experience high loads and travel very quickly. If the driver is experiencing oversteer on corner entry, an increase in right-front compression can apply more force to that tire and reduce oversteer on entry. Similarly, oversteer on the same corner's exit can be reduced by lowering the right-rear compression and letting the car squat more onto the right-rear tire.

Low-speed Rebound changes the shock in expansion. This adjustment, unlike compression, does not directly affect the tire it is associated with, and instead affects the two tires adjacent to its corner. For example, if we adjust the RF rebound, we'll see a dynamic load change on the LF and RR tires, with very little effect on the RF or LR tires. Rebound adjustments can be simplified and thought of as a see-saw (pictured below), with the adjusted corner being the low side of the see-saw (we'll adjust the RF rebound here). In order to pull up on the RF corner, the low side of the see-saw, we need to push down on the high side of the see saw, which will put a very heavy load on the fulcrum of the see-saw, in this case the LF and RR corners of the car. This is a *highly* simplified example, but it's good enough to explain what's going on with the car.

That explanation out of the way, the low-speed rebound settings are great ways to adjust the car's dynamic crossweight while cornering. Increasing the LR rebound can decrease the dynamic crossweight at corner entry (more load on LF and RR) and help turn-in. Similarly, increasing RF rebound can decrease dynamic crossweight on throttle and help rotate the car on corner exit. Should you wish to tighten the car in these conditions, increasing rebound on the RR can help tighten up corner entry, while more LF rebound can tighten up exit. Below are two diagrams showing force changes on the wheels due to rebound adjustments.

High-Speed Adjustments

At a specific shock shaft speed, usually 3.0 inches-per-second or around that, the shock will switch from low-speed to high-speed conditions, and the damping forces will change dramatically. Most of the road racing cars now have this high-speed adjustment functionality, and tuning the high-speed settings can make an unpredictable car a little more predictable in most cases.

The major thing to making high-speed adjustments is knowing when the suspension is in this high-speed condition. If the shock is moving at 1.5 in/s, the high speed adjustment will have no effect at all. Conversely, if the shock is moving at 4 in/s, the low-speed adjustment will have no effect. This issue makes telemetry a gold-mine of information, since shock/damper velocity is one of the channels available to us from the iRacing software without having to do any math channels. If you don't feel like bringing up telemetry for whatever reason, a general rule of thumb is that the shock will be in high-speed conditions over sharp bumps or whenever the driver is using the kerbs heavily. If something is happening outside the car and it is strong enough to physically move the car around, it's probably a high-speed situation.

[lightning mcqueen]

When compared to the low-speed adjustments, high-speed adjustments are much simpler. Since the forces that cause the shock to enter high-speed are usually very short in duration relative to low-speed forces, the main goal is to keep the tire in contact with the pavement during this time. In most cases, this involves having a much lower force at high-speed than at low-speed, which allows the shock to "blow out" over a sharp bump, allowing the suspension to compress quickly without making the wheel jump up from the track. Similarly, a lower high-speed rebound will help the suspension extend quickly after such an impact, keeping the tire in contact with the pavement. If a track is very smooth, in the case of newer pavement or relatively flat kerbs, stiff high-speed settings may help aerodynamically, but can be a major problem if a big bump is encountered.

Overall Considerations

The shocks should primarily be paired to the spring they're associated with. After all, the shock's primary job is to control the suspension and the spring. Softer springs will compress quickly and won't expand as fast, so the shock should be a high-compression, lower rebound shock. Conversely, stiffer springs won't compress quickly but will expand rapidly, needing a low-compression, high-rebound shock. In my Class B and Class A NASCAR setups, I have a general range of shocks for a given set of springs. For instance, if I use a 350-500lb/in spring, the shock will start with a 25c/10r shock, but a 1000lb/in spring will start with 10c/30r shock. These are typical starting points, and I'll adjust each shock based on the track conditions and how the car is reacting.

Driver-specific options can also come into play with shock settings, specifically when a high-rebound shock is in the car. A prime example is between Alex Scriber and me when our left-rear oval shocks are considered. He prefers a higher compression with a lower rebound, while I prefer the opposite. This leads to multiple changes across both cars, and is one of the reasons why we simply cannot run each other's setups. Don't be afraid to play around with the shock arrangement and see if you prefer one configuration over another!

iRacing's shock adjustments in the garage remove the more complex adjustments that real racing teams have to fight with on a weekly basis. Gas pressure, piston hole arrangement, shim stacks, and even shim preload are (thankfully) omitted. Could that appear some day? It's very possible, and I remember a time when many people though adjustable bump stops would be too complex to appear in the garage. These things are very simple but very powerful, and you might be surprised at how effective they can be on any car.

Commodore's Garage
Article #23 – Sideforce & the Gen6 Update

by Matt Holden
March 24th, 2017

I'm not sure if you guys have noticed, but lurking in the shadows of the dirt update was a small update to the NASCAR Gen 6 Cup cars. In addition to a reworking of the aerodynamic behavior of the car itself (2017 butter-knife spoiler!!), the suspension was updated with new, tighter restrictions for the numbers in the garage. These all seem like updates to bring everyone closer together in terms of setup, but they actually eliminate the massive setup exploits that have been going on in these cars for years. The spring limitations eliminate the coil-binding that effectively eliminated bump-stop setups in lieu of something less complex. The sway bar limitations eliminate the pre-binding that became a cool fad in 2016. Superspeedways now have the ride height requirements that exist in the real-world at Daytona and Talladega. The track bar and toe limits? Well that's a little more complex, and has possibly made the most drastic change to the car above all the other changes.

The Cause

I've worked with Nick Ottinger in what is now the NASCAR PEAK Antifreeze Series since 2012. It should be no surprise to anyone that the setups necessary to win one of these races (let alone lead the race) require a bit of trickery. In my eyes, there are two types of setup "tricks": 1) Tricks that get around a garage-page tech failure, 2) Tricks that aren't governed by a garage-page tech failure.

An example of #1 is the NASCAR Truck stagger trick back when the car was first released. Early in the iRacing days, you could turn your steering wheel in the garage and see how it would affect the car. By working with the truck's stagger adjustment (which was promptly removed), you could turn the steering wheel so that a car that was too low would clear the garage ride heights, save the setup, and start the race. This allowed the car to be extremely low on-track and didn't follow the ride height rules set in the garage.

The second trick is what was going on to warrant the track bar and rear toe limitations. Before the update this week, you could set the track bar to 6" on the left side and 15" on the right side and milk a few extra 1/16" of rear toe. Back in the COT days, it wasn't uncommon for Nick's car to have 5 or 6/16" of rear toe for qualifying. With the Gen 6 cars, he ran 3/16" everywhere. The speed was easy to add into the car, but the car was unbelievably difficult to drive. The problem was that there was no pass/fail check on any of these settings in the garage.

The Effect

Do you remember the High-Drag package that was used in 2015 at Indianapolis and Michigan? It included a 9" spoiler and a very long splitter, which generated (what NASCAR assumed) to be a large amount of downforce, a huge wake, and should have produced more passing. It was a failure…at first. It wasn't until NASCAR got the Indianapolis cars back to the R&D center that they discovered these cars produced *more* downforce in traffic due to the lack of air going over the lift-generating greenhouse. Still, the drivers couldn't pass at all despite the large wake and massive amounts of downforce.

[2015 win pic]

The reasoning behind the lack of passing was *sideforce*, something NASCAR hadn't really looked at but teams knew a lot about. While downforce is applied downward on the car, sideforce is applied in the direction towards the inside of the corner (left, in this case). More sideforce means a higher cornering speed, and the taller spoiler produced a ton of sideforce due to the rear-end offset. Another way to get more is through rear-end skew, either achieved by extra rear toe or a higher track bar angle. In our case, we did both.

NASCAR has since implemented a track bar rule in the tech inspection area. All cars must pass through tech with the track bar parallel to the ground, or even on both sides. If I ran a car through tech with the track bar at 7.5" on the left side, I have to lower the right side to 7.5" to go through tech. To meet NASCAR's zero-toe rule, I'd also need to change the length of the track bar to center the rear end housing. NASCAR will then tell me I cannot change the length of the bar once I'm out of the tech line, meaning the most skew I can get out of the car will be with the track bar at even heights since the body will shift over to the right as the track bar flattens out. This reduces the amount of track bar rake (difference in bar end heights) I can put in the car and still achieve an ideal amount of rear end skew. In the sim, running the 6-15 track bar produced a heavy amount of skew, and when combined with the extra toe, a ton of sideforce and a massive drop in laptimes…if you could hang onto it.

Sideforce Basics

Sideforce is a fairly counter-intuitive aerodynamic principle, and I'm not ashamed to say that I misunderstood it when I was first learning about it. Common ideas say that, as you move the rear of the car towards the outside of a corner and put the rear spoiler into the air stream, it should try to rotate back in line and tighten the car. That makes sense, right? In reality, the more the car is yawed out, the *looser* it will get. Below is a simplified pressure diagram of a NASCAR-style car in a yawed situation:

[Sideforce diagram – straight]

Here we would have the car traveling from the right of the image to the left, so airflow is moving from the left to the right. What is often unclear is the fact that a pressure change will occur at the left-front tire as well as the right-rear tire as the car is yawed out. Our net gains from sideforce will come from the two pressure areas applying a force towards the inside of the corner and pushing the car to the left as it moves through the corner. If the rumors are to be believed, 100lbs of downforce would be worth 0.1 seconds per lap. An extra 100lbs of sideforce, however, would be worth 0.5 seconds per lap.

[crossover]

The next thing to consider with sideforce is what's known as the "crossover point". Above is a graph of the forces acting on the car as it yaws out, with the red trace representing the rear pressure zone and the blue trace representing the front pressure zone. As yaw increases, both pressure zones will increase in force acting on the car. These two will add together and result in a lower lap time, so more sideforce is always desirable for a faster lap. In my graph (which isn't accurate to real-world, I just threw some numbers on it for this example), we see that the front pressure zone starts to exert more force beyond 3° of yaw. This point is known as the "crossover" point, and beyond that the car will change from relatively stable to aerodynamically loose.

In the case of the high-drag package (as well as the 2014 and possibly the 2015 packages), it's believed that the cars were beyond the tire's maximum slip angle in the corners, but were held in place by the high amount of sideforce. In a clean air stream, this meant the car was very fast. In traffic, where no air could cleanly get to the spoiler, the car generated very little sideforce and would start to spin out due to the tires being beyond their physical limit. This is why the cars couldn't pass, let alone follow, in the two races where the high-drag package was used. This is also what has led to the rules taking away rear end skew in the hopes of lowering sideforce amounts on the cars.

iRacing's Gen 6 Update

With all that known and understood we can now look at what has happened to the iRacing Gen 6 car with the rear toe removed. Here is the same pressure diagram from earlier modified in the way it may have looked with our 6-15 track bar and high rear toe values:

[high rear toe]

First, these diagrams aren't drawn to scale. Regardless, we can now see that the left-front pressure zone is much larger than the rear pressure zone, so it would behave a lot looser than the standard 1/16" rear toe settings. This also produces a higher net sideforce with a much lower lap time. Sure, it was very loose, and we had to run a mechanically tight setup, but it was so fast out of the gate that we couldn't leave it out of the car.

[low rear toe]

This is more like what we have now with zero toe. We have a net sideforce, but now the rear pressure zone is very large relative to the front pressure zone, leading the car to feel much tighter than it did on previous builds.

Chassis Setup Changes

These sideforce changes cause the chassis setups to shift focus in a different direction. For starters, we will have a much more aerodynamically tight car than previously, so we will need a mechanically looser car to compensate. This can be found primarily with lower crossweight, but springs will need to be slightly different due to the lower downforce from the 2017 aero package. Similarly, more focus will be placed on the track bar angle to get the best skew angle you can. Another change will be a shift towards aerodynamic downforce instead of sideforce. The setups we ran in the NASCAR PEAK Antifreeze Series were awful in terms of downforce but were tremendous in terms of sideforce. The gains we got from extra skew outweighed the losses we saw from downforce and mechanical sacrifices.

You may also find that the cars handle better in traffic now without the sideforce changes to ruin your momentum. We will likely see closer racing, more controllable cars, and a better experience driving the cars.

Final thoughts

Normally, I don't recommend completely starting over on setups following a build update, but in this case it will almost be a necessity. Everything was changed in the update, and almost everything is going to need to be rebuilt. If you had a spring package you liked in the previous build, give it a shot in the new build and see if you can make it work. If you could never find a spring package you were happy with, now's the chance to get it figured out.

If this is the direction iRacing is heading in for its asphalt cars, the future looks very bright. Too many times in the past we've seen sims defined by a single setup exploit that takes everything as far from "simulation" as possible and the developers just let it happen. This is the first time I'm aware of developers actively eliminating setup exploits that were ruining the simulation aspect of the NASCAR vehicles, while also implementing changes to bring the cars more in-line with their real-world counterparts. I have driven the new Gen 6 cars quite a lot since their release and I'm more than pleased with the updates. The cars are much more pleasant to drive, they feel more "in" the track, and setup adjustments actually produce a result. The future looks bright, and I hope you feel the same way.

Commodore's Garage
Article #24 – The Dreaded Push-Loose

by Matt Holden
May 12, 2017

All three NASCAR iRacing Series just finished a short-track swing, and during those weeks I saw, heard, got involved with, and was even asked about one of the most difficult phenomenon to understand in sim-racing: The "Push-Loose". The series went to four of the most troublesome tracks to set a car up for, starting at Martinsville in late March then visiting Bristol and Richmond with Rockingham stuck in the middle for the Class B series. This week, the series went to Kansas, and the Class B series will make a lone trek to Homestead next week while Trucks the Class A series are off. All six of these tracks are wildly different in characteristics, but all of them can trick a driver into misdiagnosing a tight car as loose. We're going to break away from the schedule the articles have been on and take a look at what's going on with a "push-loose" condition, why it's hard to diagnose properly in a simulator, and what to do to stop yourself from getting frustrated with this problem.

The Physics

A "push-loose" condition occurs when a car is experiencing understeer, or a "push", in steady-state cornering and suddenly begins to oversteer, "loose", on corner exit. For every tire in a given situation (speed, load, etc.), there is a maximum angle it can be turned relative to the direction of travel. Turning the tire will cause the contact patch to twist, and the angle of twist is known as the tire's *slip angle,* and if the maximum angle is exceeded for a tire it will lose traction and begin to slide. Below is a quick chart that is similar to how a tire would exert lateral forces based on its steered angle:

[chart]

While this is not an actual representation, it's not an uncommon sight for a typical tire. As we increase the steering angle we get a higher lateral force to a point where the lateral force plummets past a specific angle. This phenomenon is the cause of most push-loose conditions, real-world or simulation. During the steady-state cornering understeer, the front tires would be turned past the maximum lateral force angle causing a loss of grip. As the car exits the corners, the driver will unwind the steering wheel and put the front tires back within their working range where they'll magically regain traction. Despite the heavy loads on the rear of the car from corner exit due to banking, throttle application, and aerodynamics, the sudden switch in the front tires will break traction in the rear tires, causing a sudden oversteer condition on corner exit. It's this combination of events that produce what is known as a push-loose condition.

The Simulation Issue

In my opinion, it's nearly impossible to accurately simulate all of the sensations involved with a push-loose condition unless you have a remarkably sophisticated motion simulator. In sim-racing, the understeer portion typically goes un-noticed by the driver, and their diagnosis is a car that is loose off the corner. When I started sim-racing back in the mid-2000s, I had this same problem. It wasn't until I began racing karts in 2013 that I truly understood why it was so difficult to diagnose a push-loose condition through a simulator.

[kart pic]

The kart pictured above is an older-model Sodi kart used at the GoPro Motorplex in Mooresville, NC for arrive-and-drive rentals as well as the leagues I raced in for a few years. They were extremely heavy (twice as heavy as my racing kart), had a high CG, and narrow front tires. All of this meant they were extremely prone to push-loose conditions when pushed hard, and it was in one of these kart-yachts that I experienced push-loose for the first time. The sensations for a push-loose are primarily felt through a combination of the seat and steering. Whenever I'm driving the karts and I think that it may be starting to understeer, I'll make tiny movements on the wheel in the direction of the turn. For instance, if I'm turning to the right, I'll tug the wheel to the right. If I don't feel any increase in lateral g-forces, then the front tires have lost traction and the kart is understeering. It may not be plowing off towards the wall and could be tracking right along the inside kerb, but it's understeering nonetheless. Once this happens, I'll either prepare for the snap-oversteer on exit if it's a short corner or unwind the steering a bit if it's a longer corner. Regardless of the corner, it's much easier to manage and prepare for. On my racing kart, where the front tires are exposed and easier to hear, the sliding sound is enough to tell me I'm about to have a tremendous event whenever I reach corner exit.

[nick pic]

In a simulator, the seat-of-the-pants feel is missing, and what's left is a combination of visuals and sounds. This makes it extremely hard to diagnose this on-the-fly since the lateral g-force is missing from driver feedback. Drivers frequently mis-diagnose a push-loose as simply loose, and it's nothing to be embarrassed about. Nick Ottinger, whom I've worked with for nearly six years in sim-racing, was unable to diagnose a push-loose condition in the first couple of years while he was running in the Pro and WC Oval series. Today he can do it extremely well, and has gotten to where he won't immediately say the car is loose or tight, but will instead take a few laps to figure out *how* it's getting to those loose or tight conditions. A great example would be this week during the NASCAR PEAK Antifreeze Series race at Richmond, pictured above. The rear end being hung way out of line during the first quarter of the race was frustratingly common, but instead of telling me the car was loose, he asked for the car to be raised up at both ends because "The splitter's hitting and it's pushing in the center." When I asked if that would fix his exit, he said, "Yes". Having the splitter touching the track was unloading the front tires, creating a push in the center of the corner, then the rear end would swing around when the splitter raised up on corner exit, loading the front tires.

Diagnosing the Problem

Despite not having a lateral-g sensation in a simulator, it's possible to still take what you are given and realize you're experiencing a push-loose condition. We're still given excellent cues from audio in the form of tire scrub noises as well as the engine in some cases, visuals, and even tire data. It's all subtle, but effective nonetheless.

The most obvious visual for a push-loose condition is the point where the car begins to spin. In almost every situation where an oval driver says "It's spinning when I come off the corner", it's a push-loose condition. This is often the case due to track banking falling away and the car spinning once it's reached a flatter straightaway. Bristol and Richmond are exceptionally bad about this because of the large banking shift at both ends of Bristol and the 2° backstretch at Richmond. It's not limited to short tracks though, as Texas is well-known for its very difficult turn 2 exit with a relatively early banking transition. In the picture of Nick's event at Richmond (above), he's all the way out of the corner when his car began to step out on him, a tell-tale sign of a push-loose condition. If the car is relatively stable all the way to corner exit before it suddenly wants to misbehave, it's most likely a push-loose condition.

The other obvious sign is from the tire data itself. Below are the temperature numbers from one of Nick's NPAS cars that was showing a push-loose condition during the race:

[temp numbers]

I've averaged the tire numbers under the I/M/O trio to make it a little easier to see what's going on with the car. From the averages, it's immediately apparent that the car is mechanically tight, with a heavily-loaded right-front tire and a largely unloaded left-front tire. We've also got some pressure issues here, but that's going to take a backseat to the bigger problem. The major thing here is the near-16° difference in the right-side tires, and the 29.5° difference in the left-side tires. We have a problem here.

[wear numbers]

Probably the *only* time I'll ever record tire wear numbers is when a push-loose condition is expected. We know from the temperatures that the right-front and left-rear were very hot, suggesting a tight car. Normally, when those numbers are high like they were, we'd expect the hot tires to be worn significantly due to over-loading. However, we can see above that this is not the case, and the rear tires were worn more than the front tires.

Mind blown, right? We have hot tires that aren't worn, and cool tires that are, which breaks some rules. This arrangement is the hallmark of a push-loose condition, and is present *every* time this happens. The right-front tire is slipping, but not necessarily *sliding*, and is producing a lot of heat because of that. On corner exit, once the front tires regain traction, the rear end of the car loses traction, wearing the unloaded tires, causing them to wear.

Adjusting the Car

Unfortunately, when you're first learning to see these problems and feel what's going on, it's going to be difficult and it's not uncommon to make a few mistakes along the way. The key is to not get frustrated and simply take a look at what you are given and figure it out from there. The hardest thing to mentally adjust for is to have the car spinning on you unpredictably and force yourself to *loosen* the car up instead of tighten it up. The two most effective adjustments for a push-loose condition are found in the rear of the car. Apologies to the road-racing crowd, but one of them is limited to oval cars.

- Raise the right-side of the track bar. I heard a NASCAR crew chief say "Tight in the center, loose off? Raise the track bar one round". It works, and how much you want to raise it is up to you.

- Rebound on the inside rear tire. This magic adjustment directly affects the car's crossweight in transition cornering such as corner entry. A higher rebound value on this shock will decrease the crossweight approaching the center of the corner, freeing the car up and potentially curing the loose-off condition.

This unique condition is also the precursor to a lot of fundamental issues with your chassis setup, and it's not uncommon for us to start fixing one problem only to find five more. Unlike having a car that is just plowing on entry or trying to face the other direction in the center of the corner, there are a lot of issues combining into one giant headache here, and rarely is it going to be limited to one single adjustment. In the tire data from Nick's car earlier, I pointed out that there were tire pressure issues. There were also alignment issues, chassis balance issues, and overall spring package issues that were present in that car. I didn't outright fix his issues during the race, I simply made them more tolerable. We're taught the old adage "If it hits the wall with the front, it's tight; If it hits the wall with the rear, it's loose", but the level of over-simplification in that is astronomical. In the older days of motorsport, backing the car into the wall was usually cured by softening the rear springs and all was well. Today, when we combat the magical forces of aerodynamics and chassis dynamics, we have to start looking deeper into the problems we encounter on track to figure out what is really happening in our cars.

Commodore's Garage
Article #25 – Minor Settings

by Matt Holden
March 24th, 2017

We've gotten almost every major part of a race car covered in the articles so far. All that's left before we dive into the meat 'n potatoes is the adjustments in the garage menu that frequently get overlooked because they rarely need any attention. Despite this, they can often lead your setup process in a bad direction and cause you to go after a handling problem that can't be fixed without addressing the real issue. These components on the car are far less complex than springs or shocks, so they didn't warrant a full article to address their importance (or lack thereof). In most cases, you'll be able to set these options before anything else and leave them alone through your entire process and finding what's right for you is simply down to a quick test session.

Brake Bias

Most people know what brake bias does, and *most* people know how to use it properly. Various driving styles will cause front and rear tires to be unloaded or loaded differently on corner entry, and the various styles can lead to wheel lockup under heavy braking. It's important to understand that wheel lockup doesn't necessarily require a wheel to come to a stop relative to the car. If any of the four wheels on your car slows considerably without the other three following suit, it will begin to slide and is approaching lockup. The conditions necessary for brake lockup are fairly simple: heavy braking torque that can overcome the traction from the tire itself. The brake rotor for a wheel will produce a torque to counter forward rotation, traction will produce a torque against that. If the tire wins, it stays in contact with the pavement. If the brake wins, the wheel locks up.

iRacing's brake bias setting is a representation of the braking force going to the front braking systems in percent. If the brake bias setting is 60% in the garage, that means 60% of the force sent to the braking system as a whole finds its way to the front brakes and 40% goes to the rear brakes. On cars with equal-sized tires, having a forward brake bias (higher than 50%) is not uncommon because the rear tires are under much less strain during deceleration. For cars with large rear tires, like formula cars (especially the Lotus 79), the brake bias may be under 50% due to the tremendous amount of rotating mass relative to the front tires. There is not right or wrong bias across the board, it's just whatever is right for the car.

Setting the bias is fairly simple: Find whatever setting prevents wheel lockup in any tire under braking. This is going to be driver-specific due to driving styles and braking techniques, so what works for your teammates might not work for you. One of the more difficult things to diagnose is when brake bias is actually causing a corner-entry handling problem and is frequently confused with a bad shock or weight distribution. A quick check on telemetry for a sudden drop in individual wheel speed is a tell-tale sign that brake bias is incorrectly set.

[telemetry]

Above is a lap from Phoenix International Raceway, and we can see a mild lockup from the left-front tire under braking in the red trace, identified by the spikes during deceleration. A simple shift of the brake bias rearward was enough to alleviate the issue and helped the car turn into the corner better because it was no longer dragging the left-front tire along. It's very important to understand when the brake bias is actually a problem and when the chassis is a problem, since adjusting brake bias when a chassis issue is present is not going to fix the issue and could wind up making problems worse. Just because the car is spinning on entry does *not* mean the brakes are being applied unevenly!

Steering Ratio, Steering Lock

Racing sims have historically flip-flopped between using steering lock and steering ratio. Some even have both! Despite the confusing, both are very simple and both are driver specific. Once you've found what you like, you can leave it until you retire.

Steering Lock is how far the steering will move before it hits the "lock", or end of the steering rack. This is typically represented in degrees, but the measurement is usually arbitrary and never explicitly defined. Regardless of where it is measured, a higher steering lock angle will result in a larger amount of available steering motion. A smaller value will result in a lower amount of steering motion.

Steering Ratio is a representation of how "fast" the steering box is on a car. The simplest way to think of this is to look at the power-steering box on a car which will have an input and an output shaft. The steering wheel is connected to the input shaft, the steering rack will be connected to the output shaft. The Ratio can then be thought of as the relationship in rotation between the two shafts in and out of the steering box. For instance, if we have a 10:1 steering box, 10° of rotation on the input shaft will produce 1° of rotation on the output shaft. Similarly, 12:1 will produce 1° on the output shaft for 12° of input shaft rotation. This leads to a higher ratio producing a slower steering input when compared to a lower ratio.

Of all the options in the garage menu, this is by far the most driver-specific option. Every driver has an amount of steering they want to put into the wheel to get the car around a corner and this amount is almost a muscle-memory type of thing. If the driver turns to that angle and the car turns too much or too little, they will most likely complain until it's fixed. Nick Ottinger, who I've worked with in the NASCAR Pro and WC series for almost six years, is so specific on his steering box that he can tell when it's not right before he's even left pit road! He prefers an 8:1 box while I prefer a 10:1, so swapping setups back and forth between the two of us has led to this issue more times than I can count.

It's important to understand that neither of these adjustments will change how much steering is necessary *at the wheels* to navigate a turn. If the car needs to turn the wheels 15° to go around a corner, it will always require that much steering regardless of what the steering settings are. A 16:1 steering box will not provide any handling benefit or hindrance when compared to an 8:1 steering box since both will still need to turn the wheels 15° to get around the corner. Drivers simply see the car as being "tight" or "loose" because their expected input didn't provide the anticipated result.

Steering Offset

Some cars are equipped with a splined steering shaft that will allow for steering offset to be adjusted on the car. This is usually represented in degrees, and is essentially the same as removing the steering wheel, rotating it, and snapping it back onto the steering shaft. The amount rotated before re-attaching the wheel is the steering offset. This is most common in oval-track cars since they're going to be set up to pull in one direction only, resulting in the steering wheel not being completely centered when the car is driving straight. Some drivers, myself included, prefer to have a slight amount of steering to the right when the car is going straight, some prefer the wheel to be completely centered when it's going straight. Assuming you keep most of your caster and toe alignments the same, you can typically find a setting you like and leave it there forever. Adjusting those alignments may require a slight change in offset to get the feel you were looking for again.

Rear Steer: Truck Arm and Link Angles

Stock cars, with their solid axle rear suspension, have to have some form of locating arm to keep the housing situated longitudinally in the car. For late models and modifieds, both asphalt and dirt, this is achieved by linkages attached to the rear end housing at one end and the chassis at the other. For NASCAR's top-three vehicles, this is achieved through the use of truck arms. These arms can be set at angles relative to the reference plane (flat plane underneath the car) to control front-to-rear movement of the rear tires when the rear housing is traveling up or down in the car.

[brian]

Above is a picture of Brian Day's NASCAR Class B car, and I've place a white bar where the truck arm is located in the car. (Normally I'd have a good picture of one, but we haven't flipped a car in a while so I don't have a good picture of it. Maybe someday!) We are allowed to adjust the front mount of the trailing arm, located within the red box. This adjustment is simply a slotted mount where the front of the arm can be raised or lowered, thus changing the angle of the arm relative to the reference plane. Here's a picture I took of a show car way back in the day with the truck arm mounts visible, identified by white arrows. Sadly, I accidentally saved over the original with the rear roll-center image I made a few months ago, so those marks are on here as well. My apologies…

[truck arm mounts]

As the front of the truck arm is raised, the "longer" that arm will become as the fixed end travels upward on the rear end housing. This will push the tire the arm is associated with backwards, "steering" the rear end housing in the process. The same idea applies to late model rear end linkages, however on a much more pronounced scale due to NASCAR's truck arm configuration being installed at an angle. This is most obvious on a dirt late model, where high angles on the left-rear linkage will pull the left-rear tire forward on acceleration and impart a tremendous amount of rear-steer. This also lifts the chassis up onto the right-front tire, producing an effect known as being "up on the bars", the bars being the rear end linkages.

For asphalt cars, the amount of rear-steer will depend heavily on how large or small the track is, as well as how much grip is available. Low grip tracks will benefit from less rear steer to keep the rear tires stuck on acceleration, but high-grip larger tracks may benefit more from larger amounts of rear steer.

For dirt cars, it will depend heavily on the track state. Fresh tracks with large amounts of grip will benefit from less rear steer, or lower left-rear linkage angles, while a slick track will play nicely with more rear steer. In all situations, testing should show you what you need in the car, as well as your driver preferences to the amount of rear steer you like in the car. Some drivers like a lot, some don't, it's just a matter of personal taste.

[dlm 05 25]

These adjustments may seem simple and insignificant, but the small details can make or break your race. Get in a test session, experiment with everything, and find what works best for you. The days of "one setup fits all" are gone, and having the car tailored to your style and expectations can go a very long way!

Commodore's Garage
Article #26 – Pre-Build Prep, Post-Season Tests

by Matt Holden
June 2nd, 2017

Believe it or not, the grand event known as Week 13 is upon us yet again and with that comes a new build. By now, you've most likely run your final race of the 12-week season and it's simply a matter of waiting just over a week until you can race again. There's a lot you can do in that time to prepare for the next season, however, and by simply sitting around and letting the time pass you could very easily lose your rhythm from the season. To keep you in the rhythm of the racing grind during the off-week, here are a few tips and processes you can go through to make sure you're prepared for whatever is thrown at you with a new build.

Setup Exports

Few know what the "EXPORT" button on the Setup Garage menu does, but it's absolutely invaluable if you're in the process of developing a setup and want to take it into the next season, or simply want to record it for reference at a later date. Clicking this little button will create an XML file of whatever setup you have loaded at that time, exporting *everything* about the current setup from springs and shocks all the way to tire data (assuming you've just finished a run and have the data available). The benefits for doing this with your favorite setups are immense, especially since it will export the corner weights, deflections, and all the small details that could get changed a few builds down the road. This allows you to rebuild a setup from the past to get the details back to where they were without simply guessing.

Outside of season's end, it's not a bad idea to export any setup that you find works very well for you (or even your teammates). A notable example of an older export file that I still look at today is the setup Alex Scribner and Brian Day used in the 2015 NASCAR PEAK Antifreeze Series race at Daytona. I exported the setup file following the race, and it was referenced earlier this year when the series went to Daytona for the season opener. Not much of it was directly applicable because the 2014 car was on composite bump stops and the current car is on bump springs, but the weight distribution, alignments, and even ride heights were used as the starting point for this year's Daytona car.

We also exported the setup Nick Ottinger used at Watkins Glen in 2015 to lead every single lap in the NASCAR PEAK Series race. The car didn't change much from 2015 to 2016, and the car that was involved in that thrilling last-lap battle was exactly the same as the year before. Still, we have the setup details documented for all eternity, which will come in handy later this year when the series returns to road courses on the newly re-worked Gen 6 chassis. Exporting setups will completely eliminate those moments of "that setup was great, but it's gone forever", and everyone's had that moment in sim-racing at some point.

Telemetry "Mule" cars

When the Gen 6 car was in the development pipeline, there was a lot of confusion as to what the new car would be like compared to the departing "COT" Gen 5 car. To make the transition easier, as well as to help usher along the setup process for the new car, we built a dedicated test car specifically to log the telemetry data to compare it directly with the new car.

This is actually simpler than it sounds, and really doesn't take a whole lot of time to get ready. Moving from the COT to the Gen 6 car, we already knew that the Gen 6 would do away with the rear sway bars we'd been using for years to replace large right-rear springs. For the most part, we'd all forgotten how to run a car without a rear sway bar, so the test car was built without one. Outside of that, the rest of the car was extremely generic, with values we assumed would be easy to replicate on the newer car like equal front springs at 500 lb/in, a 2" sway bar, etc. Nothing outrageous and definitely nothing like we'd been using for races, but something that could be built quickly following the Gen 6 car's release. Once the setup was built, it was then a matter of driving the car *a lot* and recording the telemetry from the runs. We looked at how much the car dropped at speed as well as how quickly it raised up as it coasted to a stop (aero checks), how much the car rolled and pitched (mechanical checks), and, probably the most significant thing, how the tires heated up from the setup that was installed.

Once the Gen 6 car was released, it was simply a matter of installing the same setup (as closely as possible) and running the same tests again. Since the setup was the same on both cars, by directly comparing the data we were able to see how the new car differed from the old car as well as what had gone unchanged between the two. If you're not in the mood to build a dedicated car to do all this like we did, you can do the exact same thing by simply recording telemetry during a race and doing the same later on after a build has been released.

TESTING

You've just run a full season, and you probably had some ideas on how to make your cars better but may not have put the ideas to use. The off-season is the perfect time to test, and *really* test. Maybe you want to try a stiffer spring package, or a new arrangement of shocks, or who knows what? Take those ideas and apply them, then jot down what happens. By the time the season rolls around you will have a laundry list of things that worked and what didn't work, and those things can be implemented for the next season.

I want to stress that this point that *testing* is not the same as *practice* or *setup building* sessions. For a practice session, we're either tailoring a chassis setup to the track and conditions or we're simply trying to get experience on the track to help lower lap times or to simply become more comfortable. In a testing session, there is almost no interest in the stopwatch but instead about how ideas, configurations, or specific adjustments affect the car. We're looking to understand the car better, and that information will be applied to make the car faster in practice.

[test schedule]

In April of 2016, the iRacing Class B series had an extended break in the schedule. Instead of stabling the cars and ignoring them until the next race, we put together a list of all the things we wanted to try, the things we didn't understand, and things we wanted to develop further in the time off. The picture above is the first test we ran on the Class B cars in the sessions from April 28th to May 5th, including the details to be tested.

Prior to ever starting, we had a full schedule of what needed to be worked on and how it needed to be approached for the best results. For this session we went to Charlotte, specifically because of the high loads and relatively smooth track. The more outside influences we can eliminate (bumps, wind, weather), the more accurate the test results. As you can see, we started with the setup we used at Texas, but adjusted it slightly to suit Charlotte Motor Speedway. Once that was done, the car was not adjusted anymore outside the bounds of the tests. When doing something like this, the idea is to keep the car as constant as possible and only making a single change, thus the line about constant Sway Bar diameter and other settings. This test involved running all the spring rates available between 300 lb/in and 450 lb/in to see which springs, if any, bound too early to be useful and which ones gave us the best splitter height control.

Which spring was the fastest? I'm not sure because lap times were not recorded in the test. Instead, we recorded the ride heights and splitter heights for each spring and compared all the data to find out which springs gave us what we were looking for. Each test we ran had the same type of process, ignoring lap times an instead looking at what happened as a result from the adjustments. In the entire 31-page-long "report" (I kid you not), there are only three of the 14 tests that have any reference to the car's lap times following the adjustments.

Despite being just off of the lead pace in the early part of the 2016 season, the information gathered from the tests went into almost every single car from that point on, with a large amount of the information going straight into Alex Scribner's car that eventually got developed into the car he raced at Indianapolis a few months later, kicking off a series of wins and, eventually, a series championship.

Do you have to go as crazy as we did and write up a real-world-like schedule? Not at all! We simply had so much stuff that we wanted to try to understand that, had we *not* structured everything, it would have been a disorganized mess with nobody ever on the same page. Since we had up to four cars in a session at a time, we had to say "*This* is what we're doing right now, and *this* is what should be in the cars." There were also some tests that one driver did and the rest did not, usually because the test didn't apply to everyone's setup and would have been useless to the other drivers.

Keep on Developing!

The single worst thing you can do with a new season is to say "Well, now none of my setups work" and completely start over from scratch. This takes months of work and throws it out of the window with no idea where to go in the upcoming season. Maybe a new build changed your car somehow, but you should focus more on how the build changed the car and adapt from there. Sim-racing is unique in that everybody has the same car, so if my car gets an update, your car gets an update as well. What's really cool is that the update is exactly the same for all of us, so early-season advantages are going to come from whoever understands the car updates the best. For example, let's say the new build simply changes where a shock is mounted on the car. Some simple testing could result in discovering that, if I just increase compression on the shocks by two "clicks" I'm back to what I had last season, I'm already miles ahead of the driver who scrapped his setups and started over.

New updates to cars are inevitable in sim-racing due to advances in modeling, coding, or even updates to the real-world cars. The key to handling these changes is to not become overwhelmed with all the changes you see, and instead taking the data you have access to and finding out how that data has changed. There's no guarantee that a build will change whatever car you drive, but it never hurts to take advantage of the week off and make your cars a little bit better so you can hit the track running once the season begins.

Commodore's Garage
Article #25 – Bumpstop Contact in Telemetry

by Matt Holden
March 24th, 2017

Last week's build release put a stop to the coil-binding trend in the Class B series that had become widespread across the front-running cars. Eliminating all front springs below 500 lb/in took us away from binding by simply eliminating the springs that could bind under normal conditions, placing everyone firmly into the bump spring setup world and leveling the playing field across the series. This update is just on the heels of the Gen 6 update which did the same thing to eliminate both the coil-binding and pre-binding setups that were used, meaning both of the NASCAR cars are limited to bump spring front end setups, with the NASCAR Truck being the sole remnant of an abandoned setup era.

While the abandonment of coil-binding in the Gen 6 Class A series was largely a non-issue since coil-binding seemed to be limited to a small number of cars and teams, this has produced a big problem with the Class B series due to the low-rate springs of the past (less than 430 lb/in) all having the ability to bind before ever engaging either front bump spring. For many setups that utilized springs in the 300-400 lb/in range on the front, the bump spring was never engaged at all, leaving the drivers and setup builders to believe they were using the bump springs when, in fact, they were not ever touching them at all. In my Class B car prior to the update, I ran a 420 lb/in left-front spring and a 380 lb/in right-front spring. The 420 lb/in had enough travel to hit the splitter prior to binding, but the 380 lb/in would bind just prior to the splitter contacting the track, but after engaging the bump spring. Now, with two 500 lb/in springs being the minimum, most of your gains will be from tuning the bump spring rates and packer amounts, but the hurdle for this is determining when you have contacted the bump spring and when you have not.

This is the method I've used for a few years to determine the bump spring engagement point. I use MoTeC i2 Pro v1.1 and a custom Maths channel that isn't very hard to configure at all, just follow the steps I have outlined here. I have exactly zero experience with Atlas, so if you have experience with it and know how to convert this method for use with Atlas, please let me know and I'll add it into this article!

Finding Bump Stop/Spring Contact Point

First we need to know when the shock is going to contact the bump spring. Since the bump spring is mounted on the shock shaft and the height of the entire spring/packer stack is constant while on track, this means it will engage the bump spring at a specific deflection every single time. This consistency means it can be measured and traced to a fairly reliable amount, but the drawback is that it requires a bit of work and patience to get it working properly. This can be tedious, but it is invaluable, in my opinion.

We'll start with these values highlighted in the above picture: *Packer* and *Shock Deflection*. For a quick refresher, the *Packer* value is how tall the shim stack is on top of the bump spring; *Shock Deflection* is how much the shock has traveled from full length. In our case here, we have 0.625" of packers/shims on top of the bump spring, and the shock deflection is 3.45", meaning it has compressed 3.45" from full length. (*NOTE: Starting deflection is irrelevant for shocks aside from knowing the available travel to the bump stop/spring because shocks don't have a static force as with springs*)

The image to the left shows how the shock is currently arranged on the left-front (click for larger version). We have a bump spring, packers, and a bump stop gap to the bottom of the shock. This is fairly common for all cars using bump stops/springs with a ride height rule, however there are some exceptions I'll cover at the end of the article. At this point, note what your normal packer height is (0.625" for us right now) because we'll need to reset it to that point later.

The image to the right (also click this for larger version) shows what we're going to do with the packers now. We need to add packer until it contacts the bottom of the shock and physically moves the shock. Since the garage needs to reset and settle following a packer adjustment, it's best to add two "clicks" of packers at a time and then Apply the changes. Whenever we have contacted the shock, the ride height will increase and shock deflection will decrease, meaning the shock has been pushed up by the additional packers. Once we've found the point of contact, we'll reduce shims one at a time until the car is back at the ride heights we started at. The image below shows this process.

Doing this, we found that the shock was in contact with the packers when we had 3.000" of packers installed, but not in contact with 2.938" of packer installed because of the ride height change with the 3" packer height. Doing this will tell us how tall the bump stop gap is by simple math:

Bump Stop Gap = (Contact Packer Height) – (Starting Packer Height)
Bump Stop Gap = 3.000 – 0.625 = 2.375"

This means that the shock must travel 2.375" to contact the bump stop stack and engage the bump spring in the suspension. To find the value we need for the mathematical expression, we only need to add the bumpstop gap to the starting shock deflection (2.375 + 3.45 = 5.825"). We can now take this value and apply it to the channels in MoTeC to create a channel to show when the bump stop is actually engaging.

Creating the Maths Channel

In MoTeC, we're going to go to Tools > Maths (or press Ctrl > M to bring up the menu). This is where we can build channels that are based upon mathematical formulas dependent on other channels. In our case, we're going to depend on the Shock Deflection traces, named *Damper Pos FL/FR* on the stock cars. This is a different channel than is used on coil-over cars such as the Late Models, Modifieds, or all of the road racing cars, which use *Suspension Pos* channels.

Our mathematical expression is simply going to offset the Damper Position value relative to the bump stop engagement point, placing this point at 0.00" on the trace. This will create a display where any value over zero will be bump stop/spring deflection, and any negative value will be a disengaged bump stop. The expression is very simple, and takes almost no time at all:

In the Search bar under the Tools section, search for the Damper Position channels by entering "Damper Pos".
Double-click the "Damper Pos LF" channel to place it in the expression box.
Now, subtract the value we found for the bump stop gap from the Damper Pos channel to offset the trace. The final expression should look like this (using our value of 5.825" for the bumpstop contact point):

'Damper Pos FL'[in] – 5.825"

The expression will need to be made for both corners of the car, so you will need to make a second channel for the *Damper Pos FR* channel as well. The final expression should look like this:

[maths]

I've added a smoothing filter to my own channel to reduce some of the harsh bounces associated with higher-rate bump springs, but it is minor. It's not necessary at all, so it's your choice to use it or not. Once you're done with your expression, click "OK" and "Close" on the Maths window.

Creating the Worksheet

Since the data set associated with bump stop contact is small, we don't need a tremendously large amount of data on a worksheet to see what we're looking for. Realistically, all you really need are the two traces, but it doesn't hurt to add status lights that turn on when the value is greater than 0.00" (bump stop engaged) or a box to show the actual value. Further understanding can be gathered from a track map, speed trace, and even a Driver Marker channel to help diagnose issues on-track. Below is a picture of my own Class B car at Iowa showing the right-front bump spring engaged (and deflected 0.08") while the left-front bump spring is not engaged.

[MOTEC screenshot]

Drawbacks and Issues

The biggest issue with this method is the need to manually re-calculate the bump stop gap whenever either Shock Deflection or Packer height changes independently of the other. Shock Deflection will change whenever Ride Height changes, which is caused by any of a thousand things in the garage, while Packer height is a manual adjustment. This can get a little tedious if you're in the rough stages of a setup, but is very helpful for eliminating major issues with your setup.

This issue and need to re-calculate can be delayed (most times, at least) by making sure Packer changes are made to compliment a change in Shock Deflection. For instance, if the Shock Deflection *increases* by 0.250", we need to *decrease* the Packer setting by 0.250" to maintain the same bump stop gap. However, it's a good practice to always keep the setting up-to-date to eliminate the issues associated with this.

The Gen 6 Exception

As with many other things in the garage, the NASCAR Gen 6 cars have a unique exception due to their lack of ride height rules. For the Xfinity Class B cars, the bump stop gap is required to get around the ride height rules and keep the car low while on track. The Gen 6 doesn't need to do this, however, and it's often best to not have a bump stop gap at all to keep the acting spring rate as linear as possible. This can be achieved by simply adding packers until contact is made, and then leaving the packers at that height. Once aerodynamic forces take over, the car will be compressed enough that it will likely not disengage the bump spring, preventing the spring rate from changing suddenly in the center of the corner, as eagle-eyed readers probably noticed on my car at Iowa.

But this is soooo tedious…

Yes it is! As is the case with a lot of engineering problems, this requires a lot of detail-management to keep it correct for the current setup conditions. That said, you might notice a lot of things you didn't notice before, such as bump stops disengaging at the exact point you keep spinning in the center of the corner, or the left-front bump stop never contacts at all and you can't figure out why the car is plowing like a dump truck. It can be irritating to have to do this over and over in a single session, especially if you're someone who likes throwing springs all over the car, but it is an absolutely invaluable tool to setting a modern stock car up. If you don't know when the bump stops are in contact, you essentially don't know what springs you're running on! Take a few minutes and update the channel every once and a while and you'll be amazed at the results.

Commodore's Garage
Article #26 – Let's Get Started

by Matt Holden
June 23rd, 2017

After twenty five articles covering the major concepts and options in our sim-racing garage, it's time to actually start applying all of what we've learned to build a setup from scratch. Some of you may be thinking, "Finally…it's been over a year!", while others may be thinking about what a daunting task it is to start from essentially nothing and build a car that you're happy with. It's very understandable that anyone who has never worked with the iRacing Garage can be intimidated to the point of looking for a setup on the forums or going to one of the many setup databases. These are great points to start, but once you grow comfortable working on your own cars and learning how to get what you need out of a car you'll find yourself more comfortable with your car in the race, finish better, and no doubt have more fun with the entire process. Since starting these articles in June of last year (I wasn't kidding about the 1-year thing), I've had so many people send me messages talking about how more fun it is for them now that they're working on their own cars each week. Many have started winning races. A few have even won championships. How far you go is up to you!

As we go through the entire process, we might have to take a look at the finer details in our chassis setup as they come along, so expect at least a few articles in the future that take a focused look at something very specific that can get you stuck in the mud, so to speak. I also don't plan on leaving you with just theory for you to pull numbers out of thin air with this, so every now and then you might see a few screenshots of what I've got going on in my own NASCAR Truck. For whatever reason, I've always stayed away from the Truck and I'd like to go through the setup process with that vehicle as well at the same time you're going through yours.

Step 1 – Have the Plan Ready

I will start by saying that the "click buttons 'til I'm fast" method is *not* going to be entertained here. We want to be as efficient as possible to tackle as many problems as quickly as possible, which means we need a direction before we ever begin. The moment when you first open the garage menu should be after you already know where you're headed with the setup, what you are looking for, and how you're going to go about doing this. For example, I'll take you back to mid-2016 when Alex Scribner, Brandon Buie, and I all decided to race at Gateway in the NASCAR Truck Series.

[gateway]

After a half-season under our belts in the Class B series, this probably seems like child's play right? Wrong. Each of us had not raced the truck since it was updated to the current body style. In fact, aside from a one-off disaster at Charlotte one month before this, I hadn't raced a truck in over 4 years! We had no notes on the car, no idea what the fast cars were running, and no idea how we were going to do in the race. On top of that, the setup I borrowed to race at Charlotte was plagued with issues and we threw it away. Gateway was a truly start-from-nothing race.

The one thing we *did* know about the truck was that it didn't have bump stops on it, so coil-binding pretty much told us exactly how we'd be configuring the front-end of the chassis. We still had notes from coil-binding the Class B car so we already knew the process to get that sorted out. The issue was how everything behind the splitter needed to be set up. How much rake do we need? How does the truck behave with rear-end skew (rear toe, truck arms, track bar rake)? Can we run a package similar to the Class B car and see how it reacts?

Questions like this outlined our overall plan for the setup, with the only known aspect being a coil-bind front end. The rest was totally unknown, but having the plan ready before any of us ever got into the truck meant we'd eventually have enough information to build the setup to at least be competitive.

Step 2 – The "Set-and-forget" Options

The first thing you should always do is configure the settings in the garage that are going to have no real impact on the setup until later down the road when the issues have mostly been ironed out. For our truck, I left these options constant until late in the setup process:

Steering Ratio: Since this setting has exactly no bearing on how the car handles, I left it at the 10:1 I use in the Class B car. Brandon and Alex changed to their preferred box when they got the setup prior to the race.

Steering Offset: Also has no influence on the car's handling, but it can affect the driver's confidence on the straight. Since we hadn't worked with the truck before, I kept it at zero. The rear toe settings caused the steering wheel to be slightly right of center on the straights, which I prefer, so it stayed there. I don't know if Alex or Brandon changed their setting, but it's possible they did.

Brake Bias: I build almost every oval setup at 62.5% and leave it there all the time. The *only* time I will ever change this is if braking telemetry is showing a massive imbalance in front-to-rear braking force because loose entry is almost always a left-rear suspension problem, *not* a braking problem.

Shocks: Since we hadn't worked with the truck, which has a very different aerodynamic and weight balance compared to the Cars, I set the rear shocks at 32/32 and the front shocks at 16 compression/32 rebound. We wound up changing the left-front shock to 30c/25r but left the other three the same. This is a typical shock package I use for testing purposes and doesn't mask suspension issues. When we later adapted the truck to Kentucky a few weeks later, we changed all four shocks for the bumpy surface.

Rear Alignment: Maxed the rear toe at 4/16" and got as much rear camber as possible. Gateway is unique since Turns 1 and 2 are very slow and tight, so the extra rear toe was actually a hindrance at that end of the track. However the speed carried through Turns 3 and 4 made up for it. At somewhere like Richmond or New Hampshire, we would have likely backed off on it to help keep the rear end in line.

Front Sway Bar Linkage: This is unique to the truck now, but we opted for the chain linkage due to the curbs at each end of the track. The chain basically allows the left-front tire to hit a small bump without transferring the shock to the right-front tire which was very helpful on the bumpy curbs.

I also kept a few settings less strict, with the option to adjust as needed. The track bar, for instance, was initially set with a 2.5" split (right-side high) and left alone since it was the split I used at the time on the Class B car. I was carrying over the 250/1200 rear spring package from my Class B car, so the split in the track bar needed to stay the same. I wound up adding a half-inch to that setting for a 3" split at race time. I also started the left-front camber at +7.5° and played around with that throughout the setup process, but it was raced at +7.4°.

[pic of some kind]

Setting things like this to keep them constant is important in the beginning of the setup process to just keep the options limited to what really matters at this point. The early stages can be likened to 7-post tests in a real race shop, which is nothing but swapping springs and shocks out until you wind up with what you're looking for. Eliminating the things I've listed locks everything on the car that isn't a spring element or weight placement option, and keeps you limited to adjusting the fundamentals of the car first before moving to the finer adjustments.

Step 3 – Checklist, Setup "Zoning"

Once everything is narrowed down to the basics, it's time to look at one section of the car and keep it that way while having a mental checklist of what to work on. For the more complex road cars, where aerodynamics can be directly adjusted, you can have many sections to work in such as front-end, rake, wings, differential, and shocks, but our truck was relatively simple. I looked at the front end first, moved to the rear end, and then looked at the weight distribution. Those were the three “zones” for the Gateway setup, and I never moved to the next one until the current one was complete.

The front end was set up only to keep the splitter down as low as possible for as long as possible. All I did was swap the front springs out until I found the softest pair of springs that kept the car off the track and then those were left alone. I had the minimum ride heights outlined by the garage, so I set the front end to those heights, started at 300 lb/in, and added 10 lb/in to each spring until it no longer hit the track. The spring pair we wound up running was 350 on the left-front and 370 on the right-front, with the left-front height set to 5.65”. The 350 ran just above coil-bind when the splitter hit the track, so raising the nose allowed it to bind. The higher right-front spring simply allowed the front to roll over a little more and bind later, helping turn-in. Curiously, we always ran the stiffer spring on the left-front of the Class B car!

[another pic]

Once that was set, it was a matter of setting the rear ride heights so the chassis didn’t hit the track. I kept the springs at 250/1200 because that was what I was familiar with in the Class B car and knew how to adjust with them. In theory, *any* spring combination would have worked as long as the ride heights were set to keep it off the track. In fact, we even halved the right-rear spring rate and installed a sway bar in the rear to help with bite off the corners, but found it to be too finicky and went back to the more conventional pair. Not to mention that none of us had used a rear sway bar since the Cup COT went out of style 3 years earlier.

After the rear was set, it was just a matter of adjusting the crossweight, shocks, and ballast until the truck worked as planned. Keeping your adjustments in one place, while staying focused on the current problem, is the best way to keep yourself from introducing multiple issues at a time. If I had changed something on the rear end while I was working on the splitter, I would have changed how far the rear dropped in the corners. This changes how far the splitter would have been from the track, and from that alone I now have two problems to work with.

Step 4 – Rinse and Repeat

[Nick’s Sonoma car]

Simplifying the initial setup process is the hardest thing to get into the habit of when building a setup. You could be working on the front end, but you may be a half second off the fast cars and the car just doesn’t behave properly. It’s natural to want to look at fixing the issues, but it’s very important to remain focused on the problem at hand and getting that sorted out before ever going somewhere else on the car. We want to have as few hurdles to overcome at a single time as possible, and running all over the car can build the hurdles into mountains. Before you know it, you’ll take a tiny issue and grow it into a tremendous problem that is nearly impossible to overcome without reversing every change made to the car. Keep everything focused and get the big details correct, *then* move to the other stuff.

Sometimes, completing each "zone" on the car will affect another zone so it can be a merry-go-round of adjustments until the entire car is how you want it. That's okay, and it's also normal. Like I said earlier, adjusting the rear of the car will directly change the splitter, and changing the splitter will adjust the aero balance and the needed rake as a result. It can be tedious at first, at least until you get the hang of it. Still, the setup I went over in this article finished 1st, 2nd, and 4th in its only race, so even with no knowledge of how the car should work, I think it did alright!

Commodore's Garage
Article #27 – Related Adjustments

by Matt Holden
July 28rd, 2017

In the last article we looked at the process for initial chassis setup. To recap, we’re trying to replicate what a real-world team would do on a pull-down rig in the shop to get the car essentially race-ready by the time it arrives at the track. This involves cycling through various spring options to find a package we can use as the base for our chassis setup while leaving virtually everything else on the car alone. The first stage of the setup process is not really a search for handling so much as it’s a search for the platform to work off of. We want the car off the track, we want a good aero platform for the track we’re racing at, and we’re looking for responsiveness with adjustments. In the best case scenario, a shock adjustment of just two “clicks” should be enough to produce some kind of handling change felt by the driver. If you’re not at that point yet there’s no need to worry, we’ll eventually get there!

Beyond the initial spring package, there are a couple of adjustments that we can look at right now that can dial in your car a little bit beyond the baseline built from last week’s article, but these are going to be largely dependent on other characteristics in the car, and even the track itself. These are relatively simple adjustments, but are almost always changed when other things in the setup are changed.

Ballast and Nose Weight

A few of the cars in iRacing's road racing stable and the majority of the cars on the oval racing stable have an adjustment known as "Forward Ballast" which is a direct adjustment for the car's nose-weight percentage. This adjustment is simply blocks of Lead or Tungsten that can be moved forward or rearward to tune the car's handling. Every car builder has their own unique way to represent this in their chassis setup sheets, whether it is exact locations on the car in terms of a measurement or something as simple as drawing the location on an outline of the car. Our garage represents the adjustment in a measurement of length relative to some arbitrary point in the car. This point's location is unique to the car, and largely unimportant in the grand scheme, so we'll discuss any ballast adjustments in terms of the car's Nose Weight percentage, or the amount of weight situated on the front tires, *not* the value of the ballast location itself. Knowing the percentage instead of the location is extremely important when transferring knowledge from one car to the other or whenever a build update changes the car. For example, if Nick Ottinger wants to borrow notes from my NASCAR Class B car for his NASCAR Class A car, we can't use the ballast location adjustment at all since they're totally different. Nose Weight percentage, however, *will* transfer between the two cars. Plus, the ballast location changed at the start of the Season 3 build for the Class B cars, so knowing the percentage was helpful in getting the car dialed in again following the update.

[05 and 25 pocono cars]

We looked at what the ballast does in the third *Commodore's Garage* article, which can be found here: http://www.iracing.com/commodores-garage-3-centers-of-mass/. The quick refresher:

- Ballast location controls the front-to-rear location of the car's Center of Mass. More forward ballast (higher Nose Weight percentage) shifts the CoM forward, and less forward ballast (lower Nose Weight percentage) shifts the CoM rearward.
- For most applications, higher Nose Weight percentages will increase the weight on the front tires and create more understeer. Lower Nose Weight percentages will increase the weight on the rear tires and create more oversteer.
- This effect flips at very low speeds, usually somewhere under 50mph or 80kph unless the car is very efficient with aerodynamics, but the transition is very mild if it happens at all.

Since we've looked at aerodynamics and spring packages by now, we can dive deeper into this setting and how it will affect the car in more detail. For almost every track, our best performing car will be one that has very efficient aerodynamic properties. It will create a lot of downforce without a lot of drag, causing mechanical grip and suspension settings to become an afterthought in some cases. For many cars this results in a bias towards more rear downforce for traction and stability, but it can sometimes be the opposite depending on the driver. Regardless of how the car is set up aerodynamically, the aero balance is never 50-50 and one end of the car will always have more downforce than the other. During tuning for aerodynamic balance, the ballast is used to counter the inevitable aerodynamic imbalance by generating mechanical under- or oversteer to compliment the effects of aerodynamics. Essentially, if we have a lot of front downforce the car will turn extremely well but could be unstable, so ballast can be moved forward to impart stability while keeping the aerodynamic handling effects present. Similarly, a car with a lot of rear downforce may see the ballast moved rearward to impart some mechanical oversteer. This allows the car to turn, mechanically speaking, while the aerodynamics will keep the rear tires stuck to the track.

All of this considered, the Nose Weight of the car has a direct correlation to how much rake is in the chassis and how high the splitter or valence is run above the track. In one of the few "if this, do this" adjustments in the entire garage, Nose Weight should be increased with added front downforce and decreased with more rear downforce if the intent is to keep the handling characteristics the same. Doing this results in a change of the car's aerodynamic efficiency, not necessarily the handling itself. When either downforce or Nose Weight is changed independently of the other, this will result in an immediate handling shift towards either under- or oversteer. A common adjustment for oval cars for qualifying is to add grille tape, which increases front downforce, so Nose Weight would need to be increased to compensate and keep the handling from shifting too far towards oversteer as a result. This combination of adjustments, if made together, would keep the car's handling the same as when the tape was set in race trim, but would decrease the drag acting on the car. Having less drag with no change in handling is free speed!

Track Bar Height/Angle

A more oval-specific adjustment, the track bar angle is usually set to agree with the amount of rear end travel. Ideally, the track bar should be perfectly flat and parallel to the ground when cornering for aerodynamic reasons, so the bar needs to be set to wind up at that point by using a difference in height (often called "rake"). Currently, all of iRacing's oval cars that use a track bar have the left-side mounted to the rear end housing with the right-side mounted to the car's frame. If we're going for a flat bar in the corners, this means the right-side will need to be set higher than the left-side since that will be the end that moves in travel.

The amount to have the bar rake initially set to can be found through a quick telemetry check by checking the right-rear ride height in the corners, which I've shown in the picture above. Here I've isolated turn 1 at Iowa and can see that my lowest right-rear ride height was 5.85". I started the right-rear corner at around 7.75", so the rear is dropping just under 2" in the corners, meaning I should start the track bar with a 2" rake and work from there. This is rarely going to be the setting with which you race, but it's a great starting point since having the bar flat in the corners not only aids in aerodynamic performance, but also reduces the jacking force amount from the bar itself.

It should be noted that this does often have some exceptions where this isn't the ideal way to set the track bar. Typically, you'll find the most performance with this configuration at a very fast track, especially large speedways. As the track gets shorter, speeds reduce, and grip level decreases, the jacking forces from the track bar can be used to both keep the rear end up and keep weight off the right-rear tire and on the left-rear tire by simply running the right-side of the bar lower than the left-side. This does hurt aerodynamics, but the added grip can be very beneficial at slower tracks. The other major exception would be Daytona and Talladega, where running the most rake that is legally allowed will help to pull the back of the car down out of the air. If you ever get a good look at a NASCAR Cup or Xfinity car upside-down at either of these tracks, take a good look at the track bar and you'll see it way up in the back of the car with the right-side as high as possible.

The relationship between the right-rear ride height and right-side track bar setting ultimately links the track bar height directly to the right rear spring rate, so these two settings need to be considered together when making spring adjustments on the right-rear corner. If you soften the right-rear spring, it can result in extra travel at that corner, requiring a higher static right-side track bar height. Conversely, if you run a stiffer right-rear spring, the track bar needs to start lower due to less rear travel. Of course, these are simply starting points, so adjusting the track bar is necessary while tuning. This is also the reason why a lot of people choose to run the same (or very similar) rear spring packages at every track in the Class B and Class C oval series.

Other Linked Settings

There are a lot of garage settings that will rely on another specific adjustment, but most are far less strict than the two I've gone into detail with. Here's a quick list of a few things that need to be considered together when setting the car up to make sure you don't work yourself into an unsolvable problem down the road:

- Bumpstop rate and Packers: Changing the bump stop/spring rate will always require a change in the amount of packer shims you have installed. Stiffer springs will keep the front end up in the air, while softer springs will allow more travel. Keep that in mind when changing bump stop rates.
- Rear Springs and Ride Heights: Rear spring rates will change how far the rear end will travel so the heights need to be changed accordingly. Since you're typically looking for the same ride height at full travel, the starting height will need to be raised or lowered when you soften or stiffen the rear springs.
- Rear Spring Split and Crossweight: Higher rear spring splits will result in less force on the left-rear tire, especially on corner exit. With the higher split, crossweight will also need to be increased to keep the car stable on exit and entry. The opposite is true for lower spring splits.

Final Thoughts

The truth that nobody wants to hear is that everything in the garage will, somehow, affect something else on the car. It's for this reason why you need to be aware of what you're changing in the car and how it will affect the rest of the car to keep from unintentionally adjusting multiple things at once. We're all told to "adjust only one thing at a time" and this often means being aware that one adjustment changes multiple things. It's very easy to make a few small adjustments and completely ruin the car, and in many cases it can be traced back to simply not keeping up with how the car is changing before it ever hits the race track.

Commodore's Garage
Article #28 – The 2016 Class B Champion

by Matt Holden
August 11th, 2017

Setup details at the top of the sim-racing ladder are a closely-guarded secret; we all have experienced it at some point. Earlier this year, I had a conversation with the crew chief of the car that has won a few recent NASCAR Modified championships about chassis setup, and he willingly gave me the spring rates on his car for the three packages they run in a year. No questions asked, just a simple "Here they are!" Anybody who's been in sim-racing for a decent amount of time knows that this would *never* happen, so why would he just hand over his car's setup to some guy that struck up a conversation? The key is what he *didn't* tell me about the car: shocks and suspension geometry. In real-world racing, you can hand someone the springs, weight distribution, and even alignment values for your car, and it's unlikely that they're going to be able to make it work immediately. Our Modified Crew Chief friend never told me where the roll-center in the front suspension was located; he never told me how his shocks were built. Without that information, the springs, weight, tire pressures, and everything else was essentially useless to me. If you want to go even deeper, I recently went on a tour of a top-level NASCAR team and when I fired off the random question about springs, I was literally handed the largest sway bar they use as well as told the rate of the stiffest bump spring they had in the shop. All of this was totally useless to me.

Sim-racing doesn't work like that though: we all have the same suspension geometry, we all have the same chassis, we all have the same shocks. This means that if you gave me the four (or six) spring rates on your car, I could make your car in a relatively short amount of time. The speed in the NASCAR Cup Series is rooted in team secrets like flexible truck arms or track bar mounts, while the secrets to Sim-racing speed lie in the setup itself. Last year, following a team test in April, I found a bit of information on the then-current build for iRacing's NASCAR Class B Xfinity cars that went into Alex Scribner's car later in the year. This setup was used in almost every single race from his first win at Indianapolis until Homestead, and then modified slightly for the start of 2017. While we were doing some gnarly witchcraft in the NASCAR PEAK Antifreeze Series with pre-binding, Alex's car was about as ordinary as it could be. Now, just like the pre-bind setup, it's been made obsolete by build updates. A lot of people tried to say it was full of tricks and secrets, but that really wasn't the case. It was about as ordinary as could be. Let's take a look at it, and see what made it work well enough that it snagged the 2016 NASCAR iRacing Class B Championship.

Front End Springs

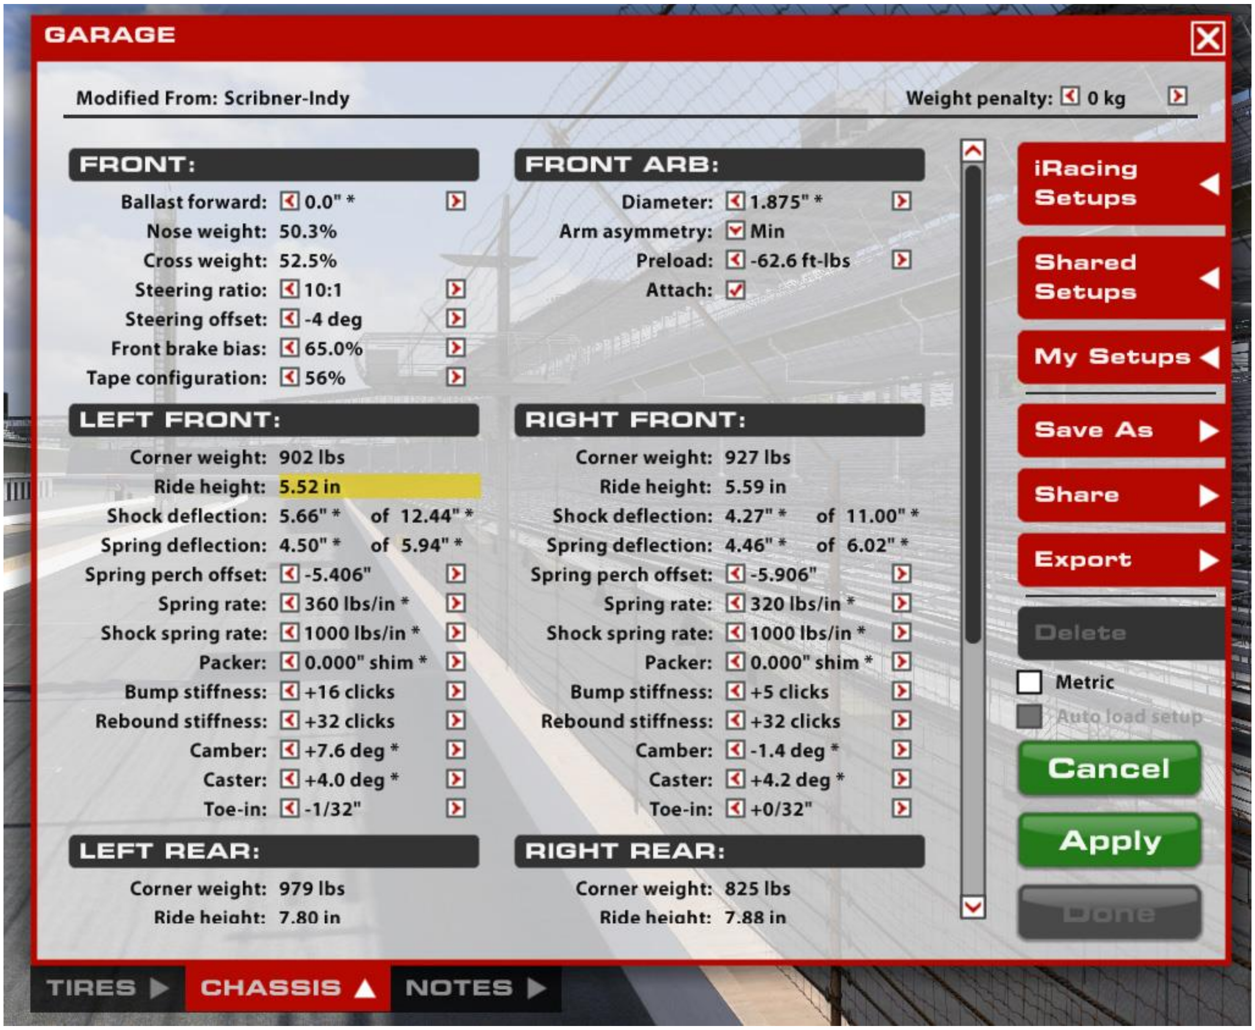

Click that for a larger version. Pretty straightforward isn't it? Outside of the car having a purely coil-bind front end despite the bump-spring update in early 2016, it's not out of the ordinary at all, but it was fast. Very fast.

To start with the driver- and setup-specific options, we can go to the Front ARB. Alex used a coil-bind setup, which meant he needed a larger bar than what I was using on my bump-spring cars. He ran this 1.875" bar at Indianapolis, where I ran a 1.625", for example. What might be a surprise is the lack of bar asymmetry, which we tried to keep as low as possible. The asymmetry adds a vertical spring rate to the front end, known as "heave" rate, which causes the front end to lift quicker out of the corners. At Indianapolis, where the straights are extremely long, this is horrible with the extra drag. Alex also rarely ran a very large amount of preload, typically under 100 ft-lbs, which was around half of what I would use (~220 ft-lbs), but that just came down to personal preference.

The springs are obviously set for coil-binding with all shims removed. The 360 on the left-front allows for a little more travel than the 320, causing the left-front to naturally sit lower on the straights than the right-front. This produces a little bit of extra turn-in due to aero and was something we carried over to my bump-spring cars with a 420/380 front-end spring package. When we got to a track, it was simply a matter of cycling through the front springs until the splitter no longer hit the track. This was done by setting the front end to a rough height with the right-front spring, then running the highest-rate left-front spring we could get away with without dragging the track. Higher rate was more travel, more travel was more downforce. Since Alex was usually running the car with the left-front bound for the entire lap and the right-front bound through the entire corner, the rate values themselves were irrelevant. They were selected for height values alone, not for handling.

Rear End Springs

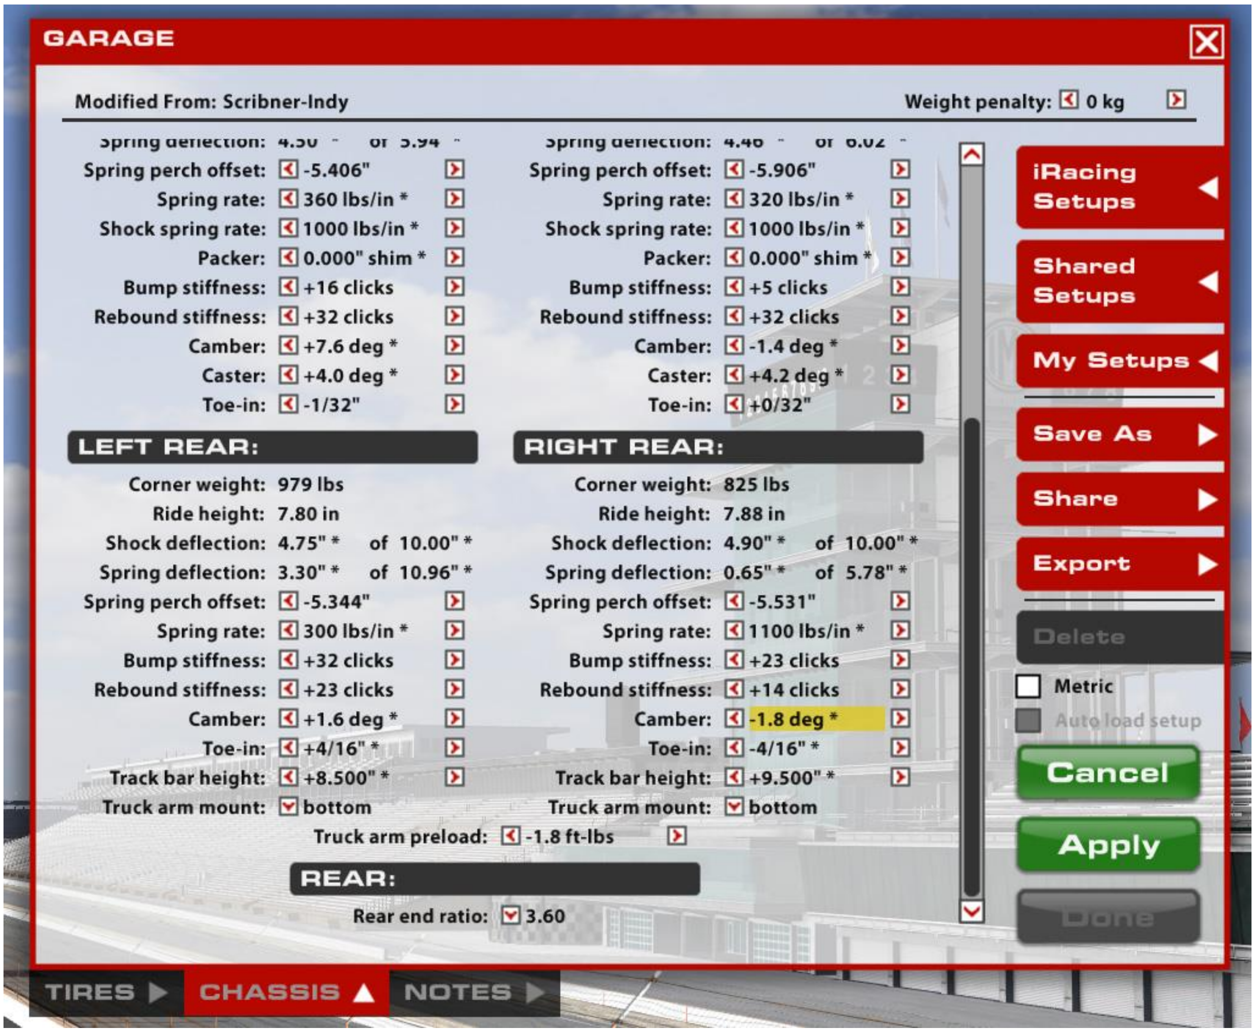

Again, we have a fairly ordinary rear end in the car, especially for a coil-binding setup. The 1100 lb/in right-rear spring was typical for Alex's cars, with a 1300 lb/in appearing in the car at a few tracks. This high rear spring rate had two purposes: First, the high rate was needed to balance the huge sway bar in the front of the car since we don't have a rear sway bar. Remember from earlier articles that we want the two suspension roll stiffnesses to be relatively close, so since we have a huge front sway bar, we need something to do the same for the rear. The second purpose was purely aero, and intended to keep the car at a tail-up attitude for the whole lap. Softer rate springs at the bigger tracks allowed too much splitter lift and killed the center of the corner. However, if the right-rear spring was set too stiff, the rear deck lid was too high and the car got very tight. Like many things, it was a balance of aero and ride heights, not handling.

The handling spring was the left-rear. The other three springs served specific purposes (mostly aero) so they couldn't be touched when handling needed some adjustment, leaving the left-rear as the one to shoulder that task. Of all the things on Alex's cars, the left-rear spring is what I know the least about, but I do know that he always ran a narrower spring split than I did. Here, he has a 300/1100 rear end while my Indianapolis car last year was 200/1200. Regardless of what he ran, this spring was the one adjusted for handling, and the *only* one adjusted for handling.

Shocks and Alignments

A lot of people ignore the shocks and say they don't work, which is completely wrong. They work quite well, but not to the extent you'd expect. Shocks were probably the most profound discovery we made during the April test, showing that we could change tire temperatures through the shock settings by themselves without having to mess around with spring perches and alignments. Alex ran high-rebound front shocks to prevent the disaster of an un-bound coil-bind spring and relatively stiff compression rear shocks. My bump-spring cars, on the other hand, ran lower front rebound to allow the springs to work and the rear compression was lower to let the car deal with bumps better. Between our two cars, there was never a race where any of our shocks were set to 32/32 or something simple like that. Every track had unique shocks for each corner.

Alignments, despite looking strange, were pulled straight from real-world data. We ran max camber on the left-front, max rear camber, max rear toe, and set the right-front camber for temperature spread. Front toe was a driver preference, and each driver on our team had what they liked. Alex typically had this zero-toe-out setting while I toed the left-front out 1/16" at smooth tracks and the right-front out 1/16" at bumpy tracks. As with everything: to each his own.

Weight Distribution

Where these cars were tuned primarily was through the weight distribution. The harsh nature of coil-binding often required a lower front weight percentage to overcome the front end skipping across the track since there was no working suspension. For those curious, this front end behavior was the reason why I only ever raced Alex's setup at one track: Kansas. I would destroy the right-front tire at any other track. The ballast was usually left pretty far back in the car since it was, mechanically, very tight.

Crossweight was not only used to tune the car's handling through the corners, but also the splitter height. At some point during the season, Alex found that changing the crossweight shifted the spring deflection for the front springs which changed when the springs would bind. Shift the crossweight one way and the left-front would bind sooner and the right-front later. Shift it the other way and the opposite happened. This was used to either raise or lower the splitter at speed, allowing the splitter to clear bumps at a track like Kentucky, but keep it very low for a smooth track like Kansas. The biggest jump forward in terms of speed was when he was able to reliably clear the splitter from the track on a week-to-week basis.

Handling

So the majority of the car was set with a specific purpose and couldn't be changed, so how did he adjust how the car handled? Well, aside from the shocks, this was achieved through track bar height. With everyone's car, we aimed for equal tire temperatures on both sides of the car (LF=LR, RF=RR), and this was tuned through overall track bar height. Initially, Alex's cars weren't very good on shorter runs, but the work to balance the car with the track bar height would pay off in droves if the run went longer than 7-8 laps.

Development, 2017 Use

This car's final use was the spring Richmond race earlier this year. In 2016, I ran this setup at Richmond in the fall with a softer rear end (850lb/in RR spring), and I carried the exact same setup over to 2017 when we returned. I started 3rd and finished 6th in the big race that week, and the car worked just as well as it did the year before with only the most minor of changes. It was made obsolete in early Summer when the car was updated to disallow coil-binding, shifting every car over to bump-spring front ends.

Why would I bring up this setup now? Aside from it being obsolete, the weekly setup process for it went through the steps I outlined in the past few articles: Set the front height, set the rear height, alignment, profit. On average, Alex's car was set up and ready to go after about 20 minutes on track. We only ever had one session per week for practice which was typically around 2 hours long. If we all weren't on long runs after 40 minutes there was a massive problem with the cars. We managed to get the time down so efficiently by carrying the exact same setup to each track each week and developing the notebook on how to adjust it. If we went to a smooth track, we knew what to change to lower the car and keep the handling in check. When we went to Kentucky, we knew how to raise it up for the bumps and what to do to keep the car from going full Kangaroo over the bigger jumps. It became simple because we knew what we were working with each and every week. It wasn't complicated; it wasn't a guess-and-click type of session. We simply had a plan each week of how to set the car up and worked through each step one at a time. Taking what was learned the week before and applying it the next week eventually developed the setup into a winning car, and that winning car eventually snagged Gale Force Racing its first championship since Nick Ottinger's 2011 Street Stock championship.

Commodore's Garage
Article #29 – Shocking!
by Matt Holden
September 15th, 2017

In oval sim-racing, shock settings get overlooked quite frequently. They're a cure-all for most cases on the road racing side of things, but why are they typically set to their extremes and ignored for oval cars? Few people ever take the time to go through the shocks and see how they might affect the car's handling and speed, leaving easy speed on the table for most cars. Is it the complexity or the belief that they don't do anything? Who knows, but fiddling with the shock settings to get them matched to your track conditions and spring package can be worth a few tenths with very little effort.

The Basics

Shock absorbers are used in passenger cars to dampen (hence the name "Damper") vibrations and oscillations before they reach the driver. This gives the car a pleasant and enjoyable ride for the passengers. From a performance aspect, the driver's comfort is rarely a concern; however the shocks will still be built to smooth out some bumps and knocks that would upset the car. In some cases, shocks may be built so stiff that they don't allow for much body movement, seen primarily in short-track oval racing.

Basic operations within the shock are very simple: A piston moves up and down through a fluid, with the valves in the piston controlling how much fluid can pass through the piston. Allowing more fluid produces a softer shock, restricting the fluid flow is a stiffer shock. How much fluid passes through the shock piston can be broken down further into what's known as "High-Speed" and "Low-Speed" adjustments, and knowing when each adjustment is necessary is very important when tuning shocks. There are situations where a high-speed adjustment is necessary, and times when a low-speed adjustment is necessary, but each are independent.

Pictured here is a simple cutaway drawing of a shock piston inside the shock body. I've added arrows to indicate the various parts we're concerned with, as well as colored arrows to indicate fluid flow through the piston. In low-speed operation, fluid will travel through holes on the outer edges of the piston ("bleed" holes), which are fairly small relative to the high-speed holes. The green arrows represent fluid flow in rebound, and the pink arrows represent fluid flow in compression. Typically, low-speed behavior takes place when the shock piston is moving less than 2.0-3.0 in/s (50-75 mm/s).

When the piston exceeds the speeds for low-speed operation it enters high-speed operation. Pictured to the right is the piston in a high-speed situation. At this point, a stack of very thin metallic shims will "pop" open and expose larger holes on the piston face, allowing more fluid to travel through the shock at the higher speeds. The science behind high-speed shock settings is almost infinite, with various shim arrangements and piston hole configurations determining how the shock behaves at the higher speed range of operation. Hopefully you can see why Shock Specialists are a necessary job position on most racing teams by now.

It's important to know that shocks do not behave *at all* like a spring does. Shocks are intended to keep the spring from oscillating out of control when encounters a bump or body movement, exerting a force to counter the oscillation. While springs operate in a linear force vs. displacement behavior (one single force for a given deflection), shock deflection is completely irrelevant in the force it will exert. Instead, shocks exert force based on piston *velocity*. For a given piston speed, it will exert *this* amount of force. For example, while a spring may exert 100 pounds at 1 inch of deflection, a shock may exert 100 pounds at 1 inch-per-second. The unit used to describe shocks is thus the pound-per-inch-per-second (lb/in/s), or Newton-per-meter-per-second (N/m/s).

Adjustments

What we have available for adjustments in the iRacing garage is the adjustment that would be available to a team while the car is sitting in the garage bay. For most cars, this is the "click" or "sweep". Some cars have both high- and low-speed adjustments, while some just have a low-speed adjustment. If it isn't explicitly stated which speed the adjustment is then it's a low-speed adjustment.

Regardless of the adjustment, the mechanics behind the adjustment is usually fairly consistent. An adjuster knob at the bottom of the shock will move a needle within the shock shaft and effectively close or open the low-speed bleed holes in the piston. Higher numbers close the valves and make the shock stiffer in that adjustment and lower numbers will open the holes and soften the shock. Some cars have a "sweep" adjuster which is a very similar internal adjustment, however the adjustment knob usually won't have clicks built in to give definitive adjustment steps on the real shock. It's not tremendously important to know which adjustment you're working with because they're the same for our situation, but it's important to understand that one "click" on one car may not be the same as a "click" on another car. Testing is important to figure out how much the car is affected by a single adjustment increment.

Spring Pairing

Like I said at the beginning of the article, a shock's primary job is to control suspension oscillations. To get this, we want to pair the shocks with the spring on that corner of the car to get the proper amount of control over the spring. If the shock is too stiff then the car can start bouncing, but if it's too soft it could be very floppy and unresponsive to driver inputs.

Luckily, this part is fairly simple when compared to the rest of the things to know for shocks. We can split it into two categories: Stiff springs and Soft springs, with everything else falling somewhere in between. Stiff springs will resist compression very easily due to the higher rates, so the shock doesn't need a lot of compression to control movement. On the other hand, a heavy-rate spring will want to return to equilibrium very quickly, and with a lot of enthusiasm, so a heavier rebound setting is necessary to keep it from violently returning to its happy place. If we flip it all around and look at a softer spring, the spring won't resist compression, so we need more compression to keep the car from slamming down onto the track. Unlike the stiff spring, a softer spring won't try to return to static height very quickly due to aerodynamics and track loading, so extremely high rebound is not necessary.

Adjusting Shocks

Shock adjustments will come in two flavors: Compression and Rebound. Compression (also "bump") is when the shock is in…well, compression. Anything that causes the shock to shorten will be handled by the shock's Compression settings. The opposite, expansion, is known as Rebound. This is anything that causes the shock to expand, typically spring rates or hill crests. Each of these is then broken into high- and low-speed adjustments which are fairly simple: High-speed adjustments control bumps and low-speed adjustments control body motion.

The most common adjustment you'll see in the garage is low-speed adjustment. These adjustments control the car in what's known as "transition", essentially whenever the car is pitching or rolling due to driver inputs. When the driver hits the brakes and the car dives, or the driver hits the throttle and the car squats, this is within the realm of low-speed adjustments. Low-speed adjustments can also handle large changes in track shape such as hills or crests, but not sharp bumps.

We'll start with the simplest option, low-speed compression, because it's about as straight-forward as it can get. Higher compression values will slow body motion and delay full weight transfer for a situation. If the driver slams on the brakes, for instance, weight is going to transfer forward to the front tires and the car will dive. If we add compression to the front shocks, the severity of the dive will be reduced, and the opposite will occur if we reduce compression. The primary concern with compression settings is to not have dive or squat happen too quickly. If this is too severe the car will change handling balance extremely quickly and could cause the car to spin or do something else silly. If the dive and squat are restricted too much, weight transfer necessary for cornering won't happen and corner speeds will suffer. In most cases, all we're looking for with compression is a smooth and consistent transition to driver inputs and not much more.

Asymmetric compression adjustments can also be used to temporarily load one of the outside tires during corner entry and exit, something more common with oval cars or on banked turns. The shock will exert a force under compression which will make its way to the tire (as well as resist the body motion), and this force can be used to add or remove load from a tire. A great example where this is common is a very high-banked oval such as Bristol, where the banking and speeds introduce a lot of load very quickly. If a driver is loose, or oversteering, on entry to the corners, more compression on the right-front can load the tire more under braking and entry to the banking to give a little more security. Similarly, a driver who is tight, or understeering, off the corners can increase the right-rear compression and put a tiny bit more weight on that tire to help it rotate. This isn't a fix-all by any means, but the adjustment is typically made for security reasons when a driver feels like the car is "on the edge".

Rebound settings become much more complex because of how they operate. Similar to compression in dive and roll, rebound should be set to keep the car from lifting too quickly with driver inputs. However, if rebound is too stiff then the wheels won't be allowed to extend as they're unloaded which could cause them to unload completely and lose grip. That's bad. Obviously there are some aero considerations with rebound as well, so keeping the aero-platform smooth is very important as well. Remember: Keep the car smooth in transition and all will be well.

Again, as with almost every oval car, rebound settings can do some funky junk. If a rebound adjustment is made asymmetrically, it will affect the diagonally-opposite tires. This means if you adjust the left-rear rebound, the load difference will occur on the left-front and right-rear tires. It's easiest to think of the entire car as a diagonal see-saw, with the fulcrum across the left-front and right rear tires. A force that pushes down on the right-front will pull up on the left-rear, and the other two tires will be the axis around which this occurs, so if we increase the resistance on the left-rear, we wind up increasing the load on the left-front and right-rear more than the load on the left-rear, where the adjustment was made. The same thing would happen if we made the same adjustment on the right-front shock as well. Below is an illustration of what happens:

[rebound cross]

This effect is very commonly used to manipulate dynamic crossweight in transition. If the car is tight on entry, more left-rear rebound will load the left-front and right-rear tires, which reduces the crossweight while the shock is moving. Similarly, if the car is tight on exit, more right-front rebound will help the car turn off the corner. If the car is loose on entry, more right-rear rebound will tighten it in the same way that more left-front rebound will tighten the car on corner exit.

High Speed Adjustment

Finally, we get to High Speed adjustments. Since these only come into play at high piston speeds, the only way to get a result from these is to hit a sharp bump. This almost never happens on ovals where we'd need to adjust for it, but it's very common in road racing over corner curbs (or Kerbs, for the rest of the world). High speed adjustments are intended to allow the shocks to "blow out", essentially reduce force suddenly, when the suspension encounters a sudden bump. The force itself isn't lower than low-speed, but it's lower than what it would have been if the low-speed linearity continued. There's not much handling influence here outside of keeping the wheel in contact without moving the rest of the car, so if a bump or kerb causes the car to jump up, a reduction in high-speed compression would help this. Similarly, we also want the wheel to return to its normal position, so reducing high-speed rebound would help out here as well. As with low-speed, we don't want to over-damp anything, just have enough to keep the tires working.

[24 hours pic]

Don't be afraid to play!

I'll leave you with a reminder: Don't be afraid to try a shock adjustment. For most cars, it's the only adjustment in the garage that won't change the rest of the car, making it the easiest adjustment you can make! The fixed setups and baselines may have min or max settings, but that's rarely the fastest setting available to you, so don't be afraid to move the shocks around and soften one corner for bumps or stiffen one corner for aero.

Commodore's Garage
Article #30 – A Whole New World

by Matt Holden
December 15th, 2017

Did anybody else notice we have a new build? I know, it's a surprise to me as well.

This year has been a year of major updates to iRacing's stock car platforms, and we got a pretty big one right at the end. To look back at this same time last year and think about what we were doing then with setups and comparing it to what we're doing now is wild. The setups we used at the end of 2016 were about as far from realistic as we could have gotten, and some of the setups my team used in this week's Pro race at Atlanta were the most realistic I've seen in a simulator. As it turns out, I wasn't alone in that thought so we must be doing something right.

Instead of going into great detail about what this build likes and doesn't like behind the wheel (and diving into how one of our drivers was managing tires so well that he passed 20 cars in one tire stint in the Pro race), I'll save that for later when we have a much better handle on what's going on. The ugly truth is that we really don't know what's fast right now from a driving standpoint because there were so many approaches to it this week, all of them appearing in the lower splits of the Pro race. In time, however, some of those strategies will start to show which are better and which should be tossed aside. What we can look at is how this build has altered the garage and what new challenges have been thrown at us.

The "Must Know" – Crossweight

Crossweight is, easily, the most useful tool in the garage. It also takes a little while to get a hang of due to the complexity involved, and I've had a few people admit that they don't understand how to change it. That's not a big deal and it's nothing to be ashamed of if you've never gotten a hang of it. I have no problems admitting that I still mess up so badly with crossweight changes that I have to stop and reload setups because I've missed something silly in the process.

Crossweight in an oval car is basically a multi-effect adjustment: It will cause multiple things to happen and all of those must be accounted for. Primarily, I will change crossweight to affect entry/exit bite with the center of the corner being ignored in the adjustment. Since crossweight is a figure of how much static weight is on the left-rear and left-front tires, it can be seen of how much those two tires are going to bite into the track while turning. Don't forget that friction is in direct relation to vertical load, and more friction yields more lateral grip. Despite its effects on the right-front tire, primary concern is going to be placed on the rear tires with crossweight adjustment: Turn in handling is influenced heavily by right-rear loading and turn exit handling is influenced by left-rear loading. More right-rear loading is more oversteer, more left-rear loading is more traction. Unlike most things we look at, this is pretty straightforward.

If we look at Atlanta, where the Pro race was held this week, we have a track with an extremely low amount of grip but extremely high speeds. Static crossweight in our Cup cars was over 62% across the board and, believe it or not, that's still probably a little low when realism is considered. Why it was that high is pretty simple: we needed traction and as much as we could get. When we take the cars to grippier tracks such as Charlotte and Kansas, we'll see it much lower since we don't need to generate the grip from the chassis. Speeds are also much higher at those tracks, meaning the car's going to want to keep going straight so we'll need a little help turning through lower crossweight.

Adjusting the crossweight is not simple, but it's not necessarily super-difficult if you take your time and look at what's going on. Since crossweight is adjusted through the spring perches (or shock collars) by altering spring preload at each corner, we also need to keep an eye on ride heights and the various torque preloads in the garage. In most cars, this is going to be the sway bar preload values but the NASCAR-style stock cars will have the rear end housing preload thrown in the mix as well. If any one of these things changes and is not corrected, then two separate settings will be changing the car's handling instead of just the crossweight. Whenever I'm helping someone learn how to do this, I always suggest disconnecting the sway bar before making changes. Preload will still have to be corrected, but it keeps things simpler by preventing the tech inspection system from causing everything to start flashing. The steps involved are simple:

- Note the starting bar preload
- Disconnect the sway bar
- Make your crossweight change (equal number of "clicks" for each corner), click "Apply"
- Once the car has settled, reconnect the sway bar
- Adjust the bar's preload back to what it was before
- (If necessary), correct the rear-end housing preload to legal limits

And that's it. If you followed those steps then you should have a new crossweight value with the same bar preload and ride heights that haven't changed very much at all. For a quick reminder on the perch and collar settings to change crossweight, see the images below:

[increase CW]

[decrease CW]

Crossweight is likely going to become a much more varied setting in the garage going forward due to the added degradation on the tire changes with the Top-3 NASCAR series. In the 2017 NASCAR Class B series I ran 55-56% crossweight on my car at nearly every race on the calendar, tuning the balance through track bar height and ballast. There were a few races where I experimented with rear spring configurations toward the end of the season, but the crossweight was largely kept constant because I didn't want to re-set bump stop gaps every week. The new build could put more emphasis on tire management, meaning spring choices are going to be more critical to long-run performance. Finding the arrangement that helps you maintain tire life throughout a run would mean keeping that arrangement from race to race, something seen *a lot* in the real-world. This would place handling in the hands of weight distribution and shocks which is, go figure, more realistic than changing springs for handling.

The only other important effect of changing crossweight is how it affects deflections in the front suspension, specifically bump stop gap and coil-bind timing. There's a fun saying "You could coil-bind any spring with enough crossweight", which references the increase in deflection on the right-front spring for increasing crossweight. Similarly, decreasing crossweight will increase deflection on the left-front spring. You can use this to your advantage to fine-tune splitter height in the Truck for coil-binding as well as fine-tune bump stop gaps to control corner loads. However, if you'd prefer those to stay the same, you need to make sure you keep the same amount of gap before and after a crossweight change. For bump stops this is dealt with using packers, but for coil-binding you may have to change the spring entirely to one with more or less travel.

Tie-Downs

We've waited a long time for an increase in shock forces and we finally have it. The forces on the Xfinity car and Truck are much higher than what we had before, and the Cup car has what seems to be a mild increase. Interestingly enough, it also seems like we have more range than is necessary! Personally I like this since it means we have to find the optimum setting instead of pegging the shocks at 32/32 and running it without a care.

None of the tracks we run on are perfectly smooth, meaning the suspension still needs to do some work (even at Kansas) to keep the tires in contact with the track. This not only means soaking up bumps with compression but also allowing the wheels back down to get ready for the next bump. Whenever the wheels don't get released after a bump we wind up with a car that starts bouncing uncontrollably, and that's *exactly* what we've seen with too much rebound in these cars. For shorter races and runs, we can get away with this because of the aerodynamic advantage of a low and steady splitter height. For longer races and lengthy green flag runs the tires need to be abused as little as possible and allowed to go with the flow, so to speak.

Exact numbers are still a mystery, and the extra rebound allows new setup arrangements that we've never been able to run before. What works and what doesn't will be largely dependent on what springs you choose and how bumpy the track is. So, while we are still in the ignorance stage of this build, it's pretty safe to say that "32 Rebound" is limited in its uses these days.

Now for some fun stuff…

Season 1 at the Commodore's Garage

The winter season offers a unique situation where none of us at Gale Force Racing are actually racing. In years past these months have been filled with baseline testing and diagnosing issues that cropped up through the year. This year we're going to do something different: Alex Scribner, Jeffrey Parker, and I are going to crash the party in the Class C Trucks this winter. Instead of just racing in the series, however, we're going to document the setup process in my truck, with the majority of the details written up the in the *Garage* articles over the next few months. Starting next week at Myrtle Beach you'll not only be able to see what's going on with our trucks, but you can run the same thing in yours while learning what went in to every decision from the sway bar to the track bar. Hope to see you on the track!

Commodore's Garage
Article #31 – Starting The Truck Project
by Matt Holden
December 19th, 2017

[truck image]

Last week I teased the next batch of articles that will cover the setup process of our own in-house Class C Truck over the Winter season. Since this is week 1 and nobody on the team has raced a truck in over a year and a half, we had to start from scratch. Since the first week is at Myrtle Beach, we also ran into the problem where none of us have raced there before, so we had that to contend with as well. Despite that, we were able to pull from our Class B Xfinity notes and eliminate most of the problems in the truck to get ready for the first week. Do we expect to win with this? I don't, but anything's possible!

As of writing, the truck hasn't been raced yet, so we don't know where it is on pace and we don't really have a grasp on how it will behave through the events. Still, that's part of the process, so if we can get out of Myrtle Beach with information and notes on what happened, we can take that straight to IRP and hopefully implement updates to the chassis for that race due to the similarities between the two tracks.

Front End Configuration

The options for the front end at Myrtle Beach were quite varied, so we had to look at the track characteristics to see which we wanted to prioritize and place higher on the “give this a whirl” list. Without bump springs in the truck, we’re really limited to four configurations possible for any track configuration:

- Double Coil-Bind
- LF Coil-Bind
- RF Coil-Bind
- Conventional

Each of these options has inherent advantages and disadvantages, most dependent on the track itself. We knew the track was relatively slow, it has a rough surface, and it’s very flat. All of these played into what we wound up choosing for the front end.

We were able to eliminate the Double-Bind option immediately due to the track’s rough surface which would cause a lot of headaches in tire life. The solid suspension in a double-bind suspension is great for speed and splitter control, however we knew from the Class B series that it was hit-or-miss on tire life. With the new build seeming to shine a spotlight on tire management, we decided not to go this route.

A conventional setup makes a lot of sense for such a bumpy track since it keeps the suspension moving at all times. Plus, the extra rebound in the front shocks seems to allow larger springs to be used without the front end rising up too far on the straights. We decided this would be a last-ditch effort if we couldn’t get the other two working.

Lastly, we have a single-bind front end configuration, which we’ve actually used before in the Class B series in 2016. When the car had bump springs added, everyone on our team bound one spring but ran a bump spring on the other so we had a *lot* of notes to pull from. We only ran these at Texas and California before figuring out how to re-work the cars for double-binding, so our notes covered bumpy tracks. Looking through these, the notes showed that cars with LF-bind setups had issues with bouncing excessively at both tracks, while the RF-bind cars didn’t have such problems. A quick test in the truck confirmed the LF-bind setup as being a little too bouncy for us to commit to, so we threw a 300lb/in spring in the right-front, a 600 lb/in in the left-front, and got to work.

Pogo Stick - Solved

We ran into some issues early with bumps towards the center of the corner, especially down in Turn 3. This was bad enough to cause the front end to lose grip and slide, evidenced by a loss of lateral G-force despite increasing steering angle. We then ensured the splitter was not touching the track at all (which could set off an uncontrolled oscillation in the suspension) and isolated the suspension to sort out the problem. We also compared damper position traces to ride height traces to see if either suspension was moving *with* the bounce and discovered that they weren’t. In fact, both suspension systems appeared locked in place through the whole ordeal. Not an easy fix, but one we’ve seen before. To work through the issue, we removed all rebound from the left-front shock (to eliminate tie-down issues) and went through these steps:

- *Spring Rates:* Changing the spring rate on either corner would alleviate or exacerbate the bounce if it was being caused by the spring. Changes to the left-front spring did not alleviate the issue at all, however increasing the right-front spring appeared to reduce the severity in the center of the corner. We settled on a 320lb/in right-front spring to allow it to open up without being bound solid through the bumps. This also required a reset on the ride heights (the spring allows more travel than a 300) and only one extra click of rebound to keep it where we had it.
- *Bar Diameter:* We swapped the solid link for a chain to keep sharp impacts from the left-front wheel from getting to the right-front wheel. It was nice but didn’t help…at all. A reduction in bar diameter stabilized the front end a little bit, but also produced no change in right-front height but huge changes in left-front bounce height. This was a big deal, a sign that whatever was happening stemmed from the left-front suspension.

What finally worked was a reduction in crossweight to preload the left-front spring more and keep the truck a little more solid. We had started the crossweight at 57% but dropped it to 55%, a change that eliminated the repeated “chirp” from the left-front tire, a sign that it was leaving the racing surface, and was worth around 6 tenths, likely from just having four tires on the track instead of three! From there, we simply set the nose height to just touch the track on fresh, low-pressure tires but have it clear the track after a few laps. Nose weight was set to simply balance the truck with these settings, and we wound up with 51.0%.

Shocks, set to just barely hold on to the front end, are 10/25 on the left-front and 20/15 on the right front. Splitter lift out of the corner is perfectly fine at a track like this where corner speed is going to outweigh the top speed, plus having less rebound on a bumpy surface will help with tire wear.

For those interested, we basically carried camber and caster values over from Class B coil-bind values, setting the left-front as high as possible and right-front to where the Inside and Middle of the RF tire match. Your mileage may vary, but this produced -4.5° of camber. Caster was 7.0/8.5, and the front tires are toed-in slightly.

Rear End

The rear-end is actually where we had no problems at all, since this is fairly straightforward. Short tracks offer the use of a rear sway bar (which we used religiously in the days of the COT), but we opted to leave that in the hauler. Instead we went with a pretty simple rear spring package, a 200lb/in left-rear and a 650 lb/in right-rear, setting the rear heights to put the skirts just above the track at full-load. The left-rear shock is a high-rebound shock with 9 compression and 17 rebound to help de-wedge the truck under braking, but ideally we'd like the rebound to be around 10 or so since this is a bit of a crutch. The right-rear shock is a low-rebound shock with 19 compression and 2 rebound to aid in turn-in when rolling off the brakes. Rear toe is full skew (4/16) on both wheels, mostly because we would rather go with what we know right now instead of running a straighter amount to help drive off from the tight corners. Hopefully we can get that worked out by week 4 at IRP as well. The track bar was full rake, as high as we could get it, so 12" and 15". Truck arms were at the bottom.

Pressures

The pressures are, admittedly, a shot in the dark. Nobody on our team has raced this truck in a year and a half, so we went with a comfortable set of pressures to suit the rest of the setup. The left sides are set to 14psi to put the center temperature in good agreement with the outer temperatures, and the right-sides are 39psi to do the same. Test runs are showing the buildup pressures are showing slightly tight, with the left-rear and right-front building up more than the other two which is expected. Should they start showing issues during the races, we can change them pretty easily during a pit stop.

Driver Tuning

This being the first race on the schedule, the drivers we have will all be on a very similar setup. The problem with this is that Alex, Jeff, and I all drive *completely* differently from each other, so one of us will likely do quite well while the other two just exist on the race track taking up space. Luckily, we can do quite a lot to tailor the car to each of us with just shocks and crossweight and noseweight adjustments. The springs, however, will likely go un-changed before IRP.

It's very likely that the three setups start going in a different direction after Myrtle Beach as we figure out what each of us wants the truck to do. Different crossweight/noseweight combinations, different shock packages, and even rear spring choices could send each of us in a different direction, but that only helps each of us in the long run by having a notebook to fall back on. Considering how this truck was built off of notes developed from the Xfinity car, that can only be a good thing!

Post-Race

We should have gotten at least one race under our belt by the end of the week, likely Thursday or Friday, so the article following the races will be about what we've found and what needs to be addressed. The schedule for the winter season consists of three flat short tracks (Myrtle Beach, IRP, and Milwaukee) that this chassis will be used on, so developing the chassis quickly and effectively will be paramount.

Of course, there is the chance that it's trash. If the setup is indeed a total bust, we just rebuild it based on what went wrong during the race by making notes of how the setup didn't perform well relative to the rest of the field. Simply saying "no good" and starting over with no information won't help at all, but knowing what went wrong will help to eliminate many issues and give a good starting point going forward.

But hopefully it's not trash!

Commodore's Garage
Article #32 – Myrtle Beach Results

by Matt Holden
December 26th, 2017

The bad news up front: We were never able to get a race in at Myrtle Beach! In hindsight, however, it probably wasn't the best idea to start all this during the week before Christmas. Still, we did manage to get into a few sessions with regular series drivers and managed to gather quite a bit of info on the setup's strengths and weaknesses. The setup file will be posted as it was when the week ended on our Facebook page, but as I've stressed before: This series is about working through the setup process and I do not expect the setup to work for everybody. Some may like it, some may hate it, but the idea is to give everyone the knowledge and tools to take something and make it work for them!

First Race Run

The first run made in the car in anger was a 20-lapper to gauge the fuel usage per lap and how the car worked the tires in over a lengthy run. Since this was the first run made with other cars (drivers who I know would be competitive), it was essentially a shakedown run because the pace necessary to run in the race was shown to us for the first time. Once the 20 laps were finished, I calculated the fuel usage:

- *Finished 20 laps with 16.1 gallons remaining, total burnoff is 2.6 gallons.*
- *2.6 gallons ÷ 20 laps = 0.13 gallons/lap*

This fuel burnoff rate would produce about 143 laps on a tank, which is longer than the race itself. It's possible that this burnoff rate would get better or worse depending on the conditions, and we have to remember that this includes a lap where I left pit road, so a more realistic number is between 135-140 laps. For this race it's a relatively useless value since the race is shorter than what can be run on a full tank, but it never hurts to have this in the notebook for later.

Next, we need to take a look at the tires:

[tire image, run 1]

Note that there are no wear numbers because it's unnecessary at this point: A short track oval with left-hand turns will always show higher wear on the right-front. Period. What we do want to look at here is how the tread is heading up from camber and the pressure buildups because they'll give us a good idea of which tires were loaded and which ones weren't. Immediately, we can see that there isn't enough left-front camber or right-rear camber, both of which would be free speed! The pressure buildups also show a slightly high crossweight for what we were shooting for (equal buildup on each side), but we'll address that later.

The biggest issue, by far, was a "loose-in" condition. It could be partially attributed to the brake bias (since releasing the brakes would settle the truck down), but it was also brought on prior to braking by simply jumping out of the throttle quickly. This is a sign that it could be shock related as well, so changes to both were made for the second run.

Once everything was looked over, these are the changes made for the second run and descriptions for each:

- LF Camber from +7.2° to +7.9°: The inside of the left-front tire was showing 8° cooler than the middle. If this didn't cure it, we'd need to drop the pressure in that tire.
- RR Camber from -1.5° to -1.7°: The Middle was slightly warmer on this tire and more rear camber is a no-brainer when you can afford to use it.
- Brake Bias from 62.5% to 64.5%: Fix brake-related oversteer
- LR Rebound from 17 to 15: Reduce the left-rear shock rebound force to add crossweight under braking

NOTE: I intentionally left the crossweight where it was despite finding an issue with it on the tire data. Adjusting the left-rear rebound will have an effect on dynamic crossweight under braking and deceleration, so adjusting crossweight at the same time could have a combined effect with the rebound adjustment. Adjusting the rebound alone would show how much the crossweight changed, or even if it would be enough to show up on the tire data.

Second Run

The second run was shorter, only 6 laps this time, because it was very obvious that the loose-in condition was gone completely. The truck, at this point, was quite stable and would easily be in a position where it could hold its own in a race. Since there was time left in the practice session, we'd simply use it as a test session for IRP in a few weeks. First, a quick look at the tires:

[Run 2 tires]

We can see from these tires that the left-front tire could stand to be a little lower on pressure, which isn't crazy considering how small Myrtle Beach is. Curiously enough, it also looks like the adjustments may have matched up the right-side tires in terms of loading, which is a very good thing. All that's left is to work on cooling the left-rear off and heating the left-front up.

Final Runs

To check for possible solutions to the left-side tire heating imbalance, I took some time at the end of the practice session to test various ideas to implement next week at IRP. Some changes to the rear spring pairing, specifically a much stiffer right-rear to match our old Xfinity springs, didn't match well with the line into turn 1 due to having to be on the brakes while steering into a flat corner. It upset the rear-end too much, so I rolled that back. The other rear-end option would be to run a rear sway bar, but I didn't have the time to reset the right-rear spring to get the same travel rates with a softer coil-spring and a sway bar. Interestingly enough, however, I still have notes from the 2011 Pro Series and 2012 NASCAR DWC season where Nick Ottinger used a rear-sway bar at small, flatter tracks so it would be much easier to hit the ground running with those notes when we go to IRP.

What took up the most time, however, was dropping the crossweight and increasing the nose weight to see how much farther I could take it and if it was possible to iron out the bumps completely. Long story short: Yes. The bumps vanished around 52% crossweight (down from 54.6%), which I had to balance with a bump in nose weight from 50.6% to 51.3%. This produced a much freer truck (as you'd expect) that would roll the center much faster, but produced an interesting effect where the rear was squatting onto the right-rear corner on exit and basically "wobbling" its way off the corners. It got rid of the bumps, but dealing with the wobble was a whole new issue in itself.

IRP Preparations

So, despite not actually being able to race, we did get quite a lot of useful information from the longer runs made while building the setup and in practice. The two initial runs in practice showed that the setup would respond to adjustments extremely well, but adjusting by a large amount could produce some unforeseen problems like the "wobble". As for speed, the setup was decent, but was around 3-4 tenths off of the drivers I know would be front-runners so there's room to work. This speed could likely be found in simple aerodynamic tweaks, since the large bumps were causing the splitter to hit the track a few times, and the rear was higher than I'd like to have had it with the softer rear springs. Running around the faster cars showed that my setup was running well under braking and on throttle, but was lacking in the center of the corners.

We have no plans to run Talladega since it would be a one-off event for the season, so the next track (Indianapolis) will see this setup return. While I've never raced anything at Myrtle Beach, I *have* run modifieds and late models at Indianapolis so I actually have an idea of where to go before rolling into town. I would definitely like to run the crossweight down where I had it at the end of practice, but I must figure out how to remove the right-rear squat on corner exit. This will likely come in the form of a stiffer right-rear spring and a higher-compression shock. To prevent the truck from becoming overly-loose in the higher-speed corners, I'll need a stiffer left-rear spring to go with it, so I will fall back on the 325/1200 combination I used in the Xfinity series for short tracks in 2016. Other than that, the front-end will likely stay the same, but I do plan to test a more aggressive left-front configuration with a high-rate coil-bind spring (probably 400 or so) in an effort to get the splitter down.

This setup, as it was at the end of the racing week, will be posted on our Facebook page when this article goes live. If you want to take it for a spin or use it as a starting point for IRP, feel free to do so!

Commodore's Garage
Article #33 – Truck Project: Lucas Oil Raceway

by Matt Holden
January 8th, 2018

Week four of the Season 1 Class C season is in the books, and this time we actually got a race in with the Truck Project! Last Thursday Alex Scribner, Jeff Parker, and I all got into the 9:45pm Eastern race and were able to shakedown the truck in a full-bore, "this matters" type of situation and we walked away with a lot of great information. While we didn't quite put the smackdown on anybody in the event, we still managed a 4th place finish with Jeff, I finished in 6th just a few feet away from a top-5, and Alex finished 11th after losing an engine following contact with a car that was multiple laps down. That result is extremely unfortunate, as Alex was in 3rd place when his power plant let go. Oh, what could have been…

I put the major changes from Myrtle Beach into a short YouTube video to show how simple spring changes and crossweight changes are without having to make a thousand other changes to correct problems, so if you want to go see that head over to our Facebook page via the link at the end of the article!

Practice Sessions

The setup unloaded quite well with the changes we made, but it still showed some minor teething issues. The worst of which was the bouncing issue that had plagued us through the entire week at Myrtle Beach, compounding into more problems due to the higher speeds producing a much heavier dependence on aerodynamics at Lucas Oil Raceway.

Something we fought all week was a push in the center, likely triggered by the left-front bounce on corner entry. I was able to get rid of most of the bounce by running the crossweight even further down, but something in the way we have it configured wasn't really agreeing with the track characteristics. We never got it ironed out by race time, however I managed to cure it with some in-race adjustments and get some grip in the center of the corner. Most of the adjustments I made in practice were simply to "crutch" the problem and mask it for just this track:

- **More preload, more bar asymmetry:** The left-front corner was hiking up in the center of the corner which would have contributed to the push-loose condition we were seeing at both ends of the track. The extra preload and asymmetry simply pinned this corner down to get a flatter splitter attitude and free up the center of the corner. This also required a drop in front ride heights due to the extra vertical stiffness (heave spring rate) from the extra bar asymmetry.
- **LF spring from 400 to 420:** This added a little more travel in the left-front corner to allow this corner to drop more under load. This also adds a tiny bit of rate to the suspension to try to absorb some of the bouncing.
- **Lower Crossweight, Higher LR spring:** Since we had the crossweight so low to iron out the bumps, we needed to increase the left-rear spring rate to raise the dynamic crossweight to something that was more balanced. This provided stability in banking transitions but likely hurt our long-run performance a little bit.
- **Lowered Track Bar:** I dropped the track bar equally to give the truck a little more stability for the higher-speed, low-banked entry at LOR.

Something I also tested in the practice sessions was a "double-bind" front-end similar to what Alex used in 2016 on his championship-winning Class B setup. While the speed was there on a short run, I wasn't able to get the long-run performance to match the setup with a stiffer left-front spring so I shelved it. It could make a return tour later, however.

The Race

Our race was fairly eventful, but it was not without minor issues that turn out to be common with first-try setups. Overall, the balance was good, long run performance was fantastic, and the handling swing wasn't overly severe throughout the run. The race started with two short runs, then a very long green flag run, and a final restart with about 10 to go. Our largest issue by far was an overheating problem. I started the tape at 55%, which is what we would have run the Xfinity car at a short track in 2017. For whatever reason, the Truck doesn't like this much tape *at all*, so all three of us spent the majority of the race running at a pace that kept the oil right at the maximum safe temperature. This put a huge damper on how fast we could run, so it's not clear what we would have been capable of doing had we been able to run as hard as we wanted to during the long green-flag run.

Jeff chose not to qualify and started 11th, but made his way to the top-5 in the first 32 laps (which included two cautions, by the way). His pace late in the run matched the eventual winner, but the middle of the race was dictated by the overheating issues. During the long run, Jeff mostly ran a few slower laps to cool the engine down followed by a few faster laps to make up time. The lap times when he wasn't cooling the engine off were very good and it's very possible that, had he qualified, he would have been running much closer to the leader for most of the race.

Alex had a similar race, however he spent his entire race in 3rd place nursing his engine. Once his engine started overheating and he had to back the pace down, he said to me "The pace and balance is great, I could compete with these two in front of me if it wasn't overheating". Unfortunately, we never got to see what Alex could have done because his engine problems got even worse when a lapped car locked up all four wheels on corner entry and slid up in front of him in the center of the corner. The damage was enough that Alex's engine expired not long afterwards due to even higher water and oil temperatures.

My race was mostly spent feeling out what the truck was doing instead of racing for anything. To get a little extra speed at the start of the race, I followed Alex's suggestion and raised the left-side pressures to 15psi. This actually wound up making the bouncing worse, so I dropped them back to 12psi during the first caution. I also made a crossweight adjustment (+0.050 LR, -0.025 RR) during this pit stop to try to get more of the bumpiness out of the chassis. It worked extremely well for about 30 or so laps before the pressures built up and the bounce returned, but by this time all three of us were overheating and a bouncy front end was the least of our problems. I did get to really hammer the setup on the final restart and see what it could do since the caution placed me one lap down with only one other car. It was quite good during that final run, and I even set my fastest lap of the race on lap 102 of 110. I haven't pointed this out to him, but my fastest lap was faster than Jeff's. At least I'll know immediately if he's read this article!

All things considered, it was a great first showing for the project. Our times were fairly competitive, and according to both Jeff and Alex, it could have competed for the win had the overheating issue been found earlier in testing. This setup should make a showing again at Milwaukee and Iowa, however I'll most likely swap which front corner will bind due to the much higher demands for aerodynamics at both of those tracks. We'll have to look elsewhere to cure the bumps, but we have enough of a notebook to know where to look and what to do in order to cure all that.

Race Tire Data

Here are the two tire sets I took off of my car during the race and a look at what could be going on. I didn't record the data from the third set of tires, which I regret tremendously because I pushed those tires the hardest.

[set 1]

This is the set of tires I started the race on. It's worth noting that the left side tires began at 15psi but most of the testing was done with 12psi, so the left front is showing overinflated. We can also see a slight issue with a little too much crossweight (might be good for that track though), but overall the tires look quite good.

[set 2]

The second set of tires is a little more telling, and this has a lot of good info we can take to Iowa and Milwaukee later in the season. To start, the left-front was dropped to 12psi and the left-rear dropped to 14psi on the first stop and I also made the crossweight change mentioned earlier. Interestingly, this crossweight change had almost no bearing on the temperatures for the tire averages, noted by the 10° difference in the Positive and Negative Cross Average for both tire sets. The right-side tires are cooler by a large margin, but it must also be noted that the longer run wore the tires more and this alone will cool the tires. This is a sign that the long-run push that developed very late is likely somewhere else in the car and crossweight won't have a tremendous effect on that.

California Truckin'

Next the series moves to California where we must build a new setup entirely to deal with the different chassis, aerodynamic dependence, and the speeds we'll see at a 2-mile track. We can still carry a lot of what we learned at LOR and Myrtle Beach to the big track, however the way everything else is configured will change. We'll undoubtedly switch to a left-front bind setup, but whether or not we bind both springs is yet to be seen. With California's surface being very bumpy and low-grip, we may opt for something similar to LOR where we have a stiffer right-front spring that will still bind at full load, but can expand and cushion the bumps and let the tires work without the suspension being too solid. As for the rear, it's very likely that the springs used at LOR will be used for California and changed based on what we see when it hits the track. Our current plans are to run the same time window that we ran this week, hope to see you there!

Commodore's Garage
Article #34 – Truck Project: Auto Club Speedway

by Matt Holden
January 12th, 2018

Somebody needs to mark the calendar because a sign of End Times arrived last night after our Truck race: Somebody was happy with the way Jeff Parker raced him.

All jokes aside, it's safe to say that our race at Auto Club Speedway was a one-and-done situation because it's highly unlikely that we'll be able to equal the result we had in any subsequent race. Consider the situation: The race had a 3822 SOF, was populated with the names we're used to racing against in the Class B Series and who have run the truck a *lot* more than we have, and we hadn't raced any trucks on a track larger than a mile since mid-2016. We started 20th, 22nd, and 24th but finished 4th, 5th, and 9th. Yours truly was bringing up the rear again, but I have a very good reason for this which we'll address later.

As I've mentioned in previous *Garage* articles, the three of us have never been able to run the same exact setup. So, following the race at Indianapolis, we each went our own way with the California setup but didn't wind up straying too far from each other. As a matter of fact, we were all within a tenth of each other in practice so that's pretty impressive in my opinion. We're going to focus on my setup alone, however, but I will point out a few points where I know my teammates' setups varied.

Setup Background

California is our first Intermediate track of the year, so the setup needed to be more heavily focused on the aerodynamic aspect instead of mechanical grip like we did at Indianapolis and Myrtle Beach. This means the heights were much more crucial and keeping those heights from changing a lot in the corner is vital. At Indianapolis we could sacrifice a bit on the body attitude for the sake of keeping the tires in contact with the track (this is why we spent so much time curing the bounce), but not here at California when we're over 180mph!

Despite the shift in thinking behind the setup, it wasn't dramatically different from the IRP setup. The big right-rear spring stayed, the rear spring split almost went unchanged (ran a little softer at California to help rotation past center), the track bar was still right-side-high for extra skew, and one of the front springs was bound while the other was not. The major difference, however, was *which* spring was bound: on my car, I bound the left-front to put a hard-stop on the left-front height just to make it easier to control the splitter's height, something that was a bit of challenge at Indianapolis last week. Alex went the opposite way with his setup, basically running the same springs reversed so that his right-front spring bound and the left-front wasn't. This isn't particularly surprising, since we did the exact same swap in the Xfinity cars at California in 2016. Just like in the 2016 race, he and I eventually raced around each other for a decent portion of the race. Fine-tuning the right-front travel was fairly simple using the sway bar diameter since it was behaving as a torsion spring. Again, asymmetry was used to flatten the splitter attitude and the bar was lightly preloaded.

Crossweight Adjustment & LF Spring Rate

The setup was largely plug-and-play for California. Once the heights were set where I wanted them, it was basically an exercise in crossweight adjustment to get the balance right. This wasn't quite as simple as it was in the previous weeks and I wanted to dedicate part of the article to making this change with a front-end setup like I was running. If you haven't seen the quick-and-easy way to make crossweight changes in the video I posted last week, head over to our Facebook page so you know how everything should work.

For aerodynamics, the left-front corner is extremely crucial to the car's performance. If it gets too high there will be a lot of drag and a major loss in downforce, so keeping it from moving around during the setup process is very important. For a coil-bind car like I had, the left-front spring is serving as a "hard-stop" for the left-front corner. Whenever I made a crossweight change to the truck the available travel in the left-front spring also changed: travel is reduced when crossweight was lowered, travel is increased when crossweight was raised. The same (but opposite) is true for the right-front spring, so it's important to know this is going on when you make a crossweight change.

I got around this by increasing the left-front spring rate by 10 lb/in every time the available travel started to decrease enough to raise the left-front corner when the spring was bound. This is something you can see in telemetry by overlaying two runs together, and simply increasing the spring rate when the splitter gets too high. This is a little counter-intuitive since it would make more sense to run a softer spring rate to allow the car to drop, right? Well things get a little funky when we're binding springs, so the way springs are designed has to be considered in this situation. As a spring increases in rate it will typically have fewer coils in the spring, and fewer coils in a spring means more travel so whenever we have a spring completely bound like I did at California, what matters the most is how tall the spring is when it's bound (known as "block height"). When the spring is binding with the front end too high, we need a spring with a shorter block height, so we need a higher-rate spring which will have more travel.

For a setup like Alex's, with the right-front bound, you would do the opposite for a decrease in crossweight adjustment. Lower crossweight would put less preload on the right-front spring and cause it to bind later, possibly causing the chassis to hit the track. For this case, you would need to *decrease* the rate to raise the front end from the track.

The Gripfest 500

Our race was not immune from whatever is causing the Truck races to be run in extremely cold temperatures. We tested in 70°F weather and our race info page said 70°F as well, but we were about a second faster than testing. It wasn't until the results page was finalized that we saw the weather was actually 45°F! Incredibly, this is not the coldest race I've seen this week.

With temperatures that low, we knew the racing was going to be extremely fast, and the majority of the opening laps for a tire run were wide-open if you wanted to be competitive. While Alex decided he was going to jump right in head-first from lap 1, Jeff and I decided to lay back from the field and give ourselves plenty of room to avoid the massive crashes that ensued in the first half of the race. It worked quite well since we had quite a lot of room to slow down and let the wrecks happen before we even got close to them, but the downside to this was that we didn't know how the trucks were going to handle when pushed harder to be competitive.

Jeff and I decided to join in the fun later in the race past halfway and quickly discovered that our setups were just fine. I managed to get past Alex on a restart and successfully annoyed him for much longer than I expected without losing time on the lead pack. The three of us simply bided our time and let everyone else run their tires off before making passes and working up to our eventual finishing positions. Jeff was quite aggressive and found his way up front, but Alex was likely kept from a better finishing position by my stubbornness. Eventually Alex did get around me when I parked myself in a remarkably bad position that he and one other driver capitalized on. Once we were settled in, the field formed a line and we hugged the apron like a low-banked superspeedway.

The balance was quite good and didn't change throughout the final run when it was driven the hardest. Checking the tires after the race, there isn't much that didn't go wrong on the setup at all:

[tire set]

If anything, the setup still has just a little too much crossweight (evidenced by the pressure buildup and temperature averages) but it's not something to be worried about with California's bumps and low-grip surface. Plus, the track was 25° colder than we practiced in so having the setup show a little tight isn't surprising. Had I made a more aggressive run early in the race, I would have been able to decrease the right-front pressure by 1psi and the left-rear by 0.5psi, giving a much more balanced pressure buildup on the car and possibly giving me enough grip to make some passes late in the race. Still, it's not bad at all.

One thing that could be changed is how much load is on the right-rear. It does show a little cool, the only tire not in the 230s across the two inside zones. This would be a simple change by just reducing the right-rear spring rate. I know Alex ran a softer right-rear spring than I did and his tires showed a much better right-rear temperature. I'll look to drop the rate at Michigan to try to balance the temperatures out a little better.

Milwaukee

Next week we'll go back to the short-track program at Milwaukee but, instead of taking the super-short setup that was used at Indianapolis and Myrtle Beach, we'll be unloading with something closer to California. The long straights and higher-speeds at Milwaukee will bring aerodynamics into play, so that needs to be a priority. It will mostly be a fact-finding mission for Iowa, since those two tracks will have very similar speeds and loads. Again, we'll be planning on the 9:45pm race on Thursday just to keep things consistent.

As with Indianapolis, I've posted the California setup I used in the race on the *Commodore's Garage* Facebook page so go check it out if you're interested! Also drop us a like and follow the page so you can always be kept up-to-date with the Truck Project. Take care!

Commodore's Garage
Article #35 – Truck Project: The Milwaukee Mishap

by Matt Holden
January 25th, 2018

I knew Milwaukee was going to be a difficult race to get figured out way back when the foundation for this endeavor was laid out. It's an old track, few series race there anymore, and it rarely gets placed on any of the NASCAR iRacing Schedules meaning there aren't a lot of people who have Milwaukee on their accounts. It's also a track that we, as a team, haven't raced on since 2012 when there was an invitational Late Model league. Many unknowns hit us right in the face last week and, despite finishing 2nd, 3rd, and 4th in the race we ran, The Milwaukee Mile turned out to be the race where we simply "missed" the setup quite severely. It happens, and I'm glad it did, because this is often the point where everyone wonders what to do next with their setup. Do you throw it out and build another one? Do you take a look at it and see what you can fix? The answer might surprise you here, and hopefully you'll get an idea of where to go whenever this happens in the future.

California Dreamin'

The Milwaukee setup itself was originally what I ran at California: Left-front bind, low crossweight, big RR spring. Milwaukee's corners are high speed, relatively smooth, and we can get some speed from making sure the aerodynamic attitude is good while sacrificing mechanical grip. By race time, the springs went unchanged (except for a stiffer left-front for extra travel) and the only things that got swapped out were shocks and weight distribution.

The idea seemed pretty good at first, but it quickly showed its ugly side in the bounce that we simply can't seem to get rid of in this short track chassis. The same problem that existed at Myrtle Beach and Lucas Oil Raceway returned in force last week and the only thing that alleviated it was some mid-race adjustments, however it did not eliminate the bounce completely.

Unlike Myrtle Beach and Lucas Oil Raceway, where I ran a setup with an un-bound left-front spring (meaning there was something to soak up some of the bounce), I flipped the front end and bound the left-front spring to see if the same thing could be achieved. The short version: Nope. It had great short-run speed but the instability eventually manifested into a severe mid-run falloff that could have resulted in a far worse finish had the field we ran in been much larger.

The Race

Ben Welvaert filled the role of third car this week since Jeff doesn't own Milwaukee, and he ran the exact same setup I ran. Alex went with a front end similar to what we had at Lucas Oil Raceway, and since he was able to gap Ben and me fairly quickly it seems that was the better choice. Still, Ben and I pulled away from the rest of the field over time, so our car wasn't as disastrous as we initially thought.

The race went caution-free, leaving room for only one pit stop to try to settle down the bounce and make an attempt to get back to Alex on the second set of tires. At one point during the race Nick Ottinger, who was watching me, said the bounce looked "violent" and that was a bit of an understatement. I eventually found a way around the corners that minimized it, primarily staying away from the inside curbing to prevent setting off the oscillation, but the downside was that it was a slower line than what Alex and race-winner Brent Day were able to run.

I've mentioned many times the importance of keeping yourself from rebuilding a setup for every race and that statement shined bright when it came time to adjust the car for the pit stop in our race. I knew the adjustments I'd made at Lucas Oil Raceway that cured the bounce completely and, since this setup was an evolution of the LOR setup, I knew the adjustments that should help to do the same thing at Milwaukee:

- **Both left-side tires were dropped to 12psi:** I started the race on 15psi because of the higher speeds compared to the other flat tracks, but it was pretty clear that this was too much again. I started the left tires at Lucas Oil Raceway with 15psi and dropped them to 12psi and it worked there, it also worked at Milwaukee. This doesn't really do much for contact patch or anything like that, it's simply dropping the tire's *spring rate*. Since I was coil-binding the left-front, this tire spring rate was the only thing that was soaking up bumps on the track, so dropping the spring rate will help that.
- **Wedge bolts to +0.075" LR and -0.050" RR:** This is just dropping the crossweight even more to put more preload on the right-rear spring and try to keep the left-front pinned down instead of hopping off the track. These values are different (higher LR adjustment) because the left-rear spring is so much softer than the right-rear spring. Adjusting both by the same amount would alter the overall rear ride height and mess up the aerodynamics.
- **Tape from 46% to 50%:** The engine was running cool. Not anything exciting going on here.

[tire data]

The tire data basically confirmed that I'd made the right adjustments but whenever there's a problem as huge as the bouncing in my car, handling and tire data is not particularly important. I was more concerned with getting all four tires to stay in contact with the track than I was getting everything balanced properly. The only thing that is interesting is the right-rear tire being the coolest on the car, which can be addressed very easily in the future.

The Future

So the setup was junk and, had the field been full-sized, would have likely finished very poorly. Do we throw out the setup and start over or do we adjust the problems out of it? This is a common question in sim-racing, and most people will just start over at the next track. This setup, however, will be unloaded at Iowa instead of thrown in the trash bin, and the California setup will go to Michigan and Rockingham. It's easy to take a bad setup like this one and just throw it out, but starting over erases everything we've learned so far. More likely than not, any setup that gets built in its place is going to be built in the same way and the problems will likely still exist because nothing was learned that could overcome the problems.

[pic]

But let's say that I rebuild the setup and the problems *do* go away. I've gotten rid of the bouncy problem, but *how* did I get rid of it? With this way of thinking, there's nothing that explains what was causing the bounce or what went into killing off the problem, so I have no way to fix it in the future. It will simply return in the future and then I'll have no information on what to do outside of just starting over on a vicious cycle. I'll be the first to admit that it's very difficult to take a setup that was utter trash and keep working on it, but had we not done just that at Gale Force Racing, Alex Scribner would never have won the Class B Championship in 2016, and I wouldn't have been fighting for wins in the same series in 2017. Both of us had races where we would have been perfectly happy to take cars and put them in tree-chippers, but we didn't. We took the problem setup to the next track and worked until the problem was gone. Then when the problem came back a month or so later, it was a matter of minutes until the problem was gone. Eventually we knew how to work on our cars to the point where weekly practices took 30-45 minutes and we were rolling on cars capable of winning races. Had we restarted every time we ran into a problem, we'd go back to square one every single time.

Options to cure the Bounce

There are a few options available for getting rid of this irritating issue that's plagued all our setups at short, flat tracks this season. Luckily, running California's setup at Milwaukee unveiled something very interesting in the way the Truck works: This bounce is only present in the Short Track version of the car. I basically ran the same car at both tracks, and only one had a bounce. Sure, California results in a lot more speed and downforce, but the hallmark signs of the death-bounce were not present in that chassis. It was smooth as butter.

With only three races left, only Iowa will run on what we call the Short Track chassis, and this track is configured in a way that allows multiple setup options to exist. We can flip it back to the LOR configuration by binding the RF corner and opening the LF spring up. We can run a fully conventional setup with some bigger springs and high-rebound shocks. Or, we can run the same thing and go up on the left-front rate even further to try to get the spring to open up and absorb the bouncing. There are many options, but all center around one thing: Get the left-front suspension to work better. Simply understanding what needs to be done is a huge step towards fixing it because it places a target to aim for. We already know from LOR and Myrtle Beach that running a larger left-front spring and un-binding it will reduce the bounce considerably, but it's now a case of figuring out which method will work best for each driver.

If we were going to another flat track again, such as New Hampshire, there is one change that can be made immediately to help the car. I mentioned the right-rear tire being cold relative to the other tires, and that's a pretty simple fix: Soften the right-rear spring. It's really not necessary to run such a stiff right-rear spring at a flat track like Milwaukee or New Hampshire since the vertical load is so low. Softening the spring would put some heat into the right-rear tire and possibly even help to subdue the bounce on the left-front when it was transferring over to the right-rear.

It looks like a tall order to try and cure a problem that's plagued every short-track car I've put on the track this season, but this was the point of the Truck Project in the first place: Demonstrate how to work through problems without having to rebuild setups each and every week and eventually wind up with a setup that works well everywhere and does what the driver wants. While it may look like we've gone nowhere in 4 weeks, the truth is that we really have a lot of information we can work from to get the truck working properly. It's not an overnight success story and it never will be, but it looks like the sun's starting to rise and we're headed in the direction we want to go.

Commodore's Garage

Article #36 – Truck Project: Out on a High Note

by Matt Holden

January 30th, 2018

When we left Milwaukee it would be an understatement to say that Alex and I were disappointed in how our cars performed. We finished 2nd and 3rd, which looks nice on paper, but when you consider that both of us got lapped it's not something to be proud of by any means. The bounce that fought all of us since Myrtle Beach made us look like pushovers. Considering all we'd done to try and get rid of it, it seemed like we'd carry that same problem to Iowa, get steamrolled again, and end the Short Track portion of our Truck Project at rock bottom with our hands up in the air in surrender.

That was not the case. Not only did we fix the bounce, we were competitive for the first time all season against drivers we knew were the class of the field. Instead of leaving the final short track with our heads down, we left in the best mood we've been in all season.

Curing the Bounce

Everyone who's followed this series since day 1 will know that we've found a bounce in the front end of the car every lap we've had on the short track chassis. It's forced us to run the races in a way that will save the tires and keep the tires loaded instead of being able to attack when necessary. Going to the final race of the season on this short track package, it was paramount that this bounce was eliminated. Iowa is unique in that it's bumpy in turns 1 and 2 while being very smooth in 3 and 4, so if the bounce was still around by the time we raced our options in turns 1 and 2 were limited to whichever line didn't set off the bounce. At a progressive track like this, having options and the ability to move around are vital to how you perform in the race. If you get locked into a single line because of car handling, it becomes extremely difficult to pass other cars because they can just run the line you're locked into and keep you behind them.

To start, I stuck with the California setup that was used at Milwaukee. I ran a shakedown run to make sure the bounce was still there at Iowa, which it was, then went over the options to settle the car down. There were a few options learned at previous tracks that we knew would work, and these were implemented into the truck immediately:

- **Right-front Bind:** The Indianapolis setup had the mildest bounce of the three short-track chassis up to that point, and it was run with a right-front bind front suspension. Plus, since the Milwaukee truck used a left-front bind and it was a disaster that needed to go immediately. I went back to the 300 right-front spring used at Indianapolis, and paired it with a 550 left-front spring to handle the extra vertical loads from the banking. I also didn't want it to bind, so a higher rate was necessary.
- **Lower Crossweight:** Normally, Iowa seems to like a higher crossweight for the extra bite when the track gets hot and slick. It's not flat-track numbers like I'd use at New Hampshire or Martinsville, but it's usually higher than an intermediate track.

However, it was well known that the lower the crossweight was on the right-front bind setup, the less bounce the suspension produced. The crossweight at Indy was 51.2%, the Iowa crossweight was 52.7%.

- **Left-side Tire Pressures:** It's impossible to forget how bad the first run at Indianapolis was on higher-pressure left side tires, so this was another case of "How little air can I put in them?" Minimum air pressure was the only option, so 16psi is where they were set.

These were implemented prior to the second run on the track. The majority of the runs afterwards were spent getting the heights correct and the splitter flattened, which was the usual song and dance between crossweight and spring rates. Once the truck was back down, the answer to the bounce was found almost immediately, and believe it or not, it came from Milwaukee's notes: During practice, Alex took the asymmetry out of his sway bar and commented that the bounce was gone, but he couldn't get the truck back to a flat attitude. Removing the asymmetry at Iowa had the same effect on the car. It took a lot of detail work to get the car flat in the corners, but this is eventually what got rid of the bounce.

The Asymmetry Problem

Knowing that this is what was likely causing the bounce issue, it's easy to say "just never run asymmetry, right?" Not necessarily, because we've used asymmetry in both the Xfinity and Cup cars to great success. The difference between those cars and the truck was that we knew how to use it in the cars, while we were just going off of what we'd used before on this truck. The problem lies in the nature of arm asymmetry itself, specifically with what happens when one wheel sees a bump while the other doesn't. Asymmetry shortens one of the sway bar arms without changing the length of the other, meaning the two suspension systems (right and left-front) will act on the sway bar in different ways.

Still with me? Good. For the Xfinity and Cup cars, we run bump springs which leave an active spring rate in the suspension even at high travel values. Even if the car is pinned down to the track, there is still suspension working on both corners. The truck is different, since coil-binding produces an *infinite* spring rate on whichever corner is binding. As you can probably imagine, applying a sudden force to a corner with an infinite spring rate will send it straight into the chassis and to the other corner of the suspension. At Milwaukee, I had the left-front coil-bound and was hitting the bumps with that corner first which shook the chassis and the suspension couldn't dampen the oscillation. At Indy, I bound the right-front (but still ran asymmetry) and the bounce was less severe. At Iowa, I removed the asymmetry and bound the right-front, getting rid of the bounce completely.

This doesn't mean asymmetry has no uses, because I still split the ride heights very heavily (which also produces asymmetry effects) it just meant I didn't have it implemented properly. Completely casting aside an idea because it causes problems for one configuration is not a great way to go about this, since it can work in other configurations. If you eventually get into the mind of eliminating options such as having ballast as far back as possible, no rear spring split, or no bar asymmetry, you're eventually going to get to a point where you can't break yourself out of the box you built around your own setup and you can't fix underlying problems because your mind won't let you go outside of the box. Alex said the asymmetry at Milwaukee kept his car flatter, and this was the biggest problem I had after removing it. We've used asymmetry in every single one of our team cars since 2016 because of how it keeps the nose flat, but I didn't write off using none at all because it could still be used in the right circumstances. When we go to Michigan tomorrow night, we'll likely have asymmetry in the bar because it worked at California, and my Xfinity cars will undoubtedly have it this year at most tracks. However we now know for certain that asymmetry, in the wrong configuration, can produce some nasty results. That kind of information is invaluable and will undoubtedly serve us well in the future.

The Race

The race, believe it or not, went caution-free despite a full field. We were also going to be racing against Brandon Buie, a former teammate of ours at Gale Force Racing, and someone I have a tremendous amount of respect for. Brandon, in my opinion, *is* the standard to which we were measuring ourselves against in that race. Anybody who follows the news on our team page might remember that I built a truck setup in 2016 to race at Gateway, and Brandon drove that car to victory with Alex running second. Now a year and a half later, it was a good opportunity to see where we stacked up. Plus, Brent Day was also in our race. Brent is the driver that lapped us at Milwaukee, so we literally had two yardsticks to compare ourselves to, and I think we did extremely well to finish 5th and 6th. Plus, Brent didn't lap us so that's also a step in the right direction.

A caution-free race is always fun since it really shows which cars are built for longer runs and makes the pit stop adjustments so much more crucial. I was planning on stopping around the race's mid-point, but leading up to that lap I was actually running Alex down for 5th place so I decided to drop to pit road just before he did to undercut him through the pit cycle. He also had the same idea and dove for pit road when I didn't expect it, but I had enough time to barely make it onto pit road right behind him. We both slid through our pit stalls like true professionals, and the only adjustment I made was 2% more tape because the engine was running cool. When you're used to a car bouncing like a pogo stick and you suddenly don't have that anymore, it's pretty easy to leave it alone.

[tire data]

Because I was trying to hard to beat Alex in this race, I completely forgot to record tire temperatures. I was, however, recording telemetry so we can take a moment to look at these numbers and see what can be made better on the truck in the future.

Both sets of tire temperatures were taken the laps right before a tire stint ended and in the middle of turns 3 and 4 at the highest point of load. This data looks just about perfect, with camber spreads looking dead-on what I would hope for, and tire temperature averages looking okay. It's important to note that these temperatures are tire surface temperatures (not carcass temperatures), so they won't show heat retention in the tire as well as the garage values. The right-rear tire is cooler again, so it's worth looking into that to see if it could be warmed up with a softer spring rate. The only major issue is a heavy bias towards the front tires, showing a need for less nose weight than what I had in the car (53.1%). Since these camber values look pretty good, we could also experiment with adding more camber to both front tires and see the effects since we know a good set of baseline values.

Moving Forward

Iowa was the last race we'll run on the short tracks with the Truck Project. The final two races at Michigan and Rockingham will use the intermediate chassis, and Michigan's setup will work from the California setup since the tracks are so similar. That said, it's possible that the asymmetry stuff learned at Iowa could find its way into the Michigan car as well if it shows to be a better option. We might be done with the Short Tracks, but these tracks make up the majority of what we know about the Truck so far so it's silly to ignore what we've learned here. In fact, it's very possible that we can take the California setup (which was pretty good to begin with) and tailor it to Michigan. As with previous weeks, we'll race on Thursday night. We all hope to see you there!

Commodore's Garage
Article #35 – Truck Project: Farewell from Michigan

by Matt Holden
January 25th, 2018

For the entire Truck Project series of articles, I've tried to stress the concept of carrying one setup from track to track and making adjustments to suit the track instead of trying to reinvent the wheel each week. While sim-racers have historically spent hours and hours building a setup each week for a new track only to toss it aside for the next week, it doesn't mean that's how things have to be. We split the Truck Project into two mini-branches: Short Track, and Speedway. The Short Track setup dominated this series and rightly so, since it was the most problematic of the two setups from the get-go. Simply adjustments each week took it from a chassis that was largely uncompetitive in week 1 to a chassis that was actually fighting for spots by its fourth time on track. We never built something new and we always pulled from previous weeks to figure out what to change and how to go about correcting the bounce problem that plagued that setup for three weeks. Admittedly, we did stumble across the answer at Milwaukee when Alex made a change on a whim, but that's part of the process! Sometimes you'll find the answer accidentally, sometimes it just takes time.

Michigan Speedway

Michigan was the second time we've pulled the Intermediate setup out this season following the race at California. Coincidentally, California and Michigan are incredibly similar, so instead of updating the setup with wholesale changes I wanted to make minimal changes just to suit the track and see how far off the truck was in the race. In the end, all I changed was the left-front spring, some weight distribution, and sent it on the way.

These two tracks are the same shape, have similar grip levels, but the major difference in the two is the amount of banking at either track. California only has 14° of banking while Michigan has an extra 4° of banking at each end, which adds quite a bit of vertical load to the car at Michigan. At California I ran a lower-rate coil-bind spring on the left-front to keep the splitter down, but I had to run a little bit higher ride heights to compensate for the extra vertical travel. This meant a higher-rate spring was needed to take up the extra travel, and this was basically the only major change that was made. There was also a change in rear spring rate to resist the vertical movement as well, but that was a relatively simple change compared to the front end. By making only simple changes to keep the setup the same, it's a good way to gauge the weaknesses for the setup against competition, which we'll look at later.

Race Day

The race was, by far, the most competitive race we've put a truck in this season. In terms of SOF, it was similar to Iowa but the names in the race were some of the big guns. Where Iowa tested the Short Track setup against some of the best, Michigan would test the Intermediate setup against some of the best. It's probably worth noting ahead of time that my truck was about 7 tenths off of the fastest lap of the race, so there's definitely some short-run speed to be had with it. Long run pace, however, was exceptionally good. I started 20th and finished 11th with only one 3-lap caution in the entire 45-lap race.

These long runs really bring out weaknesses in cars that seem good during practice, and this race produced an interesting behavior with this particular setup. At California on the longer runs we had, the chassis shifted towards understeer as the run progressed. At Michigan, despite having very a very similar setup configuration, produced almost no handling shift as the run progressed. This includes a near-30°F difference in air temperature between the two races. This is significant because it shows the handling issue at California can be mostly attributed to the track temperature itself instead of the chassis producing a handling change. If the chassis was the problem it should have trended towards understeer at both tracks, regardless of the temperature. True, the California setup would be tighter anyway due to the lower track temperature, but the truck showed zero signs of a tight condition at Michigan. If anything, it was trending towards loose at the end of the first run which is typical for a hotter track.

[coleman picture]

The most significant problem with the chassis is a lack of short-run speed. At California I was able to hang with Chad Coleman later in the race, getting in position a few times to set up a pass on him. At Michigan, however, Coleman blew right by without even waving to say hello and I never saw him again. I was in the same car, driving against the same driver, but a dramatically different result at a different track.

In the #1 Truck's Dust

At this point, I'm sure there are some who are thinking, "Well of course you weren't able to race against Chad, you didn't update the setup!" but this actually gives a pretty solid direction on where we would take the setup in the future. We know the handling is good because it didn't change much and we know the long-run performance is quite good because I was making passes through the entire race. Races like this can often leave you complacent with the idea that your setup is good, but details like having Chad pass me without any ability to challenge show that the setup could have been a fluke at California. Chad made most of his ground up in the corners, but not the straights, meaning my car is likely too slippery while lacking cornering speed.

Had Chad never been around me during the race, I likely would look to get some more straight-line speed in the truck while sacrificing a bit of corner speed. After all, I was having no trouble racing the guys I was around for most of the race, but these weren't guys who were challenging for the win. Chad had the #1 truck, marched all the way up to 7th just behind Tyler Hudson, all while leaving me in the dust. I remember Chad from the Pro/WC series, and Tyler's a former DWC Champion, so seeing how Chad's truck behaved around mine is extremely significant. Without seeing that I need to look at corner speed instead of straight-line speed, I would have likely widened the gap to the leaders by adjusting for the wrong thing.

Pay Attention To Your Competitors!

If that pass hadn't happened during the race, I would have been struggling to put together an article because I wouldn't be able to make the point I want to make. Too often, people look at their own car and say "it's not doing this" or "it's not doing that" and adjust for something they're unhappy with and wind up making up zero ground to the front-runners or (in worst case) lose ground to the front-runners. It's extremely important to be able to diagnose your own car and adjust to what you need from the car, but it's equally as important to pay attention to what the front-runners are doing in relation to your car. Do they have higher speed in the center of the corner like Chad did? Do they have a lower corner speed but have a higher straight-line speed? Are they simply backing the corner up while you're sending the car off into the corner like your hair's on fire?

It's impossible to make progress if you don't have a goal. For our Short Track setup, our goal was to get rid of the bounce and load the left-front tire. We did that and picked up a ton of ground on the competition. For the Intermediate setup, the California car performed so well that there was literally nothing that said "This needs to be better". Instead of trying to fix what wasn't broken, I left it alone and ran it again. Because of that, I managed to see what was wrong with a setup that looked quite good at first glance and that produced the goal to work towards. This is one of the situations where the handling is great but the pace is off, which is extremely hard to diagnose and work into the plan of changes in the future without some kind of comparison to a better car.

The End of the Truck Project

In a little over a month you've followed along as I, with the help of my teammates, have gone from nothing, in a car we really knew nothing about, to two setups with few major issues and show a lot of speed and are fairly competitive. Admittedly, we didn't actually start from "nothing" because we have a lot of past knowledge of coil-bind setups with the Xfinity cars (and, technically, the 2016 Cup car as well) but the similarities stop after saying the two cars were on coil-bind setups. Without over-complicating things by rebuilding setups constantly, we simply took issues that showed up in the past and made changes to cure one problem at a time. It wasn't a huge time-sink either, with about 90 minutes each week spent on the setup. Going in with a goal and ideas of how what to do give the practice time a purpose and direction, preventing any waste of time.

I hope that this mini-series of articles was able to convey what I've wanted to show for so long: We're not working with just a bunch of numbers that change a handling factor on the cars, we're working with a simulated system. Our short track setup had a problem that was literally caused by two components fighting each other dynamically, something previously only ever seen on real-world cars.

Unfortunately we only scratched the surface of this fascinating topic, but with the upcoming season in the NASCAR iRacing Series looking to be a bigger challenge than last year, I'm sure we won't be short on things to talk about. I've enjoyed this little off-season project and enjoyed racing against those of you who said "Hello!" during the races. I want to end by saying thank you to every one of you who reads these articles and offered your kind words during this series, you guys really made this enjoyable. For now, we'll roll our trucks back into the garage and trade coil-binding for bump springs. I hope to see you all on track this year, and as always: Best of luck, keep the shiny side up.

[CG Cup Car]