WISCONSIN RACING

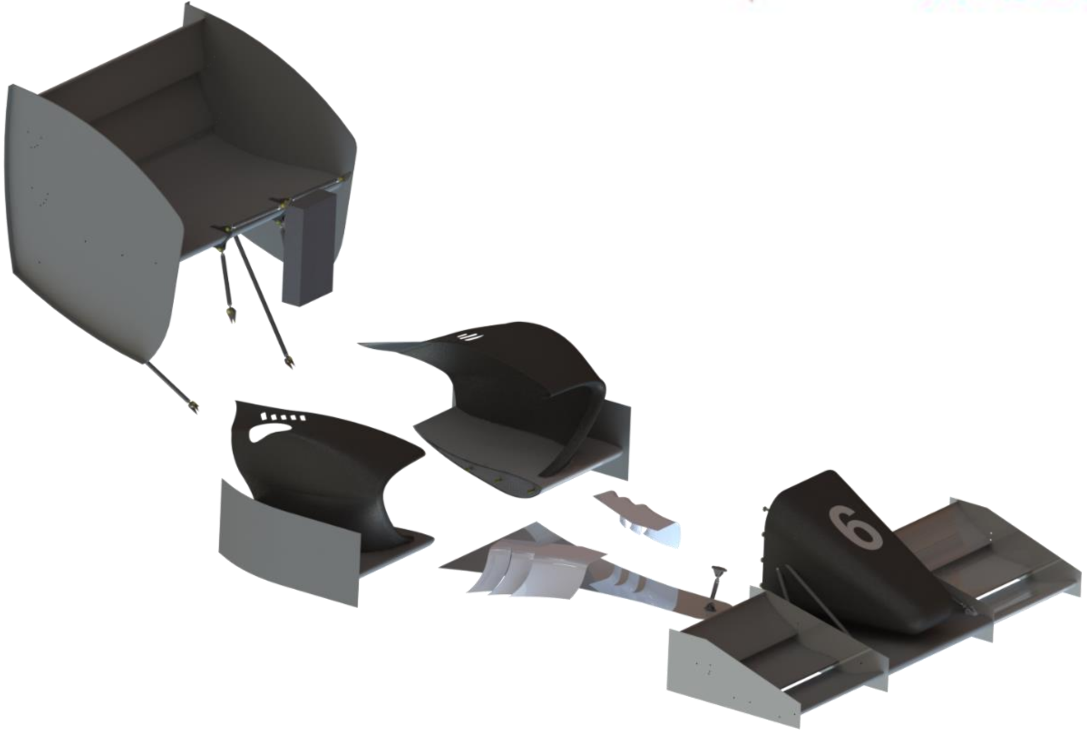

Aerodynamic Package: Design, Manufacturing, and Validation

By: Matt Hanson

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Opening Statements:

The purpose of this informal report is to shed light on how we develop our aerodynamic packages. Those reading this may feel free to use our techniques or reach out to us for further details. With that said, this document will not dive deep into the fundamentals of vehicle dynamics (especially tires), fluid mechanics (such as Bernoulli), nor composites.

Justification for Aero:

We at Wisconsin Racing believe that the fastest cars absolutely require aero. Look at the top cars from competition year in and year out and over 90% have aerodynamic packages. Simply put, the added benefit of downforce outweighs the negative effects of added mass and added size (the front wing likes to eat cones so make sure your drivers are used to the protruding corners). Furthermore, the added benefit of downforce outweighs the increase in drag. For E-cars, this becomes a slightly bigger deal since batteries are not energy dense, but for combustion cars simply add a couple more drops of fuel and you're good to go. We ran a spread of C_L and C_D, Lift and Drag coefficients, through our Lap Sim model (See *Powertrain Architecture, Simulation, and Controls* Document) to compare the effects on lap times. We found that lap times increase nearly linearly with improved L/D ratios. This tells us that it does not matter how much drag you're creating, if you're putting out more downforce, and have an energy source sized out for your competition needs, <u>your lap times will improve</u>.

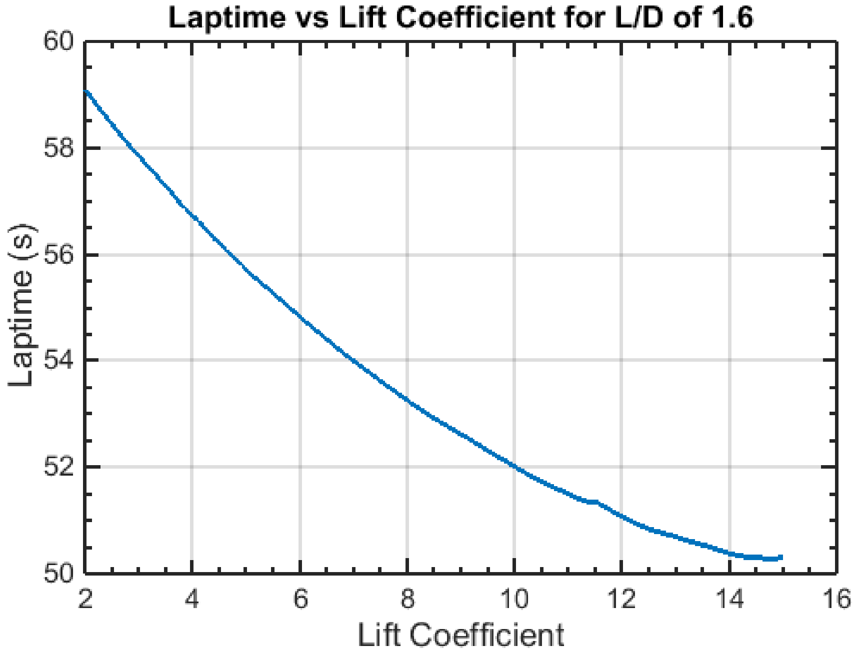

For a constant L/D, we have justification to pump as much downforce into our car as possible.

** Side note: There is also justification to *not* use aero. Aero is an added complexity and works proportional to v^2. If your car does not move, there is no use for aero. Make sure you have a working car and a large enough team dedicated to manufacturing and fixing your package. It is a very labor-intensive subsystem, and if not properly installed will be just one more thing to lose you points during competition.

Design of the System:

Our system is designed to be exactly that: a system. Every aerodynamic device affects the devices downstream of it. Every minor tweak to the front wing changes the airflow to the rear wing. We first design our parts on the individual level, then integrate them into the full car model for further iteration and testing.

We start out by setting overall performance goals for the year. The values are usually determined by % increases from the previous year's Downforce and Drag, and the % is based on how ambitious we think we can be. We take CG estimates from the car's mass model and then set the aerodynamic balance rearward of that. It is necessary for the balance to be behind the CG for sake of stability, otherwise the front tires carry more grip and are more inclined to spin the car out at the slightest movement of the wheel. Next to these, we also look for ways to improve manufacturing to decrease the weight of the package.

Airfoil Selection:

The first step in designing the package is selecting airfoils for the front and rear wings. For sake of ease of manufacturing, we use the same airfoils throughout the car. We started off by looking through airfoil databases for high-lift low-Reynolds Numbers airfoils. We compared the L/D ratios and the overall lift values of many before ultimately settling on the Selig s1223rtl because of its high lift capability.

Multi-element Wings:

Using multiple airfoil's in series allows you to effectively create a high cambered and attached wing. Where one large airfoil would stall past 15°, multiple elements can remain attached to an almost vertical profile. We run three elements in the front and rear, however four and even five element wings are not uncommon. Our reasoning for just three elements and not more is no significant performance gains for the added amount of manufacturing needed.

To determine the size and orientation of each element in the rear wing, we created a parameterized 2D sketch in SolidWorks of three Selig elements in a wing liked shape. We parametrized every dimension (chord length, angle, X-position, and Y-position) of the elements and developed an in-house genetic algorithm to set the final position. The algorithm cycles the sketches through 2D CFD analysis in ANSYS CFX with a bias towards L/D and downforce. It is a lengthy process, but it converged on a set up we have been running for several years now. We have since gone back and swapped out the Selig for numerous airfoils all in the same positions and yielded almost identical results. We just stick with the Selig since we already have the tooling for it and it does not have any difficult features like hooked edges or divets.

The front wing was a little different. We tried a similar method except with the addition of a rotating tire behind the elements. The GA gave us base results, but from there we had to make our own tweaks. It is important to maximize the effectiveness of the front wing, but the strongest front wing set up may be the weakest for the car. If all the air is shooting over the rear wing, then the balance is too far forward and the overall downforce is below what it could be.

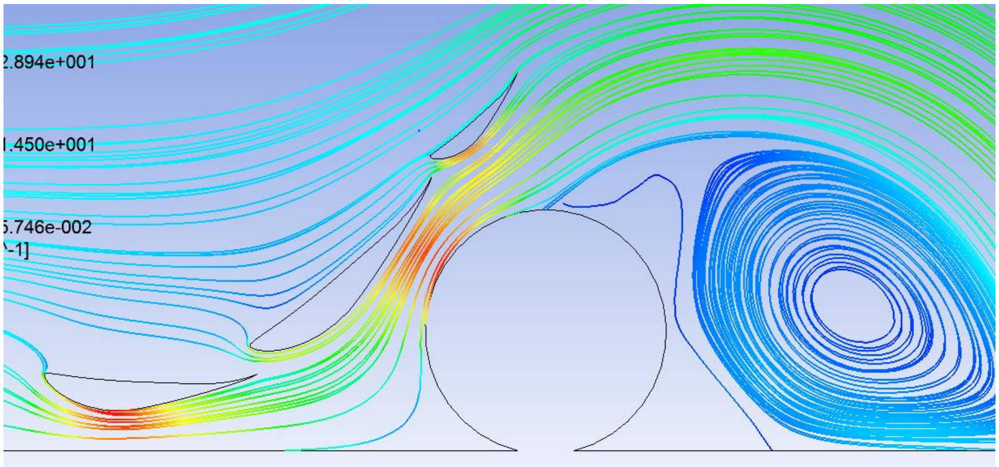

The elements are simulated with the rotating tire to determine their layout. This test shows the limitations of 2D, however, as these elements would be just inside the tire on the car.

CFD Analysis:

We heavily rely on steady-state CFD for the design of the car's aero package. Our preference is ANSYS CFX because of its user-friendly GUI, but other programs are starting to catch up. All our base CFD is run at 35 mph (the measured average speed of our car during autocross) for the purpose of repeatability. CFD will not directly compare to Windtunnel or On-track testing, but if it's always compared to itself it can be considered valid for design. We start off with each component in 2D to rapidly work through parametrization and optimization but understand that 2D has limitations. For starters, vorticity is not captured in 2D, and this is incredibly important for racecar aerodynamics. Once a device has been optimized in 2D, we bump it up to 3D analysis for further iterations. The 3D simulations are run with symmetric boundary conditions and mesh sensitivity analysis to speed up the process. 3D is still much slower so do not anticipate running nearly as many trials of iterations unless you have access to a super computer (Wisconsin Racing does not). For the 3D runs, we look at end effects and vorticity. This of course leads to the creation of fun devices such as gurney flaps, vortex generators, and endplates.

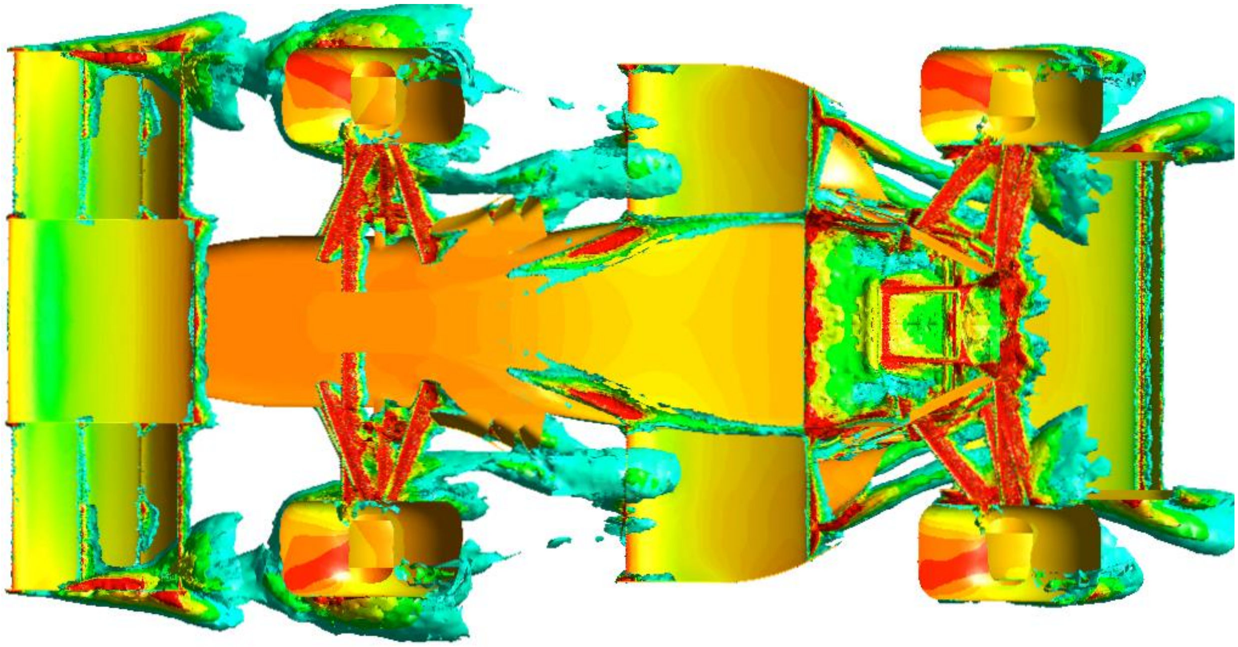

This is the vorticity of the underside of the vehicle. The various endplates and turning vanes work to create vortices to seal off tire wake and hit downstream inlets.

Once an entire device has been designed, say the rear wing, it is put into our full car 3D model and run. The full car is simplified from our actual full car assembly in SolidWorks. Only the necessary surfaces are saved (this dramatically improves simulation time). Again, every simulation is completed with the same boundary conditions and set up to benchmark against itself. The mesh requires nearly 7 million elements, so it is a time costly simulation. Over the past several years we have been squeezing out improvements by constantly adjusting angles and positions of devices. So for early packages… do not get overwhelmed with the constant fidgeting in SolidWorks and CFD. Make your package early, then go back and spend time in CFD finding the perfect orientation of your wings.

**The after-development fine-tuning makes for great new-member projects. It's less design intensive but introduces them to the analysis tools and shows them the sensitivity associated with aerodynamics.

After the base tests are complete, begin testing for sensitivities due to yaw and pitch. The car very rarely drives in straight, level lines. Rather, it is always accelerating and turning. Pitch analysis is important because the front wing will choke flow under the car if it is lowered too much during braking (or from max downforce at high speeds, pushing the car lower to the ground). Similarly, if the car accelerates fast and squats back to far, the front wings could be too high and lose all its ground effect. Begin finding the positioning (height and AoA) for the wing where it is least likely to be drastically affected. Yaw is much more difficulty since you can no longer run with a symmetric boundary. Things to look for in yaw are how effectively you are washing out the dirty air from the tires. If the dirty air is sucked under the car, it can destroy the performance of your diffuser or other ground effect devices. To seal off the underside, employ vortices to channel the air to other parts of the car, or around the car entirely. Additionally, make sure air is being directed cooling inlets. Again, vortices are effective for channeling air to these and keeping dirty air out.

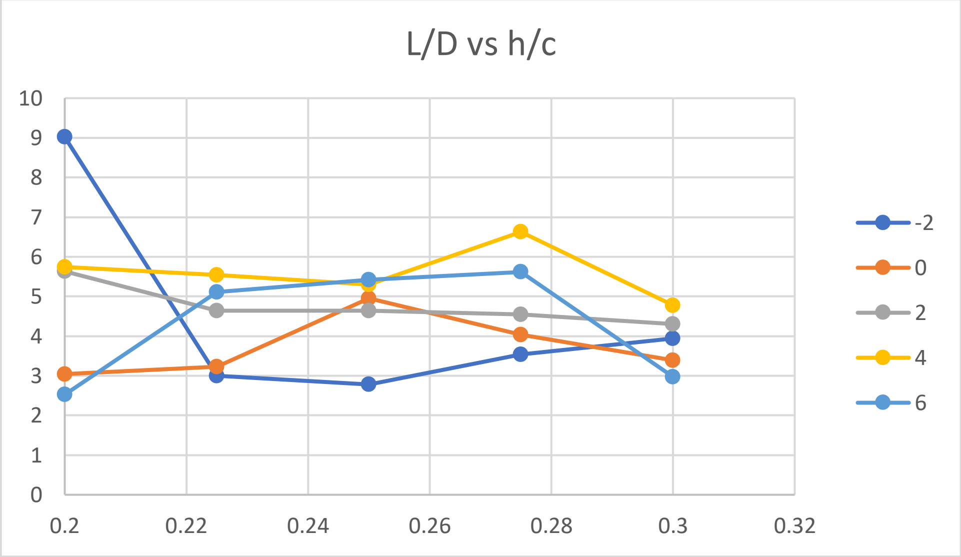

This is an example of yaw-sensitivity for the font wing's main plane. The best L/D is when it's lowest to the ground, but the moment that option raises it loses all its gains.

Manufacturing of the System:

Our aerodynamics package is made from carbon fiber to keep weight down. Do not plan on winning Cost with an aero package. We produce our wing skins, in top and bottom halves, through vacuum infusion and glue them together clamshell style. We are in the works of a new manufacturing method that should help us save cost, time, and weight, but we will not get into that until it is fully implemented. Make no mistake, developing new parts is incredibly time sensitive on the manufacturing side. Once the tooling is in place it is easier to just mass produce wing skins, but for a first-year team understand the challenges you will face to develop everything at once.

<u>Tooling:</u>

The reasoning behind our tooling methods is strongly limited by available materials. We use high density foam tooling board for our wing skins. We get the stock sponsored but must prepare it ourselves. For deep molds we glue together multiple boards before machining. If the part is small enough we can machine the mold in-house, but if not, we ship out to a sponsor with large enough machines. Regardless of it's a plug or a mold, all our tooling board is then sanded to approximately 240-400 grit. We then spray it with a duratec surface primer and continue sanding until it's a polished 2000-4000 grit tool. Your part will *always* reflect the surface of your tool. Scratches will carry over so take care when prepping your tool. At minimum, each tool takes two-days to prepare. But if your labor force does not show up because they do not want to sand, it will take much longer.

From here, we either make carbon or fiberglass molds off the plugs, or we directly layup the final part. We make carbon and fiber glass molds for the curvier parts (sidepods) that get made of TC275 prepreg. The tooling board warps at high temps so it is not reliable for oven cures, and it is too heavy/bulky/easily destroyed to be used in multi-part molds that need to bolt together. Either way, the second molds or direct parts are vacuum infusions.

Wing Skins:

Our wing skins are spread tow dry-carbon. We like the spread tow for aesthetics, but also because there is less crimping from tighter weaves. Our prepregs typically have some form of “pinholing,” but we can get a class A surface from infusions. We start by selecting the spread tow, the underlying twills for support, and the epoxy. Key factors when selecting resins are their gel temperatures (especially important if the part is being used as a mold for oven cures or if its being placed in a hot zone such as by a muffler or motor), their potlife (you don’t want it to cure before you’ve even penetrated the layers), and their viscosity (less viscous resins have difficulty fully penetrating really thick parts such as molds). We use the same wing skins for both front and rear wings which speeds up the manufacturing times. We can layup wide parts and then trim them to multiple parts. Infusions usually take several hours of active labor, and then require up to a day to fully cure so do factor all this time in when trying to make a package last minute.

Wing Elements:

Finally, the wing skins can be glued together to complete the wing. We start by placing the bottom half skin in the mold. Depending on the size or demand of the element, we glue a carbon spar to the ¼ chord. Then we glue aluminum in-house machined endcaps to the very sides of the wings. These are integral for the element to maintain its structure and attach to the final assembly. We’ve used carbon endcaps before in the past, but they’re more difficult to mass manufacture. Next, we glue pieces of foam, machined to the profile of the element, behind the leading and trailing edges. While these help to provide rigid shape, they’re mostly used as surfaces to glue to. Lastly, we glue the top half skin in place and press fit everything until the glue dries. This process is not perfect, and it contains too much glue and internal pieces for our liking, but it does provide a near flawless surface finish and rigid element ready to withstand being beaten by cones. If the leading or trailing edges are not fully aligned, a commonality given our method, then we apply a PCR carbon paste and wait until it hardens. We then sand until smooth and unnoticeable. The PCR paste is essentially bondo but for carbon parts and it’s black so it’s not noticeable.

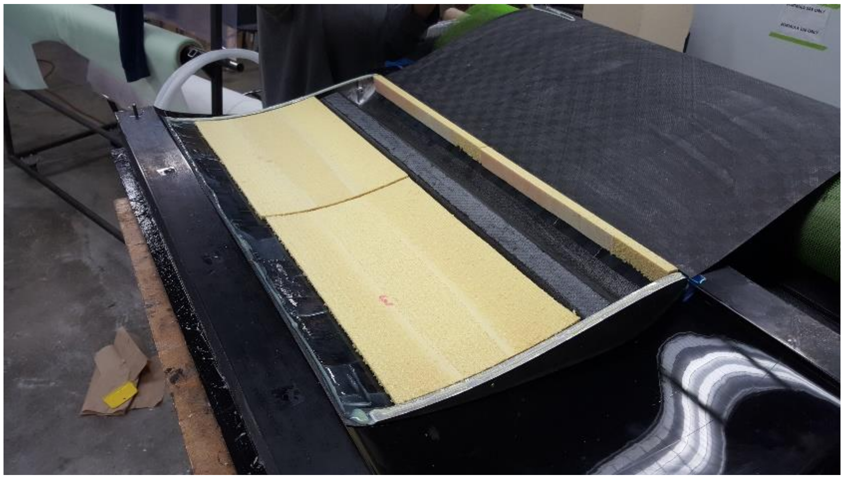

This is the gluing assembly. One the sides are the aluminum endcaps, and interior are the foam pieces and carbon spar. The top skin is about to be folded back onto the element.

<u>Wings:</u>

The wings are then constructed from the wing elements. We create jigs that contain bolt holes for the wing caps to screw into. We align the jigs on the endplates, drill out the holes for every orientation and trim of the elements, and then bolt everything together. It's the same process for both the front and rear wings. Strong word of caution is making sure the endplates themselves are strong. They are pieces that hold everything together and should not be an afterthought. Also, make sure the holes through the endplates will not wear away. Carbon does not like being drilled into so use big washers or inserts. This will keep them from eroding and causing your wings to become shakier than anticipated. We mount everything using tie-rods and rod ends. These allow for easy adjustability, especially during tech-inspection.

This is the endplate assembly jig. The holes are lined up and drilled out. All the different trim settings are on the jig.

Validation of the System:

CFD is a fine tool when compared against itself, but it has no real world meaning. Judges do not care about your CFD. You must show ways that show that your analysis matches up with your car.

Windtunnel:

We are fortunate to have access to a static windtunnel in southeast Wisconsin. To make the results more accurate we place the car on an elevated platform made of plywood and 2x4s. The elevated platform trips a new boundary layer to closer to the car, which better resembles what you'd find on the road. We place scales under each tire and use simple statics to determine the downforce and balance. We also have a diagonal setup attached to the jackbar to measure drag. We make the suspension rigid so that we can measure the exact ride heights for each run and then sweep varying velocities. The windtunnel is useful for quick flow viz analysis, quick ride height sweeps, quick yaw sweeps, and quick swapping of devices and aero set ups. In short, once you get a windtunnel reserved, make sure you check every possible thing on your car because it is much faster to get results than checking each individually in CFD. Leaving the windtunnel, we have complete aeromaps for ride heights and speeds. These can then be cross-correlated in lap sim to determine which heights we will ultimately use.

The racecar is placed onto four scales even with the platform. The goal here is to develop numbers relative to the CFD to prove the simulations are on track.

To validate our CFD program with windtunnel, we recreate the same test in CFD. For this we model the exact dimensions of the windtunnel and platform, and no longer run the CFD with a rolling road. Your numbers will not directly match up which is fine. Just make sure you're able to tell the judges why and where your discrepancies come from.

On-track Testing:

On track testing is the most important form of validation because it’s how the car is actually performing. It is sometimes the hardest to pull off because it requires various sensors and is often noisy to post-process. We use two forms, laser ride-height and coastdown test. The laser ride height measures the change in height of the front and rear axles during a lap. At a constant velocity on a smooth road, the height should not change. But as the car increases speed and downforce, the springs of the suspension get compressed and the car lowers to the ground. If you also know how much your tires compress at different speeds, you can back track the amount of downforce you’re creating. You can also find how the aerodynamic balance changes by comparing the relative delta in front and rear heights. To find drag, we conduct the coast down test. This requires several runs of popping the car in neutral from a certain speed and letting the car coast until it eventually stops. By comparing the velocity over time, and knowing your rolling resistances and transmission resistances, you can back track the drag forces on the car. Both tests are noisy and imperfect so run them several times and then compare them your CFD results. Again, know why they are different and where the discrepancies come from.

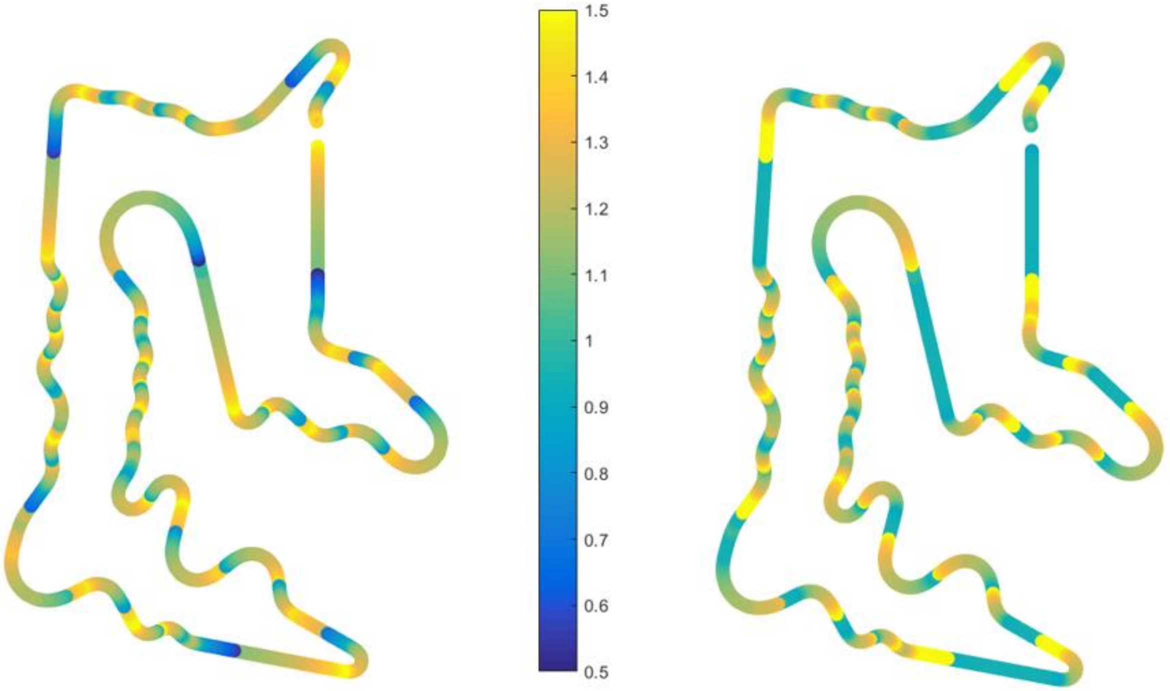

The ride height [inches] for the front (left) and rear (right) axles can be used to validate both CFD downforce results and the lap sim estimations for ride height through an endurance. These results can further be used to determine suspension and wing set ups.

Conclusion:

In conclusion, aerodynamics is what makes slow cars average, average cars good, and good cars great. Making radical changes is not always the best solution, but rather settling on a complete design and then tweaking and optimizing over time. It will make your life a lot easier if you constrain your degrees of freedom rather than trying to redesign everything at once. Make sure you design as a complete system and understand the sensitivities every part has. It all starts with your front wing setting the airflow for everything behind it. Don’t design anything that you won’t be able to manufacture. Manufacturing composites and getting all your surfaces to be perfect demands a lot of time. If you cannot build it, then do not design it. It is more important to get it made and on the car, then forever putting off deadlines and looking for the perfect CFD optimizations. And lastly, it is more valuable to have real world proof that your package performs the way you designed it to, than having one hundred pretty CFD streamlines and plots.