cervélo

AERODYNAMIC DEVELOPMENT AT CERVÉLO

CONTENTS

1. Aero-Centered Design 2

2. Evaluating Drag Performance. 3

2.1. Describing Drag Performance 3

2.2. Drag Versus Yaw 4

3. CFD Simulation: Considerations for Evaluating Bike Aerodynamic Performance. 7

3.1. Mesh Size 8

3.2. Distribution of Mesh Cells 8

3.3 Mesh Partitioning 9

3.4. Simulation of Unsteady Flow 9

3.5. Cost Versus Accuracy for Unsteady Simulation 12

3.6. CFD Validation 12

4. CFD as a Visualization and Design Aid. 12

4.1. Why Do We Use CFD? 12

4.2. CFD in the Design Cycle: A Case Study 13

4.3. CFD Analysis Methods 15

5. Wind Tunnel Testing. 18

5.1. Introduction to Wind Tunnels 18

5.2. Wind Tunnel Validation 20

5.3. Cervélo Tunnel Testing Practices 24

5.4. Wind Tunnel Testing as a Development Tool 34

6. Real-world Validation. 38

7. Conclusions. 41

Contributors. 42

1. AERO-CENTERED DESIGN

By Richard Matthews

Aerodynamics is at the very core of what we do at Cervélo. Phil White and Gerard Vroomen founded the company with the goal of making cyclists faster, and using aerodynamics is the best way to achieve that goal. The success of their approach is obvious when you look back at the 1994 Barrachi, the first Cervélo ever made (Figure 1). It looks like a modern aerodynamic superbike, even more than 20 years later.

Figure 1. 1994 Cervélo Barrachi

At Cervélo, we started making bikes and riders more aerodynamic more than 20 years ago, and we remain leaders in this area. For us, aerodynamics is not just another feature of a bike; it is the heart of our design process. Years before starting the formal design of a new bike, we spend months researching potential aerodynamic ideas. Many of these ideas never make it onto a bicycle, but in this process, we learn. It's not just one person or group in the lead; everyone in our design team is involved. In this way, our entire team learns to design with aerodynamics as a focus, rather than following the lead of a single "expert" who tells everyone else what shapes to use. The merging of a diverse range of experience and views fuels our progression and ensures we continue to grow.

We have been testing in wind tunnels since the 1990s, spending hundreds of thousands of dollars every year to improve the aerodynamic performance of our bicycles. Over time, it became clear that testing a bicycle alone is not enough, because the interactions between bicycle and rider form a large part of the design. Thus, in 2007 we created DZ, our rider mannequin. We started using computational fluid dynamics (CFD) in the early 2000s, performing our own simulations in-house beginning in 2009 and expanding our capabilities every year since then.

All of these factors contribute to our leadership in bicycle aerodynamics—and this leadership has been proven in hard data; it's not just a marketing line. Over the years, Cervélo has been recognized throughout the industry, by our partners and competitors, as both the aerodynamics leader and the clear target for other manufacturers. Independent testing by the likes of *TOUR Magazin* and *VeloNews* has corroborated our designs, as well as our test results.

All of our culture centers on innovation in aerodynamics. This paper describes some of the techniques, skills, and attention to detail behind our success in making the fastest bikes on the planet.

We use three major tools to develop the aerodynamic performance of our bikes: computational fluid dynamics, wind tunnels, and real-world validation. Each has a different role and is used in a different way during a development project. How these tools work, how and where they are used, and their advantages and disadvantages are described in this paper.

In applying these tools, our goal is to minimize aerodynamic drag. But what is drag, and how do we measure it? Both questions can be tricky to answer. Because so much of our work depends on a correct understanding and execution of drag measurements, it is worthwhile to start with some fundamentals.

2. EVALUATING DRAG PERFORMANCE

By Howard Buckley and Richard Matthews

2.1. Describing Drag Performance

Because air is, in physical terms, a fluid, it exerts a force on a moving bike and rider. When we talk of drag on the bike and rider, we are referring to the component of the air's force that is opposing the bike's direction of travel. To keep the bike from slowing down, the rider must produce an equal and opposite force to balance the drag force. To confuse matters, literature in the bike industry reports drag performance using several different measurement units.

Drag performance is most often reported in grams-force or "grams." One gram of force is the force due to gravity acting on a mass of one gram. To get a physical sense for a drag performance gain that is quoted in "grams," imagine the force your hand would feel if you were holding up an object that had that mass in grams. Drag performance improvement can also be described in terms of the rider's power output using units of watts (a measure of power). Power is supplied by the rider's legs to keep the bike moving. When a drag performance gain is quoted in watts, this value is the extra power the rider's legs must supply to move a slower bike at the same speed as a faster bike.

Another, though less common, way of describing drag performance is in terms of drag coefficient (C_D or, more simply, Cd) and frontal area (A). The drag coefficient is more commonly used in academic literature on cycling aerodynamics. It provides a way to compare the drag performance of aerodynamic shapes independent of their size and velocity. To understand the drag coefficient, it is instructive to look at the equation that defines it:

Eq. 1

$$C_D = \frac{\text{Drag}}{\frac{1}{2}\rho V^2 A}$$

Where ρ is the air density and V is the air velocity. It shows that Cd is obtained by normalizing the drag on an object by its frontal area and velocity. We can rearrange Eq. 1 to obtain an expression for drag as follows:

Eq. 2

$$\text{Drag} = \frac{1}{2}\rho V^2 C_D A$$

Equation 2 shows us that to reduce the drag experienced by the bike and rider at a given velocity, we must improve the aerodynamic efficiency of the bike/rider shape (reduce Cd), and/or reduce the area of the bike/rider seen by the oncoming air (reduce A). Because Cd is a dimensionless quantity and area has units of square meters, Cd × A has units of square meters (m^2).

Wind tunnel drag results (or CdA or other values) are typically plotted against yaw angle. This convention reflects the conditions of real-world riding when the wind is not coming straight on to the bike and rider. But what is yaw angle?

2.2. Drag Versus Yaw

Yaw refers to the angle measured between the longitudinal axis of the bike and the effective wind velocity vector. An example will clarify this concept. In perfectly still air, riding a bike at 50 km/h in a straight line results in an effective wind of 50 km/h coming straight towards the bike and rider. If there is a headwind of 20 km/h straight on, the effective wind speed can be determined by adding the bike speed and wind speed (50 + 20 = 70 km/h), as shown in Figure 2. Thus, we see that if the wind direction is straight into or behind the bike's direction of travel, the bike speed and wind speed can simply be added together (or subtracted).

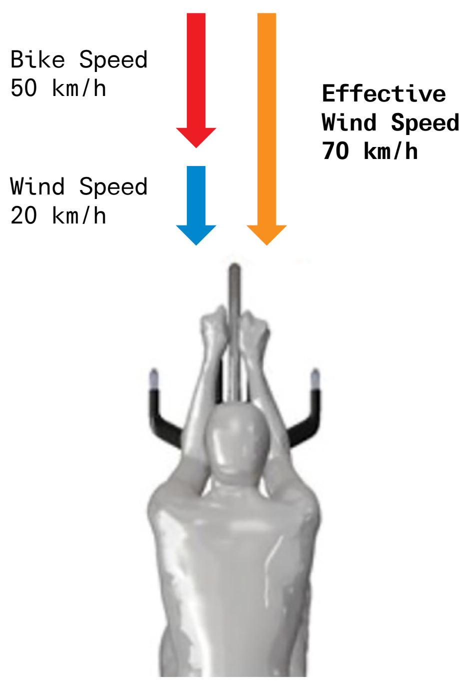

Figure 2. Effective wind speed with a headwind

Now, if the same rider is going 50 km/h but there is a side wind, perpendicular to the direction of travel at 20 km/h, how do we determine the effective wind speed? In this case we have to use vector addition—that is, we add the speed and direction, as shown in Figure 3. If we represent the bike speed and wind speed with separate arrows whose lengths match their relative speeds, by lining them up we can then draw a new arrow that joins the two. This new arrow gives the direction and speed of the effective wind. The same can be done using trigonometry, for those who are mathematically inclined. In the case just mentioned with a perpendicular wind, we can see that the effective wind is at a 21.8-degree angle to the rider and at a speed of 53.8 km/h. This resulting 21.8-degree angle is what we call the yaw angle. It is defined as the angle of the effective wind relative to the bike's direction of travel. A zero-degree yaw means the bike is travelling straight into the wind.

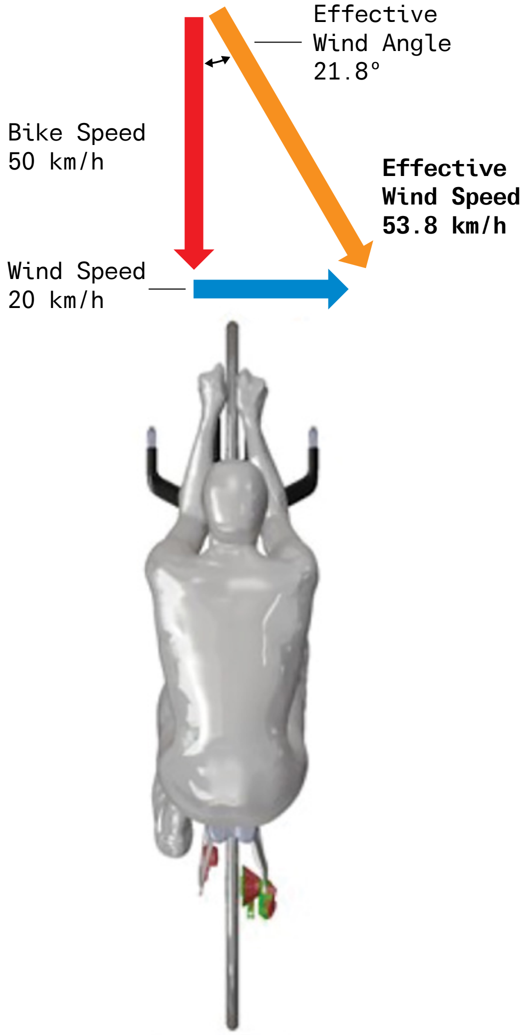

Figure 3. Effective wind speed with a side wind

In the real world, a rider could experience an effective wind at any yaw angle in a 360-degree range on any given ride. In reality, the probability of seeing any given yaw angle varies with the speed and direction of both the bike and wind, including the prevailing wind conditions in different geographic locations, and many other factors, as discussed in Section 5.3.8.

For wind tunnel testing, the wind always comes from the same direction (straight down the tunnel). To simulate a different effective wind angle, we don't move the wind; instead, we move the bike and rider. To do this, we rotate the entire balance assembly (the structure the bike sits on) relative to the wind direction. (The balance assembly, which measures the forces on the bike, is discussed in Section 5.1.) We define the yaw angle as positive when the wind is hitting the drive side of the bike first, and negative on the other side of the bike, as shown in Figure 4.

It is important to note some nuances about how drag and wind speed are measured and reported. The drag force is always measured parallel to the bike travel direction; this is just how the balances work. It makes sense, though, because the drag force reported is then the force that is slowing the bike down in the direction of travel. The wind speed generated by the tunnel is measured in the direction of flow in the tunnel. However, when testing at yaw, the flow in the tunnel becomes the effective wind from the perspective of the bike. Thus, because of the yaw angle of the effective wind, the speed of the head-on air flow seen by the bike and rider is not exactly the same as the tunnel speed. For this reason, the measured drag values are adjusted based on the yaw angle so that the reported test wind speed (50 km/h in most cases) is relative to the bike direction of travel. This correction is standard in the industry.

Figure 4. Sign convention for yaw angle

3. CFD SIMULATION: CONSIDERATIONS FOR EVALUATING BIKE AERODYNAMIC PERFORMANCE

By Howard Buckley

Computational fluid dynamics (CFD) is a method for simulating the flow of fluid around objects, such as bikes, so that aerodynamic performance can be evaluated, analyzed, and visualized. As the name suggests, computers are used to execute the vast number of calculations required to obtain a solution. When considering bicycle aerodynamics, the "fluid" of the name is air.

In CFD, the motion of the fluid is described mathematically by the Navier-Stokes equations, which are often referred to as the governing equations. As you can imagine, a fluid's motion is complicated by the fact that it can deform continuously and that its particles move somewhat independently. (Describing the motion of a solid object is relatively simple in comparison, because all of its particles move together.) In a CFD code, the region being studied is divided into cells, which together are called a mesh. Approximations of the governing equations are evaluated at discrete spatial locations defined by the mesh cells. When simulating an unsteady flow, the governing equations must also be evaluated at discrete instances of time, as the flow evolves. This is known as spatial and temporal discretization. Figure 5 shows an example of a 2D mesh surrounding an airfoil, where the mesh lines define cell boundaries.

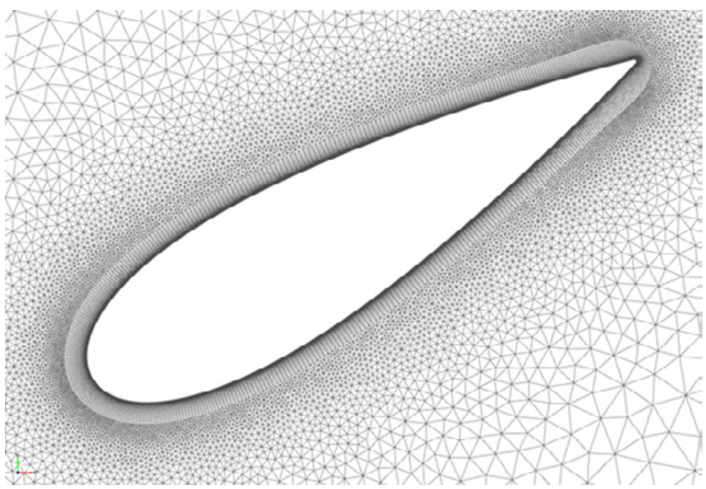

Figure 5. A 2D unstructured mesh surrounding an airfoil

A spatial discretization technique commonly used in modern CFD codes is the finite volume approach. Within each mesh cell, the solution of the governing equations yields the local mass, momentum, and energy of the air—conserved quantities that can be used to calculate aerodynamic forces, such as lift and drag. A solution that satisfies the governing equations at each cell indicates that the net inflow and outflow of these conserved quantities through the surfaces of a cell are equal to their time rate of change within the cell. The set of equations representing all the mesh cells within the fluid domain are solved simultaneously using an iterative approach based on Newton's method until a converged solution is obtained.

In addition to predicting lift and drag forces, CFD can be used to present the solutions visually. Insights about flow behaviour gained from such visualizations can help the designer understand sources of drag and may inspire design changes to improve performance. Figure 6 shows a source of drag associated with a large wake (shown in blue) behind the rider's leg.

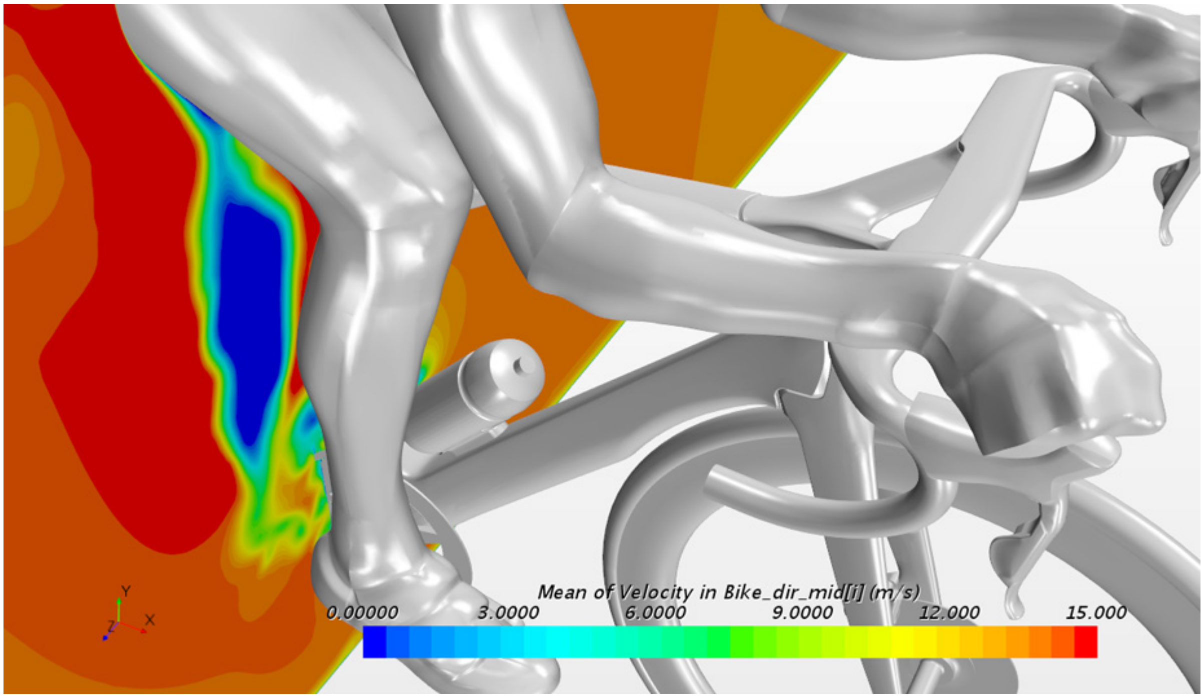

Figure 6. Contours of velocity showing separated flow in wake region (dark blue area) behind rider's shin

The validity and usefulness of a CFD simulation is highly dependent on several key factors affecting the accuracy of the solution and the computational expense required to obtain it, as discussed in the following sections.

3.1. Mesh Size

The solution error associated with the spatial discretization of a CFD code is directly related to the spacing between mesh nodes and size of mesh cells. A more accurate solution can be obtained by reducing mesh spacing, but with finer spacing, more mesh cells are required to fill the space. The increased accuracy comes at a cost, because for each additional mesh cell, additional equations must be solved, requiring additional computational resources. For this reason, Cervélo has conducted mesh convergence studies to determine the mesh spacing that minimizes computational cost while producing solutions of sufficient accuracy for design purposes. Mesh sizes are normally between 20 and 30 million cells for a typical bike-rider simulation.

3.2. Distribution of Mesh Cells

While it is true that a more accurate solution can be achieved by increasing the number of mesh cells within the fluid domain, solution accuracy can be optimized for a fixed number of cells by employing an intelligent cell distribution strategy. Use of such a strategy requires an understanding of where fine mesh density is required to resolve small-scale flow features.

3.3 Mesh Partitioning

The total time required to obtain a solution on a large mesh can be significantly reduced by partitioning the mesh into groups of cells that are assigned to designated compute cores within a high-performance computing cluster. The cell calculations within each mesh partition are performed in parallel, then brought together to form a complete solution. The degree to which using cores reduces the computing time depends on the scaling efficiency of the code. A CFD code that scales perfectly (100% scaling efficiency) reduces solution time by a factor equal to the number of cores. That is, using 10 cores decreases the solution time by a factor of 10 compared to using just one core.

In reality, however, efficiency decreases as the number of cores increases (Figure 7). When using relatively few compute cores (less than 100), it can be seen that the scaling factor is almost 1.0, indicating that the solution time is approximately N times faster when executed on N cores. When using more than 100 compute cores, the scaling factor becomes progressively smaller. At 280 cores, the scaling factor is slightly reduced, but remains high at 0.89, indicating that the solution time is 0.89N times faster. This CFD code is still using the computational resources efficiently. Well-written, modern CFD codes can be expected to achieve scaling efficiencies of 70–80% when using hundreds of cores.

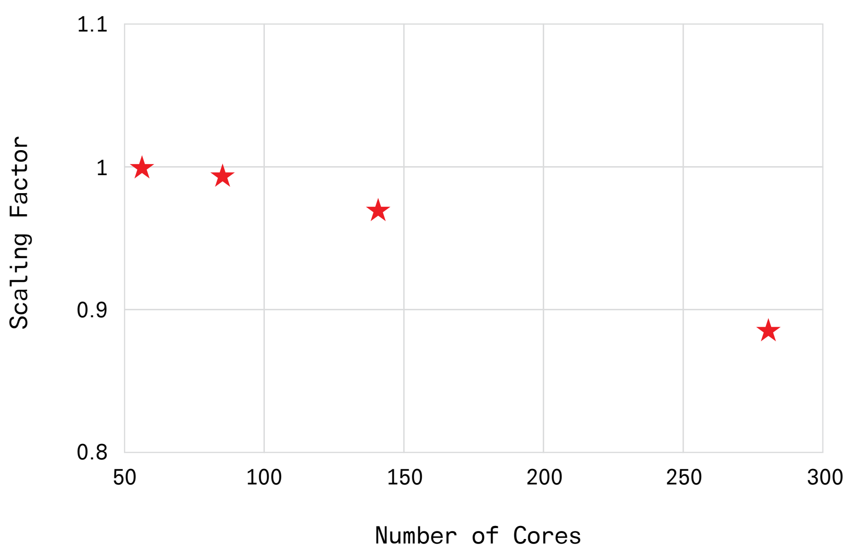

Figure 7. Parallel scaling efficiency of a flow solver for CFD simulations run on a high-performance computing cluster (simulations by Cervélo used a mesh size of 24 million cells with Star CCM+ v11 on Intel E5 2690 v4 processors)

3.4. Simulation of Unsteady Flow

When evaluating the aerodynamic performance of a bike with rider, it is important to understand the time dependence of the flow being simulated in order to determine what type of CFD simulation is required. Steady flow conditions do not change with time, whereas unsteady conditions are continuously changing. As an example of an unsteady flow, imagine a flag blowing in the wind. If you were to take two pictures of the flag, each a second apart, the flag would have a different shape in each picture. The deformation of the flag fabric (Figure 8) provides a visual indication that the complex interaction between the flag surface and the wind (speed and direction) is continuously changing with time.

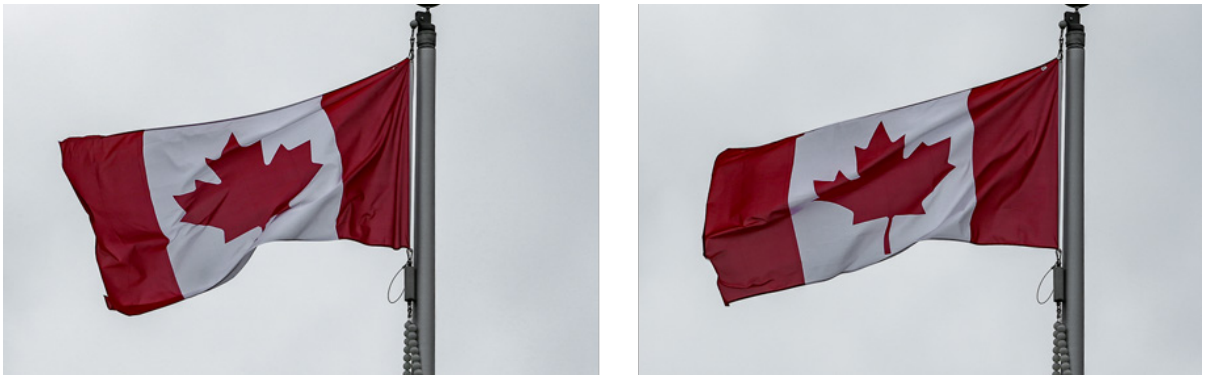

Figure 8. A flag blowing in the wind as a visual illustration of unsteady flow conditions

With respect to cycling aerodynamics, air moving around a cylinder is a more relevant example of an unsteady flow: The swirling eddies in the wake of a cylinder are continuously changing shape, size, and direction. In contrast to the flag example, the unsteadiness of the flow is not visually apparent to the casual observer. However, there are other ways to visualize and quantify unsteady flow. Figure 9 illustrates fluid flow around a cylinder. The top half of the figure shows “laminar” flow around the cylinder, characterized by flow pathways (streamlines) that deform in a uniform way as they pass over the cylinder surface and cleanly reattach into a smooth flow behind the cylinder. The laminar flow pattern around the cylinder doesn’t change from one instant to the next; therefore, it is considered to be a steady flow condition. The behaviour of a flow is related to a parameter called the Reynolds number, which is determined by several physical characteristics of the flow. The laminar flow regime exists up to Reynolds numbers of around 10,000. Beyond Reynolds numbers of 10,000, the flow transitions to “turbulent” flow, as shown in the bottom half of Figure 9, where swirling flow features called vortices or eddies produce a wake that trails behind the cylinder. The shape and size of the wake region, including the flow features within it, are continuously changing, indicating an unsteady flow condition. During typical bike riding conditions (25–35 mph), where Reynolds numbers range from 100,000 to 300,000, a bike and rider will experience unsteady, turbulent flow conditions qualitatively similar those in the bottom of Figure 9.

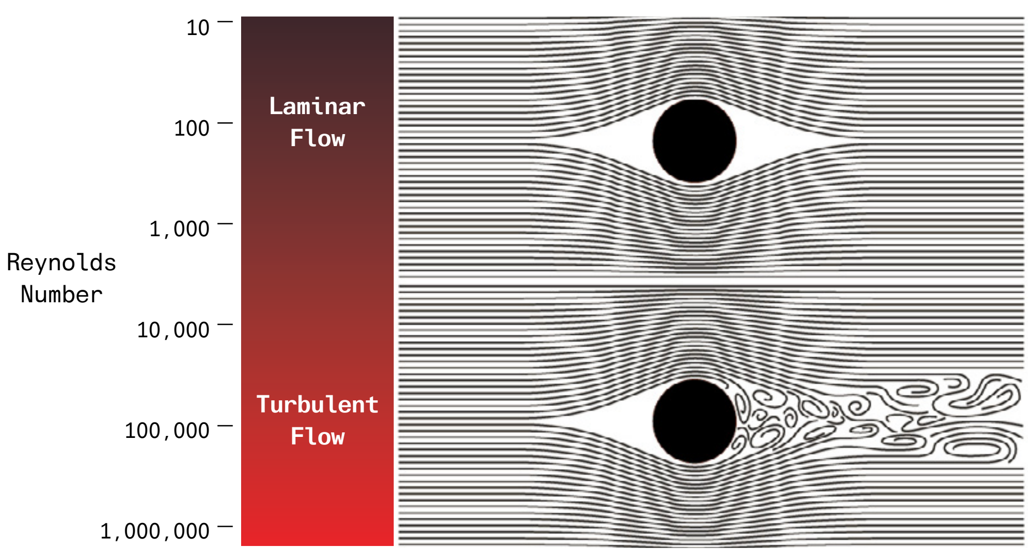

Figure 9. A comparison of laminar and turbulent flow structures around a cylinder

Most fluid flows of practical engineering interest have some degree of unsteadiness. In certain cases, an approximation of steady flow is reasonable if the fluctuation of fluid quantities and/or aerodynamic forces on an object is small compared with their mean values. Figure 10 shows drag values changing with time for a bike-and-rider simulation at a wind speed of 30 mph. The instantaneous drag on the bike and rider fluctuates significantly by ±6% about the average drag, indicating a substantial degree of unsteadiness. The unsteady nature of the flow around a bike and rider is not surprising, considering that the rider geometry consists mainly of large cylindrical and spherical objects—shins, thighs, torso, arms, head—all of which produce large unsteady wakes.

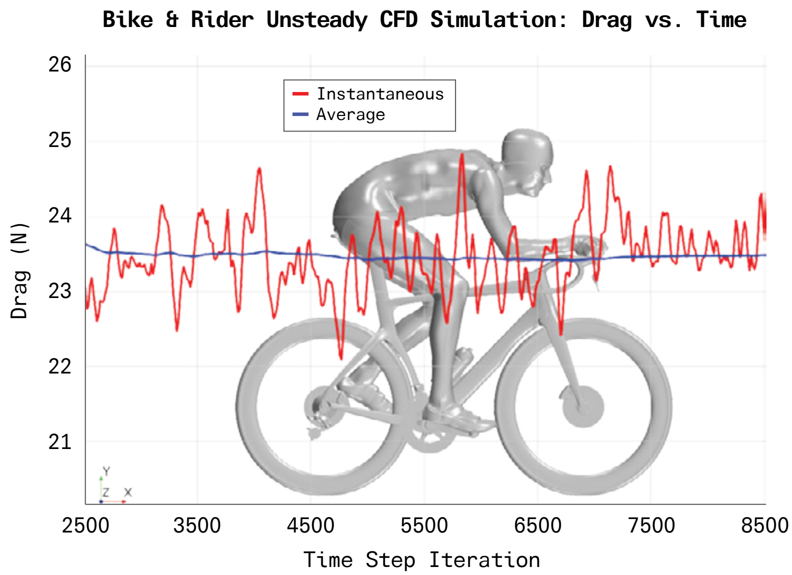

Figure 10. Fluctuations of instantaneous drag values over a period of 1.2 seconds

An important consideration for evaluating aerodynamic performance with CFD is whether to simulate steady or unsteady flow conditions. For attached flows where an approximation of steady conditions is reasonable outside of the turbulent boundary layer (the air layer very near to the surface of an object), it is appropriate to use a steady-state CFD simulation based on the Reynolds-averaged Navier-Stokes (RANS) equations with turbulence modelling. Although this approach is intended for steady-state conditions, it is possible to modify the parameters to obtain solutions for cases that have small regions of unsteady, separated flow. This approach requires less computation, but it distorts the effect of the turbulent regions on the overall flow properties and adds error to the solution. The air flow experienced by a bike and rider can be expected to have moderate- to large-sized regions of separated flow, however. Attempting to simulate such a flow with a RANS-based approach would produce solutions with excessive error. Therefore, when studying cycling aerodynamics with CFD, simulations must address unsteady conditions.

To balance computational efficiency and accuracy, Cervélo uses a hybrid approach, called detached eddy simulation (DES), in which the simulation strategy is adapted according to the scale of the flow features. For regions of unsteady turbulence with small-scale features, such as the boundary layer, a computationally less expensive RANS method is used. A more intensive approach, known as large eddy simulation (LES), is used only in regions with large-scale flow features, such as those typically found in the wake regions of a bike and rider.

3.5. Cost Versus Accuracy for Unsteady Simulation

For an unsteady simulation, both the accuracy of the solution and the cost of the simulation (in terms of the computational resources required) depend on two factors: the size of the time step and desired level of convergence of the solution. Cervélo has performed studies to find appropriate settings for these parameters that will produce results with sufficient accuracy while minimizing computational cost.

3.6. CFD Validation

For CFD to be a useful design tool, the user must be confident that the results from a CFD simulation accurately represent the aerodynamic performance that would be experienced by a cyclist in real-world riding conditions. A simulation method is validated by running it on published, standard test cases. Validation methods differ depending on how a simulation is being used, whether (1) to predict absolute values of aerodynamic forces and moments, that is, how close the simulation comes to predicting performance as experienced by the rider during on-the-road riding conditions or (2) to identify relative differences in aero performance.

During the aerodynamic analysis phase of a new bike development project, CFD is most often used for evaluating relative differences in aero performance between competing component geometries and bike configurations. Validation in this regard provides assurance that CFD will not only identify the faster bike configuration but also quantify the advantage. (In this case, the "true" aerodynamic performance values are of secondary importance.) To support such comparisons, Cervélo has performed extensive validation of DES simulation results against relevant benchmark aerodynamic test cases. These test cases range in complexity from simple cases that simulate quasi-2D separated flow conditions, to cases with 3D cylinder flow conditions, to a comparison between a 3D simulation of fully detailed bike and rider with wind tunnel results of the same configuration.

4. CFD AS A VISUALIZATION AND DESIGN AID

By David Killing and Howard Buckley

4.1. Why Do We Use CFD?

Because Cervélo's reputation and brand are closely tied to the performance metrics of our products, it is important that we have the right tools to assess this performance throughout the design process. The earlier we can get feedback, the earlier we can avoid going down the wrong path.

CFD is an essential tool for getting this feedback early in product development. It has several key advantages when used in conjunction with wind tunnel and real-world testing.

- It provides immediate results (compared to setting up a wind tunnel test).
- It can cost less than building a physical prototype and testing that in the wind tunnel.
- It allows us to visualize—and thus learn about—air flow in ways not possible in real life.

- It is effective for testing very subtle changes in design.
- It allows for use of optimization code to automatically refine a design.

4.2. CFD in the Design Cycle: A Case Study

The value and importance of CFD is best illustrated by walking through our design cycle.

4.2.1. Iterative Concept Design

At the beginning of a project we define the product goals and the scope of the project. For example, designing a custom stem and bar may be within the scope if we can get a significant aero advantage out of them.

From there we brainstorm around a wide variety of ideas, usually stemming from intuition, that are contributed by all our designers. Usually these are vastly different design concepts, which are initially built as simple computer-aided design (CAD) models. We initially apply CFD to these simple models, which helps us iterate quickly but also prevents us from getting inadvertently hung up on measuring the effect of small fillet details or surface gaps. At this point we just want to understand which paths show the most potential and the order of magnitude of the predicted improvements. We don't have time to fully develop and detail hundreds of design ideas, so narrowing the search based on quick results for simplified forms helps us get going in the right directions.

For example, early in the process of developing the 2019 S5, we were trying very large shape changes to learn about the potential of different concepts (Figure 11). We tried a range of ideas, from split seat tubes, to large gussets, to unique bar configurations. It was in this early phase of development that we first discovered the potential of the V-stem configuration, which led to the stem seen in the production S5.

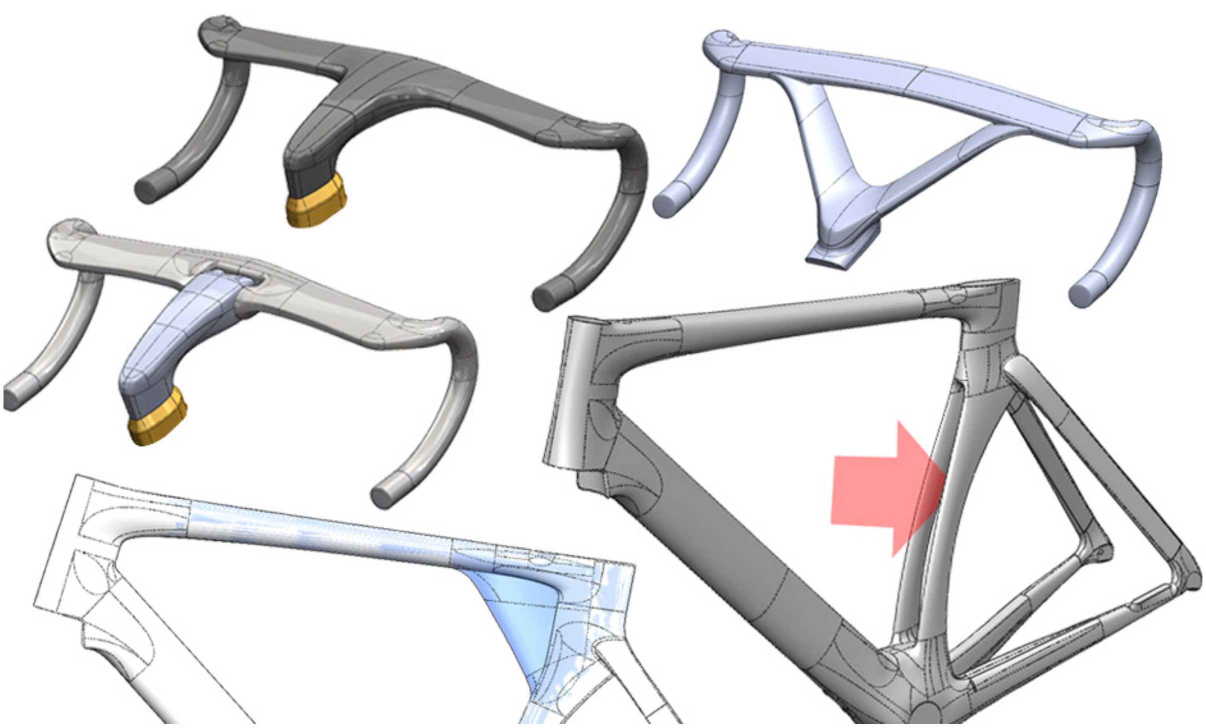

Figure 11. High-level design concepts early in the S5 development process

After running each concept in CFD, we may choose to look further into a result and try to understand why it did or did not perform well. Images of wake areas, flow direction, and pressure contours, examples of which are shown in Section 4.3, help us understand why a design is performing a certain way and allow us to create educated guesses about what to try next. We continue in this cycle until we have a design (or set of designs) with enough potential to move into detailing.

We continue to retest ideas in CFD as we progress or change the design, but the frequency of the tests will diminish as we move farther into the project timeline.

4.2.2. Optimization as a Refinement Tool

When we build a CAD model for a design, it is created parametrically. That is, we can go in after the model is complete and revise a dimension, and the design will update automatically. This approach permits the use of optimization software. In this process, the software will automatically vary a set of dimensions in the CAD model, run the resulting shape in our CFD software, record the aero drag value, and then intelligently adjust the dimensions to try to improve the aero drag. It repeats this process over and over, finding the lowest drag value it can.

For example, on the drop bar for the 2019 S5, we wanted to determine the best angle of attack for the cross section of the bar. We had seen in early CFD images that the air flow over the drop bar was not perfectly horizontal, so there was the potential that the airfoil pitch angle could be tuned. One approach would be to manually save off the CAD file with a few different pitch angles, run them in the CFD simulation, and then choose the best angle. However, this testing doesn't necessarily lead us to the absolute minimum possible drag. Instead, for the handlebar, we set up an optimization in which the angle of attack of the blade section was a parametric dimension defined separately at the root and the tip (which allows twist along the blade). The optimizer iterated through combinations of these two variables until it converged on a minimum drag condition. Figure 12 shows pressure distributions over two sections of the optimized drop bar where the lift is minimized, contributing to drag reduction.

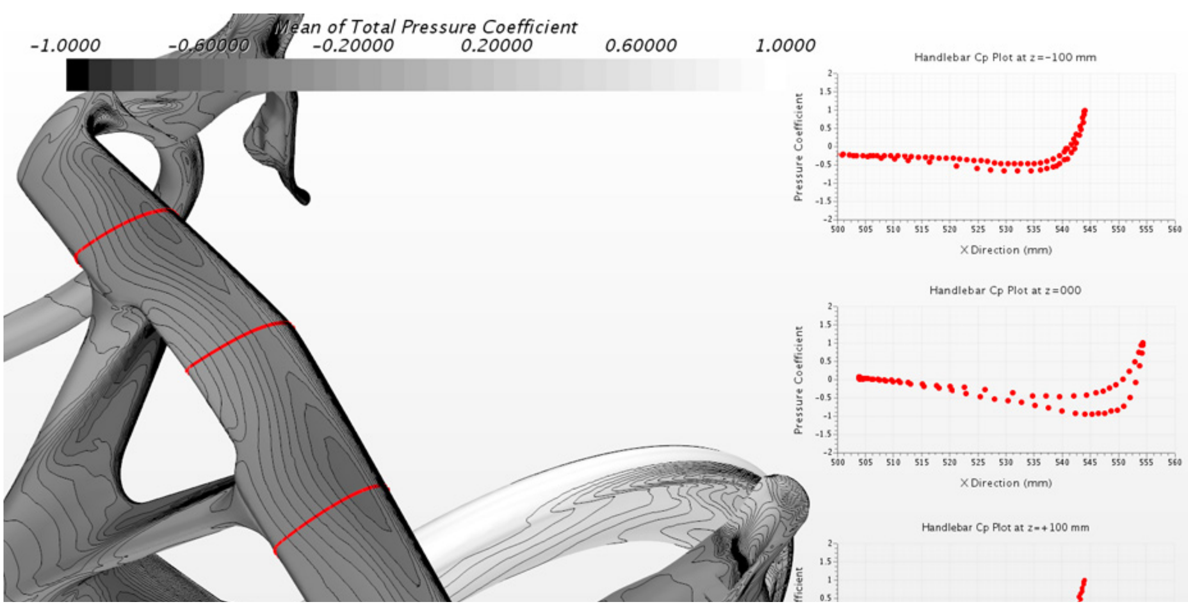

Figure 12. Pressure distributions at two sections of the optimized S5 drop bar

The process didn't end there. We took the design output and validated its potential by testing a 3D-printed part in the wind tunnel. However, it was the CFD work that led to the design as we see it today (Figure 13). When you look along the length of the top section of the AB08 bar, you can see how the angle of attack changes from the root to the tip, optimized to reduce drag when the hands are in place on the hoods.

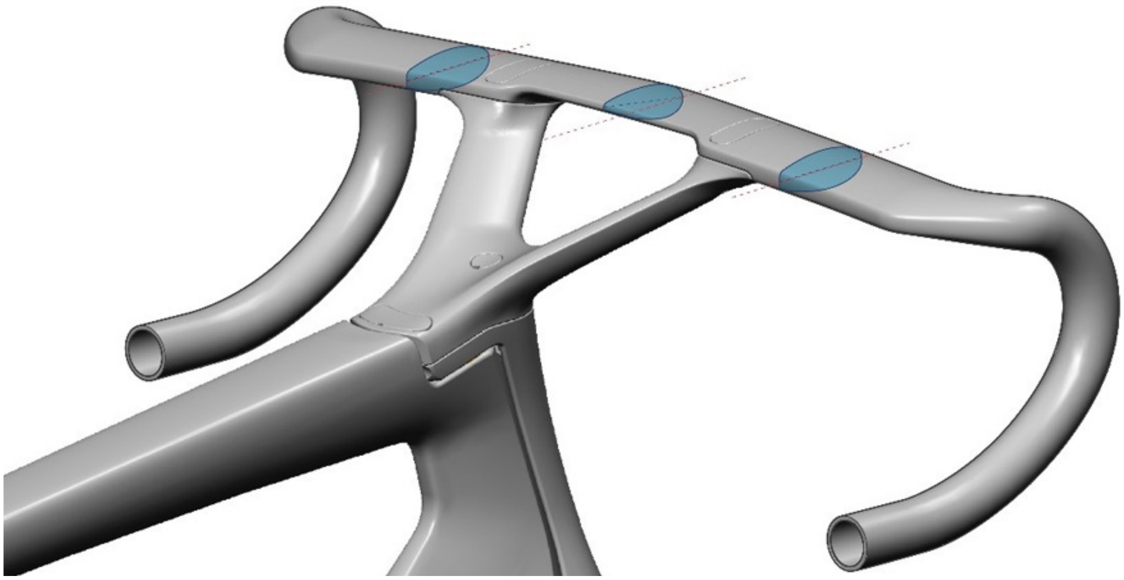

Figure 13. Aerodynamically favorable twist in new handlebar for 2019 S5 (AB08)

4.3. CFD Analysis Methods

The output from a CFD simulation is a large data set that describes the fluid properties of the air surrounding the bike and rider. The CFD data can be used in many ways to visualize the flow and analyze quantities related to the aero performance of the bike and rider. The breadth and versatility of CFD data presents a distinct advantage over wind tunnel testing, where analysis is limited to quantities that can be measured directly, such as forces and pressures, and physical visualization techniques, such as smoke, tufts, laser sheets, surface oil flow, and Schlieren photography. The following are some of the powerful analysis techniques unique to CFD that are commonly used by Cervélo engineers and designers.

4.3.1. Total Drag

A typical CFD analysis of two bike configurations begins with an evaluation of the total drag of the bike/rider system over a range of yaw angles, as shown in Figure 14. Here we see that bike A has lower drag than bike B at all yaw angles from –10 through +10 degrees. At yaw angle +10 degrees, bike A has significantly less drag than bike B. While it may be interesting that bike A achieves a significant drag reduction at yaw +10 degrees, the total drag values do not give any insight to explain why drag has been reduced.

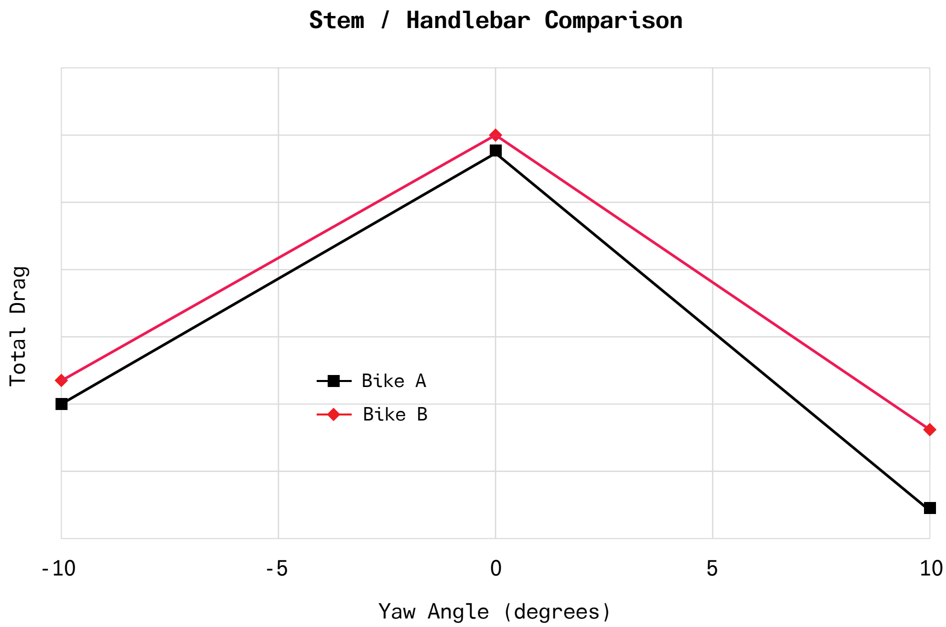

Figure 14. Typical CFD simulation of drag vs yaw performance for two bike configurations

4.3.2. Drag Breakdown by Component

A first step to understanding why two bike configurations produce different total drag values is to compare the drag produced by the individual components of each bike configuration, including the parts of the rider's body (head, arms, legs, etc.) and the mechanical parts of the bike (frame, fork, wheels, handlebar, etc.). Figure 15 shows a breakdown of total drag at yaw +10 degrees comparing the components of bike A and B. This breakdown shows that the majority of the drag difference, 48 grams out of 60, is due to the thighs and torso.

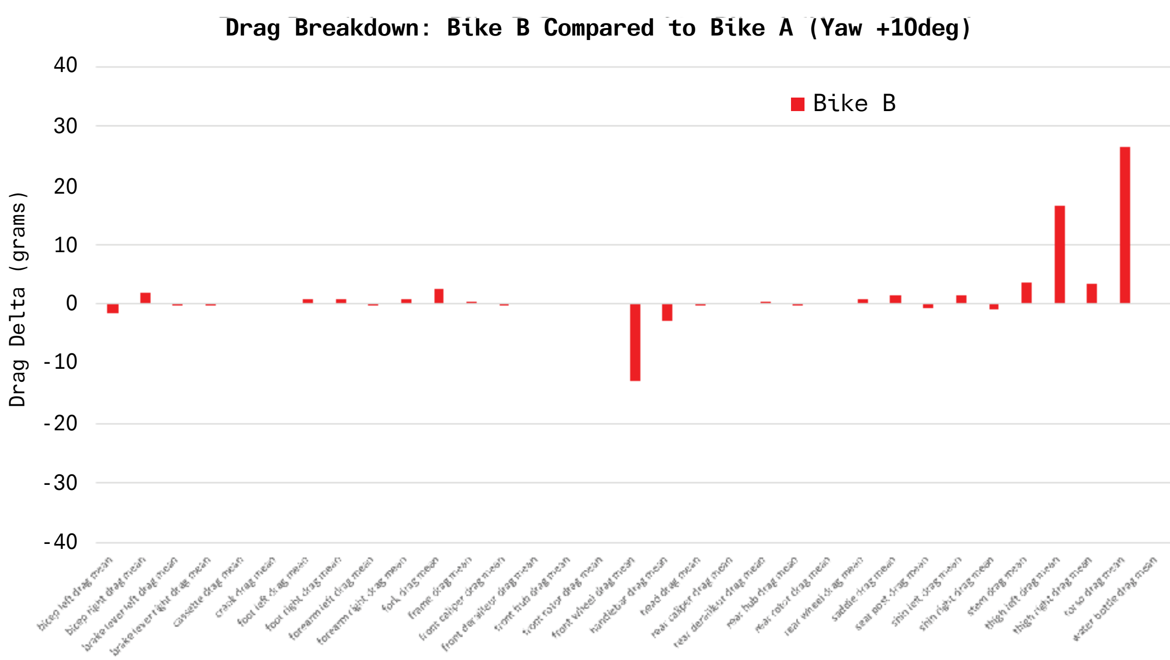

Figure 15. Breakdown by component of the change in total drag of Bike B relative to Bike A

4.3.3. Velocity and Pressure Contours

Velocity and pressure contour plots are examples of flow visualization in which colour-coded maps of the flow field are used to indicate changes in these quantities that are directly related to drag. A drag breakdown comparison such as that in Figure 15 may suggest areas of the flow field to focus on when analyzing visualization data. The contours of velocity shown in Figure 16 illustrate how the size of the wake region (dark blue area) behind the rider's thighs differs between Bike A and Bike B. The smaller wake region for Bike A supports the conclusion from the drag breakdown in Figure 15 that the majority of the drag difference between the two bikes is attributed to the thighs and torso.

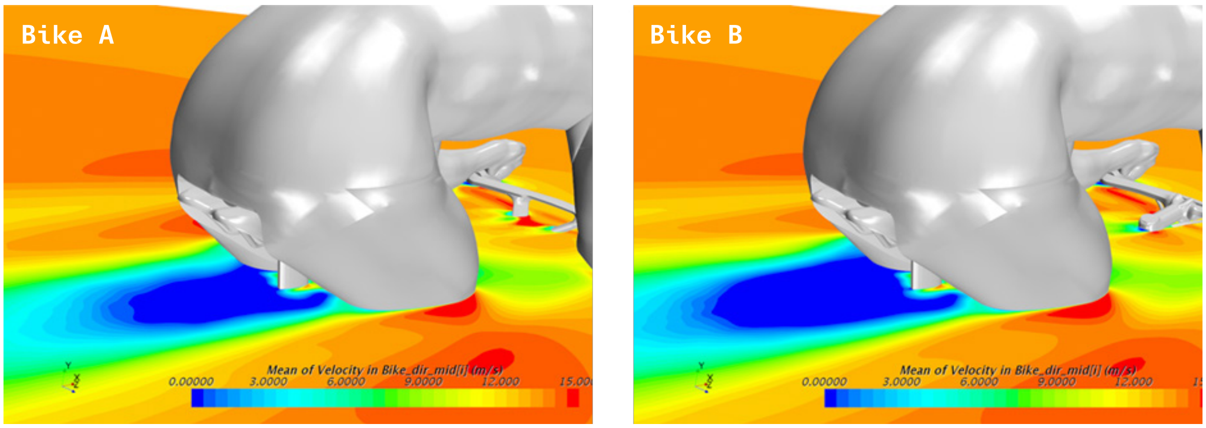

Figure 16. A comparison of velocity contours showing wake regions behind the rider's thighs

4.3.4. Streamlines

A streamline is the path traced by a particle as it moves with the flow around an object. CFD allows us to compute streamlines around the bike/rider geometry, which is another useful technique for understanding how changes in bike geometry change the drag performance of the bike/rider system. A comparison of streamlines around two different stem/handlebar configurations is shown in Figure 17. The streamlines illustrate the flow patterns originating upstream at the stem/handlebar that affect the size of the downstream wake regions behind the thighs, which results in the drag difference observed between Bike A and Bike B.

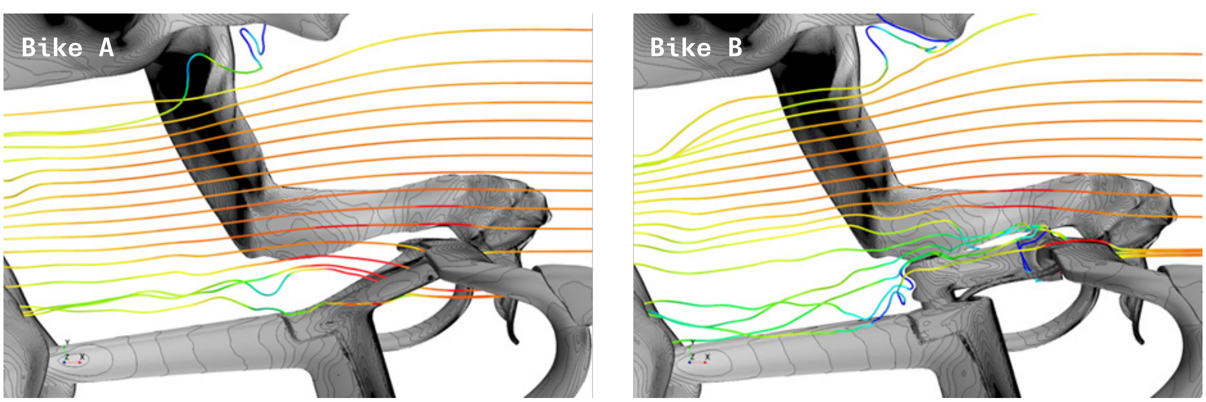

Figure 17. A comparison of streamlines around two different stem/handlebar configurations

5. WIND TUNNEL TESTING

5.1. Introduction to Wind Tunnels

By Richard Matthews

We have discussed CFD and its use as a tool to help us design faster bikes. However, it is only one tool in the development process. To validate findings and solutions discovered through CFD analysis, Cervélo relies on a second major tool: wind tunnels. A wind tunnel is a very specific test facility designed for performing aerodynamic tests. Developed for the aerospace industry to test aircraft (or more often models), wind tunnels come in many types, sizes, and shapes, but they all have a few key components in common. These are the fan, test section, and force balance.

The fan in a wind tunnel is exactly what it sounds like: a device for pushing (or pulling) air into the tunnel (Figure 18). The fan must be adjustable, so that different air speeds can be created in the test section.

Figure 18. Wind tunnel fan at RWDI, Guelph, Ontario, Canada

This test section is where the object being tested is placed (Figure 19). Typically, the test section is in an area of the wind tunnel called a contraction. In the contraction, the cross section of the tunnel shrinks, helping to speed the air up to the correct test speed while also making the air flow more uniform. When speed is the project goal, as it is for us, the most critical job of a wind tunnel design is to provide this uniform air flow. If the flow is rough (i.e. turbulent) tests will not give consistent results. Inside the tunnel itself, there are usually many different features which help to make the flow uniform, including diffusers, turning vanes, and others. All of this technology works together to provide a consistent, controllable air flow at the test section. Within the test section is the setup and instrumentation used to measure data on the test object (in our case, a bike and rider) during testing. The speed and uniformity of air flow are measured by using pressure probes (pitot tubes) or anemometers. The different forces applied to the test object are measured by using a force balance.

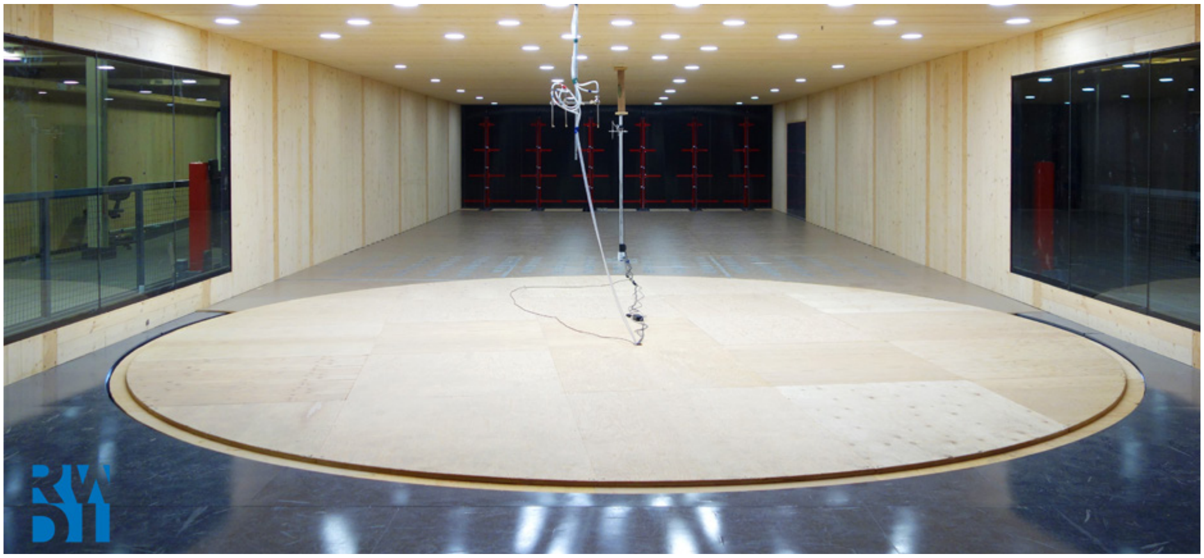

Figure 19. Test Section of Wind Tunnel at RWDI.

The force balance (Figure 20) is an extremely complicated mechanical device which is able to both hold the test object and also measure forces on it in up to six different directions (three directions plus three rotations). A force balance for a bike/rider setup is very specific and designed only for this purpose, so only a few wind tunnels have them. Such an instrument can cost anywhere from a few hundred thousand dollars up to a million dollars—that's a lot of money to test only bikes! Most wind tunnels will have different balances which can be installed for testing different things (such as bikes, cars,

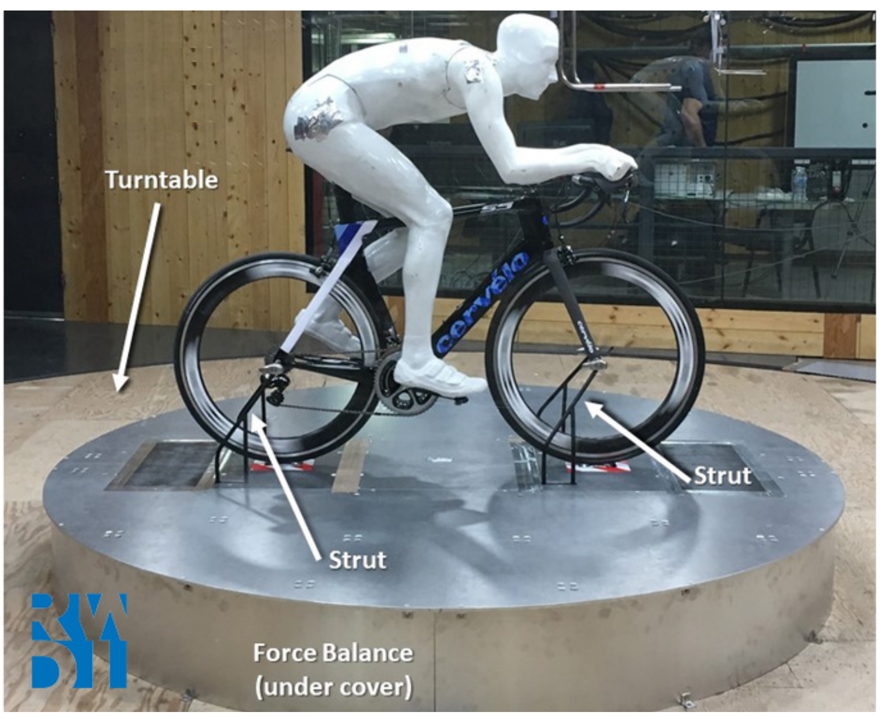

Figure 20. Force balance at RWDI

or planes). An important part of the force balance is the struts. These are metal supports which attach to the front and rear wheels of the bike and connect to the balance. Struts serve two purposes: (1) they support the bike/rider test object in a fixed position during the wind tunnel experiment, and (2) they transfer aerodynamic forces exerted on the bike/rider to the force measurement device within the force balance system. Below the struts are rotating wheel drums (rollers) on which the tires sit. These drums are powered, and they rotate to spin the wheels at the same speed as the air flow in the tunnel, simulating real-world behaviour. The entire force balance system is mounted on a turntable which rotates relative to the air flow direction to simulate different yaw angles. Often this whole assembly—turntable, force balance, and struts—is called the test rig.

The remainder of this section will go into detail on how and why we test at the wind tunnel and how this tool works in conjunction with our other tools to help us design faster bikes. First, we discuss the steps we take to ensure the results are reliable and repeatable:

- Strut tares, to quantify the drag contribution of the struts that hold the bike
- Validation tests, to confirm that the operation and conditions of the wind tunnel are consistent across data runs
- Test protocols, to ensure that results reflect only the differences among test objects, not variation in other factors.

After this review of our test procedures, we discuss how we use these quality-controlled results in our development process.

5.2. Wind Tunnel Validation

Wind tunnel validation (WTV) is the process of ensuring that the wind tunnel measurements (1) match the expected real-world results and (2) are consistent and repeatable. Validation is essential for two reasons. First, wind tunnel testing is expensive, and it is essential to know that the company's money is buying valid results. Second, only by having properly validated results can tests be compared against each other within a session and over time. The Cervélo WTV procedure is executed at least three times during each day of testing to ensure that the data can be relied upon to make informed design decisions.

Specifically, the objectives of the WTV procedure are to:

- Confirm that the tunnel is producing drag measurement data within the general range of values expected based on past experience
- Quantify the measurement uncertainty associated with results from the same test session
- Confirm that results from the same test session are consistent (within the uncertainty for that session)
- Over time, quantify the difference between current results and results of other sessions, to develop a database for statistical evaluation.

For the WTV procedure to be a useful tool, it must minimize and control uncertainty associated with installation of the bike on the test rig in such a way that problems with the wind tunnel measurement system will be apparent immediately. An important output from the WTV procedure is a clear indication of whether or not the wind tunnel is functioning correctly for the planned testing. This section describes the factors considered in validation and then presents a step-by-step outline of our protocol.

5.2.1. Sources of Uncertainty in Wind Tunnel Drag Measurement

Because the chief goal of the WTV protocol is to control uncertainty, it is valuable to review the concept and sources of uncertainty. The traditional definition of measurement error is the difference between a measured value and its "true" value. Because we usually don't know the true value of a physical observation (such as drag), it is useful to place bounds on an interval which we believe to contain the true value. The term "uncertainty" is used to refer to a possible range of values that an error may have.

Within the context of wind tunnel testing of bikes, uncertainty in drag measurements can be traced to two sources:

1. Inconsistent installation of bike on force balance
2. Issues with measurement instrumentation.

During consecutive installations on the test rig, the orientation and position of the bike, rider mannequin, and mounting struts varies slightly with respect to the tunnel geometry and direction of air flow. Variables associated with installation uncertainty include:

- Yaw, roll, pitch, and steering angle of the bike
- Position of rider mannequin with respect to the bike
- Length of strut-mount adapter and thus number of exposed adapter threads
- Configuration/orientation of struts (height, wheel-base separation, orthogonality with respect to each other and wheel rollers).

Uncertainty associated with instrumentation is due to inherent errors of the devices used to measure drag and monitor wind tunnel operating conditions. Examples of instrumentation that contribute to measurement uncertainty are:

- Force balance (the device that measures aerodynamic forces exerted on the bike/rider)
- Devices used to measure properties of air (velocity, temperature, pressure, density).

5.2.2. Strut Positioning Jig

One strategy for minimizing uncertainty is the use of a strut positioning jig. Prior to each day's testing, wind tunnel staff use this jig to verify the alignment of the struts. By ensuring consistent positioning of the struts, this step (in essence a calibration) helps limit uncertainty associated with strut height, strut orthogonality, and wheelbase spacing.

5.2.3. Strut Tares

Once the struts are correctly positioned, the next variable to be controlled is the drag contribution of the struts, which is measured in a strut tare. A typical strut tare run consists of measuring the drag over a range of yaw angles with no bike/rider mounted to the force balance. The measured drag at each yaw angle is due only to the aerodynamic forces exerted on the struts and reference strut mounts. The struts-only drag is subtracted from the total drag measured during production runs (i.e. runs with the bike/rider mounted to the force balance). Figure 21 shows the difference between tared and untared drag profiles from a production run.

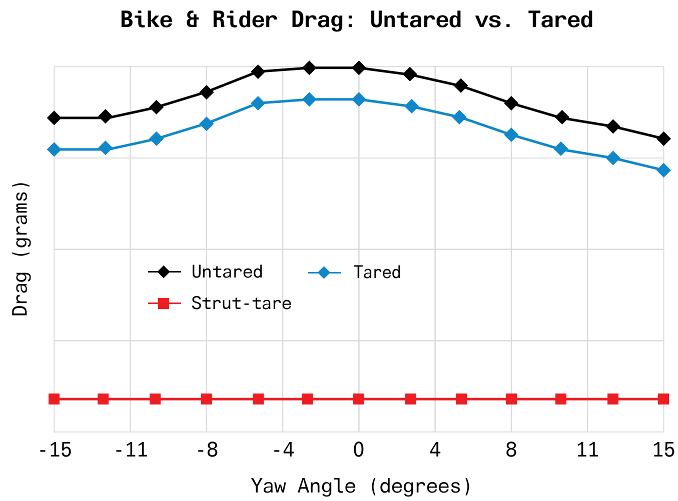

Figure 21. A comparison of bike and rider drag values from which strut drag has and has not been removed (i.e. tared and untared, respectively)

5.2.4. Wind Tunnel Validation Tests

Cervélo uses two-stage validation: A strut tare test (described in Section 5.2.3) is followed by a symmetric bike test.

The strut tare test compares current drag values from the strut tare run against those obtained during previous wind tunnel trips. This comparison gives the first indication of wind tunnel performance by showing its ability to measure drag of a simple configuration consisting of only the mounting struts and reference strut mounts.
If the strut tare test is satisfactory, we conduct the symmetric bike test. The "symmetric bike" is a simplified configuration of the Cervélo T1 (Figure 22). The T1 is used because its geometry is symmetric about its longitudinal axis. This test verifies the ability of the wind tunnel to measure a higher range of drag values (beyond those reached in the strut tare test). Again, the results of the current symmetric bike test are compared to the results obtained during previous trips.

Figure 22. A simplified configuration of the Cervélo T1 used for the symmetric bike validation test (RWDI)

5.2.5. GO/NO GO Test Criteria

The ultimate outcome of the validation protocol is a decision about whether to proceed with the planned test program. Working from the cumulative historical data collected from the symmetric bike test, a statistical treatment has been used to produce basic "GO/NO GO" metrics, such as the mean drag and standard deviation for each angle in the drag–yaw sweep. At the "GO/NO GO" decision point, current drag values for WTV tests are compared against these statistical criteria. Figure 23 and Figure 24 illustrate examples of successful (GO) and unsuccessful (NO GO) WTV tests. For the successful (GO) test example, all the current drag values lie within three standard deviations (±3σ, ~99% confidence interval) of their historical mean drag values, whereas for the unsuccessful (NO GO) test example, one or more drag values lie outside of the expected 3σ limits.

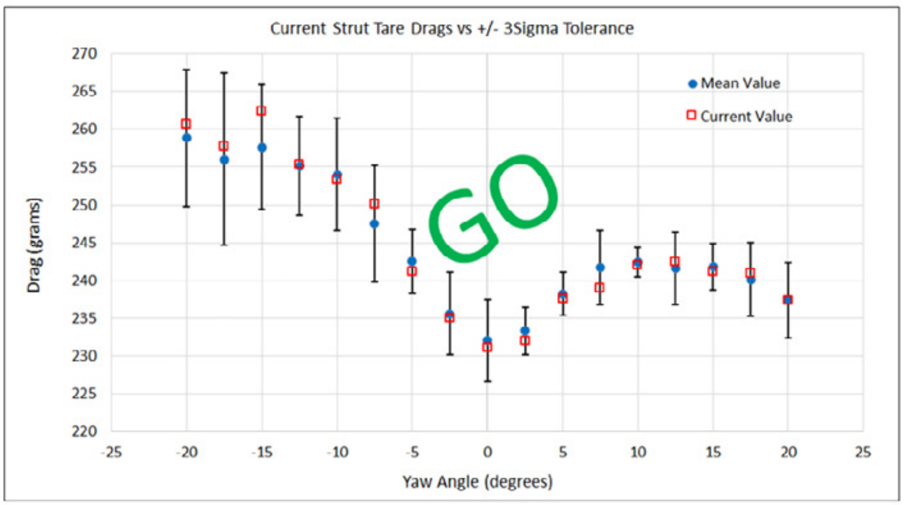

Figure 23. An example of a successful (GO) wind tunnel validation test

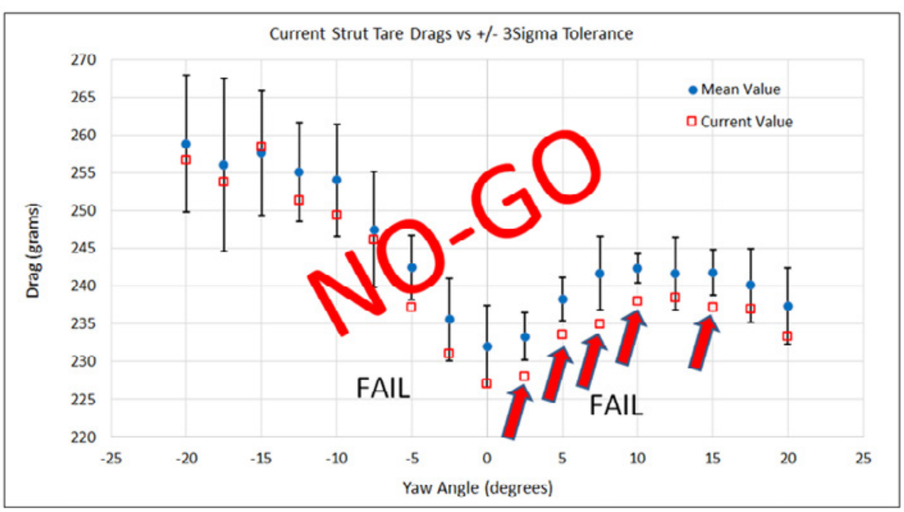

Figure 24. An example of an unsuccessful (NO GO) wind tunnel validation test

5.2.6. Wind Tunnel Validation Protocol

The following sequence describes the step-by-step protocol for setup and validation of the wind tunnel force balance in preparation for a session of wind tunnel bicycle testing.

1. Complete installation and verification procedure for bicycle force balance rig. (Each tunnel operator has its own installation protocols and testing procedures.)
2. Install strut positioning jig to verify alignment of struts and rollers after installation of force balance rig.
    a. Adjust struts and rollers as needed.
3. Run strut tare test with reference strut mounts.
4. Compare strut tare results with statistical data and make GO/NO GO decision.
    a. If GO, add this test data to stats history and proceed to next step.
    b. If NO GO, discard test data, diagnose/adjust rig setup, and repeat strut tare test.
5. Run symmetric bike test with simplified Cervélo T1 bike.
6. Compare symmetric bike test results with statistical data and make GO/NO GO decision.
    a. If GO, add this test data to stats history and proceed to planned testing.
    b. If NO GO, discard test data, diagnose/adjust rig setup, and repeat symmetric bike test.

Steps 3 through 6 are done at least three times during each test session (usually one day of testing), at the beginning, middle, and the end of the day. A revalidation may be done if unexpected results occur repeatedly, as a way of tracking down a problem with the tunnel.

5.3. Cervélo Tunnel Testing Practices

By Robert Pike

So far, we've reviewed the equipment in the wind tunnel and described what we do in advance to ensure that the wind tunnel and associated equipment are functioning properly. This section concerns what we do to prepare for and execute the actual tests. In this phase we have two priorities. First, as in the validation protocol, our constant focus is to eliminate all sources of variability other than the feature we're testing. Second, we must be absolutely efficient in how we use the time at the tunnel. Meticulous planning is required to make sure we do the right tests, as fast as possible, without compromising the quality of the results.

5.3.1. Testing the Rider

We've discussed how we demonstrate the validity of data when a bike is tested by itself. But testing a bicycle alone does not show what it will do in the real world. Adding a rider changes the frontal area presented to the air flow and creates complex aerodynamic interactions with the bicycle. In wind tunnel tests, we need to mimic those two effects by including a human shape on the bike. Validating data for a bike by itself is hard enough, but adding a rider makes it even harder to get repeatable data.

5.3.1.1. Why We Don't Use Human Riders

Cervélo has been using a purpose-built mannequin to simulate a rider during tunnel tests for more than 10 years. Long ago, we discovered that using a live person in the tunnel was not appropriate for design development, for two key reasons.

The first relates to the repeatability and stability of the rider's position during the test. We need the rider to be very consistent both during a run and from one run to the next; that's the only way we can be sure that a difference in drag is due to the bike alone. It's very difficult for a rider to perfectly hold a position through the yaw sweep of a tunnel run. During the yaw sweep, the turntable starts and stops and the wind blows from different directions; these changes can make the rider uncomfortable. Even if an elite rider can be very stable and consistent in a single run, imagine all the variation when the rider gets off the bike and onto it again…and again. When we are studying small changes in the shapes of our frame tubes or bars or seat posts, even microadjustments by the rider will throw off our drag comparisons. A moving live rider overshadows the results we need to see.

The second issue for design development is the effect of a rider pedaling during the test. The measurement balance naturally experiences some oscillation as it rotates through the yaw sweep; the measurement variation caused by this oscillation is called the error band. The movement of a rider's legs widens the error band, creating more opportunity for misinterpretation of the data. As with changes in position, a large error band can overshadow the actual comparisons we are looking for.

Any test has inaccuracies, but the goal of the testing team and the tunnel operators is to minimize the opportunity for inaccuracies. For help, we turn to DZ.

5.3.1.2. Introducing DZ, Our Model Rider

A description of our work would not be complete without introducing a key member of our team: our wind tunnel test mannequin—"DZ," as we affectionately call him. DZ is a highly accurate, life-sized model of American Time Trial specialist Dave Zabriskie (Figure 25). In 2007, Dave was laser scanned on a bike in his time-trial position, and his twin was machined from dense foam in the position needed for testing our time-trial bikes. The foam was encased in a hard epoxy layer to help keep the air flow smooth and offer some protection against damage. DZ's arms and legs can be removed for storage and transport, and he lives in a couple of foam-padded travel cases.

Figure 25. DZ rider mannequin (LSWT, San Diego USA)

DZ has threaded inserts on his feet that allow them to be screwed to the pedals during testing. His knees and elbows are fixed, and his legs and arms are also designed to stay in one position. During our bike setups, we ensure that DZ's position on the saddle is constant from test to test and that his arms are located and positioned securely in the arm pads. These three contact points offer fantastic repeatability in our setup.
We do allow a certain amount of travel in DZ's hip joints, as we need to have his back angle change for different bike setups. For example, for a time trial bike we may put him in low or high position; for a road bike setup, two or three position adjustments are often required. It's permissible for his position to change from one bike style or fit to another as long as his position is exactly the same within a series of runs that will reference each other.

5.3.2. Preparing for a Wind Tunnel Trip: Planning for Consistency

Thorough preparation prior to each wind tunnel trip is critical to the success of the testing sessions and the accuracy and value of the data collected. Tunnel time is expensive and limited, so efficiency in every aspect of the test day is very important. Something as simple as not having the correct stem spacer or the correct tire size can disrupt the workflow and compromise the data.

5.3.2.1. Components and Assembly

As we prepare for a wind tunnel trip, our first task is to put together a run list (described in more detail in Section 5.3.3). The run list is the key to the whole trip, but in the planning period its function is to tell the setup team which bikes to prepare and how to configure them. Thinking through the fitting requirements in advance is critical as it helps to ensure that the bikes are set up in such a way that the data we collect will answer the question we're asking. If, for example, we want to compare our new frame design with one of our other frame designs, we need to ensure that every other aspect of the bike build is identical. We make sure to use the same wheels, tires, derailleurs, brakes, handlebars, stem, seat post, saddle, and so on. We also make sure the bikes are assembled in exactly the same way. Some examples of such adjustments are the height, setback, and angle of the saddle; the position of the handlebars and the levers on the bars (or, in the case of a time trial bike, the position of the arm pads and extensions); and the stack and reach of the front end. We even ensure that we have a consistent chain length: Details that small can affect the position of the rear derailleur and as a result change the drag data. By controlling all these aspects of assembly, we can isolate the source of any change in the frame drag numbers.

5.3.2.2. Position of DZ

Creating identical setups is even more critical when it comes to mounting our test mannequin on the bikes. Just like a human rider, DZ represents a significant portion of the drag of the system. Thus, the greatest opportunity for error in the drag comparisons is having DZ out of position. To isolate the effects of the change we're studying, we need to make sure that DZ sits in exactly the same position from one bike to the next. This sounds simple enough, but in the tunnel, when the clock is ticking and we're trying to make bike changes as quickly as possible, there are a lot of opportunities for setup variation. This is why it's critical to plan exactly how DZ will be moved (and, if necessary, readjusted) from one bike to the next.

Such planning is especially important when we're comparing dissimilar bikes, either from other Cervélo series or from other brands. Variations that can make a significant difference in the setup include bottom bracket drop; seat post positioning and angle; stack and reach positions for the frame; and even available stem length, spacer increments, and bar sizing. It becomes even more challenging when we want to compare "Out of Box" builds, which means we need to use all the components that come with the bike. In most cases, the best we can do is set up the bikes as close to each other as possible and then tweak the positions to ensure that DZ is in the same position for each test.

To ensure that DZ's feet, hands, and seat (his touch points) will be in the same position on each bike to be tested, we use a custom setup jig. The jig fits to a bike at the points where DZ would touch it. The procedure for ensuring positioning consistency is as follows:

- The reference bike is set to the desired dimensions (as discussed in Section 5.3.2.3). The position of components (e.g. the bars, saddle, shifters, etc.) are specified relative to the bottom bracket, following industry practice for fit.
- DZ is mounted on the bike, adjusted to the desired position, and removed from the bike.
- The reference bike is set in the jig, and sliding stops are adjusted to match the touch point locations.
- Each bike in the test series is placed in the jig and adjusted so that the touch points match the jig and thus are in the same location as on the reference bike. If the bike can't be adjusted completely, the necessary adjustments to DZ are noted.

This check is conducted at Cervélo as we prepare bikes for a trip to the wind tunnel, saving time at the test session.

5.3.2.3 Stack and Reach Comparisons

Comparing bikes with different geometries is a tricky business because we can't hold all the variables constant. Whether we're comparing bikes within the Cervélo range or with other brands, we need to ensure that our approach to the comparison is appropriate and meaningful, despite the variation. We make these tests as accurate and fair as possible, especially when comparing to other brands. At Cervélo we rarely, publish our comparisons to other brands, but rather choose to compare results to our previous offerings. We do test our bikes against other brands, and we have done so for years, but generally our purpose in those tests is to ensure that our offerings remain the fastest on the market.

One of the challenges of testing different geometries is the range in stack and reach points and in stem lengths and stem spacer options across the industry. How do we deal with this? Our first priority is to get DZ in the same position on each bike. We identify, among the bikes to be compared, the bike with the highest "slammed" position (i.e. stem snugged up to headset). That bike in that position becomes our "reference bike," against which the rest of the bikes are adjusted. We choose the highest position because our only option for matching stack is to add spacers and move stems up. Unfortunately, because our top-end bikes offer a lower stack than many other brands, we usually have to add spacers to our own bikes when testing. If another brand's offering is closer to ours in frame stack and reach point, we'll create an additional reference position and make a separate comparison. Using a different reference point in this way is appropriate, as we are most interested in the change between drag numbers, rather than absolute values. As long as each test series has a baseline reference point, then the comparisons are fair and true.

5.3.3. The Run List

The backbone of every day at the wind tunnel is the run list (Figure 26). Its primary function is to list the bikes we want to run in the day, but it also includes a detailed schedule for the day and a reference for every aspect of each bike, including build, setup, and test protocol. Thus, the run list is a comprehensive map of everything we plan to do; that's why starting to develop run lists is one of the first steps in preparing for a wind tunnel trip.

- Run lists help us complete the following tasks:
- Review the needs of groups within the company (e.g. product development, marketing)
- Define a logical and efficient test sequence, taking into account, for example
- Bike style
- Mounting technique
- Positioning of rider mannequin (DZ)
- Test duration
- Data needed to inform later runs
- Document build details
- Simulate usage scenarios
- Communicate for travel and setup.

Armed with a well-thought-out run list that anticipates potential delays, we have a better chance of conducting an efficient test session that produces accurate and useful data.

The run list often evolves over the course of the day or the test session, as Figure 26 shows. Sometimes we have to change the sequence on the fly, based on results coming out of the tunnel. Other changes include annotations, such as changes to DZ's positioning, that may help us interpret the data when we return to the office.

Series #	736	736	736	736	736	736	736
Session Run #	28	29	30	31	32	33	34
Daily Run #	1	2	3	4	5	6	7
Description	CALIBRATION	Bike 1 Storage Config 1	Bike 1 Storage Config 2	Bike 1 Single Bottle BTA	Bike 1 Single Bottle CT	Bike 1 Single Bottle Rear with Pouch	Bike 2 Storage Config 1
Frame Configuration							
Notes	Bike Change	Bike Change					Bike Change
WB Change		Yes					Yes
DZ Prep Back Angle		-2.9					-2.9
DZ Tunnel Back Angle		-2.9					-2.9
Beta Schedule	89	87					87
Done?	End 8:55	End 9:32	End 9:49	End 10:10	End 10:27	End 10:45	End 11:01
Frame	T1 Aluminum	Bike 1					Bike 2
Size	56	54					Medium
Wheel Front	HED Ardennes	Enve 7					Enve 7
Tire Front	GP4000 25	GP4000 25					GP4000 25
Wheel Rear	HED Ardennes	Enve 8					Enve 8
Tire Rear	GP4000 25	GP4000 25					GP4000 25
Brakes Front		Eng Disc					Eng Disc
Brakes Rear		Eng Disc					Eng Disc
Fork		Model Spec					Model Spec
Bar		Model Spec					Model Spec
Cables		Yes					Yes
Extensions / Stem		Ex11					Spec DZ
Cranks		Durace 52/36					Durace 52/36
Cassette		Eng Spec					Eng Spec
Saddle		Eng Spec					Eng Spec
Post		Model Spec					Model Spec
Water Bottle		See Notes	See Notes	See Notes	See Notes	See Notes	See Notes
DZ Position		TT Low					TT Low
Time per Run	25.00	19.00	19.00	19.00	19.00	19.00	19.00
Time Total (min)	25.00	44.00	63.00	82.00	101.00	120.00	139.00
Time Total (hr)	0.42	0.73	1.05	1.37	1.68	2.00	2.32
Scheduled Start Time	8:30:00	9:14:00 AM	9:33:00 AM	9:52:00 AM	10:11:00 AM	10:30:00 AM	10:49:00 AM
NOTES		BTA 750 Bento 200/600 DT Bottle Round 750 Rear Bottle 750 & Aero Pouch	BTA 750 Bento 200/600 DT Storage Aero 500 2 Rear Bottle 750	BTA 750 Bento 200/600	Bento 200/600 DT Bottle 750	Smartpak 200/600 Rear Bottle 750 & aero pouch	BTA Bento DT Bottle 750 1 Rear Bottle

Series #	736	736	736	736	736	736	736	736
Session Run #	35	36	37	38	39	40	41	42
Daily Run #	8	9	10	11	12	13	14	15
Description	Bike 2 Storage Config 2	Bike 3 Single Bottle SC	Bike 3 Single Bottle DT	Bike 3 Storage Config 1	Bike 3 Storage Config 2	CALIBRATION	Bike 4 Storage Config 1	Bike 4 Storage Config
Frame Configuration								
Notes		Bike Change				Bike Change	Bike Change	
WB Change		Yes				Yes	Yes	
DZ Prep Back Angle		-2.9					-2.9	
DZ Tunnel Back Angle		-2.9					-2.9	
Beta Schedule		87				89	87	
Done?	End 11:17	End 11:34	End 11:51	End 12:09	End 12:26	End 12:58	End 01:17	End 01:35
Frame		Bike 3				T1 Aluminum	Bike 4	
Size		Medium				56	54	
Wheel Front		Enve 7				HED Ardennes	Enve 7	
Tire Front		GP4000 25				GP4000 25	GP4000 25	
Wheel Rear		Enve 8				HED Ardennes	Enve 8	
Tire Rear		GP4000 25				GP4000 25	GP4000 25	
Brakes Front		Eng Disc					Eng Disc	
Brakes Rear		Eng Disc					Eng Disc	
Fork		Model Spec					C/W BC	
Bar		Model Spec					Model Spec Low	
Cables		Yes					Yes	
Extensions / Stem		Spec DZ					Spec DZ	
Cranks		Durace 52/36					Durace 52/36	
Cassette		Eng Spec					Eng Spec	
Saddle		Eng Spec					Eng Spec	
Post		Model Spec					Model Spec	
Water Bottle	See Notes	See Notes	See Notes	See Notes	See Notes		See Notes	See Notes
DZ Position		TT Low					TT Low	
Time per Run	19.00	19.00	19.00	19.00	19.00	25.00	19.00	19.00
Time Total (min)	158.00	177.00	196.00	215.00	234.00	259.00	253.00	272.00
Time Total (hr)	2.63	2.95	3.27	3.58	3.90	4.32	4.22	4.53
Scheduled Start Time	11:08:00 AM	11:27:00 AM	11:46:00 AM	12:05:00 PM	12:24:00 PM	12:49:00 PM	12:43:00 PM	1:02:00 PM
NOTES	Bento Rear Bottle 750	Bento SB SC with Bottle 750	Bento SB DT Bottle 750	BTA Bento SB No SC DT Bottle 750 1 Rear Bottle	Bento SB Rear Bottle 750		BTA Stealth Bento Gels DT Bottle 750 1 Rear Bottle & Aero pouch	Stealth Bento 1 Rear Bottle & Aero Pouch

Figure 26. A typical run list for a wind tunnel session

The following sections elaborate on some of these preparation tasks:

5.3.3.1. Defining an Efficient Test Sequence

Once the list of runs is built, we start looking at how they relate to one another and how best to order the runs within a visit for maximum efficiency. For example, if the tunnel requires a fairly time-consuming change to convert from mounting quick-release wheels to thru-axle wheels, then in the interest of efficiency, we'll group the two types together to avoid repeated setup changes. The same applies to tire sizes and wheelbase changes. Our mannequin has two arm styles, which take time to change out, so as much as possible we group the bikes with this in mind.

Our preference is to group bikes as much as possible within a style series. For example, if we are comparing a set of road bikes with the same setup, we prefer to test them one after the other rather than scatter them throughout the day. This can't always be done, for the reasons just mentioned, but whenever possible, we do so. This approach improves the consistency of data recording. But more importantly, it helps our setup team get into a groove with a particular setup style, which also improves consistency. And of course, we prioritize the runs, placing the less critical runs toward the end of the session, in case we run out of time and some tests need to be dropped until the next visit.

The run list reflects a pretty accurate estimate of the timing for the day. This estimate is based on the actual testing period for a particular run, which varies with the number of yaw angles we're testing (see Section 5.3.7), and on an estimated average time for changing out the bikes. The change-out time is tricky to predict, partly because of the variability in what needs to change, from full bike swaps to a simple water bottle change, and partly because small problems in the tunnel can chew up valuable time. We are constantly looking for opportunities to make the changes faster and more efficient while ensuring accurate setup.

5.3.3.2. Documenting Build Details

The run list is a master document that outlines every part that is built on the bike: the derailleurs, brakes, wheels, cranks, seat post, bars, stem, chain, tire size and style, tire pressure, and more.

Builds are specified in two general categories. When we keep certain parts consistent from run to run, we are using what we call engineering specification, or Eng. Spec. By keeping these items consistent, we can isolate the change that is due to the parts we are studying during shape development. Although components do differ for different styles of bikes, we keep test subjects within a series as consistent as possible. The alternative to Eng. Spec is the “Out of Box” build. In this case we are comparing bikes as the consumer receives them, so we can rank the bikes aerodynamically. We compare Out of Box builds both within our own series and with competitors' offerings.

These two categories of specifications require two styles of testing, as the requirements for the data vary. When developing shapes and studying interactions, the development team is looking for small gains and telltale signs of improvements. The team's goal is to search out and validate reductions in drag that can be combined in the design of a faster bike and rider. To reveal these small changes, the data collected has to be as clean and as accurate as possible. The Eng. Spec helps ensure a clean and simple comparison. With the Eng. Spec, we usually build the bikes without cables or hoses.

When the product cycle shifts from early development to validation and launch support, we start looking at Out of Box testing. This approach emphasizes the bike as a complete package. At this stage, cables, bar tape, and so forth begin to figure in the picture. The comparisons here are broader, looking at overall gains that are a sum of all the finer development gains, so we can accept a little more "noise" in the results.

5.3.3.3. Simulating Usage Scenarios

Another aspect of testing is to consider usage scenarios. As standard practice, we test our road bikes with water bottles, as we want to design the best bike for the way it will be used. Such testing is even more important when we're studying triathlon bikes. We spend a lot of time looking at how triathletes use their bikes and what they need to carry with them. A good example is the concept of the integrated storage on the P5x. When preparing triathlon bikes for tunnel testing, we work hard to set them up the way riders store their gels, bottles, repair kits, and so on. And of course, we do our very best to be fair when comparing to other brands. We've even contacted other companies to ask how they suggest we set up their bikes, to offer their fastest setup to our test.

5.3.3.4. Communicating for Travel and Setup

The run list serves as a sort of pre-flight checklist. It serves as a valuable reference for ensuring we pack everything we need, which is especially important given the time sensitivity of testing. It also helps in gathering prototype and newly produced parts in timely way. The tunnel staff use it to ensure the support will be available and to prepare for data collection.

5.3.4. Photo Reference

Once we're at the tunnel, a key tool for verifying the positioning of the bike and DZ is photography. We set up a digital camera in a fixed location at the tunnel; the view is through a window into the test section, perpendicular to the bike under test. A small whiteboard showing the run number (from the run list) is placed in the frame of the picture (Figure 27). The photo visually documents exactly what was tested in each run, giving us a reference in case there are any test anomalies or documentation problems. Probably the most important function of the photographs is to help verify DZ's position on the bike. If we overlay two pictures from different runs (which is possible because the camera is in a fixed position), it is possible to compare the exact position of bike and rider (DZ) for the runs. We do this on the fly at the tunnel to help confirm position and also later when comparing data from different runs.

Figure 27. Run photos showing P2 bike configuration with and without water bottle (RWDI)

5.3.5. The Tunnel Team

A very important part of testing at the wind tunnel is the testing team. For any given trip, there are usually three to six people involved (Figure 28). One person is designated as the test leader. That person ensures that everyone knows what needs to happen during the day, communicates with the tunnel personnel, manages timing and schedule, and tracks and updates runs on the run list.

Figure 28. Tunnel team at work (RWDI)

In addition to the test leader, the team usually includes two or three more people who handle the bikes during the testing. They are responsible for taking the bikes and DZ in and out of the tunnel, as well as for setting up each bike and preparing it. There's more to a changeover than simply moving the bikes. In many cases, the team needs to switch parts between bikes, tape on 3D-printed prototype parts, run cables in frames, and see to many other such details. They may even have to quickly change a blown tube. Yes, it's possible to get a flat in the tunnel!

Actually, changing out the bike in the tunnel is an art in itself. Two people work together to carefully remove DZ. Another person (or two) removes the bike and then replaces it with the new one, all while someone holds the 40-kilogram DZ. All of this needs to be done very carefully (applying too much force to the balance can physically damage it) but also very quickly (every minute has a price tag). Some changes happen very quickly; for example, adding a bottle and cage can be done in less than a minute. On the other hand, a full bike change can take 10 or 15 minutes.

As part of developing the run list, we also plan what will be happening outside the tunnel. While one run is happening, people are setting up the bikes or parts needed for the next run.

5.3.6. Onsite Adjustments to Bike and Rider

The bike is held to the balance by two sets of struts—front and rear. Different tunnels have slightly different struts and axle interfaces, but they all generally mount to a dummy axle, which then replaces the normal axle in the bike. In the past this was a pretty common setup, since all road bikes used standard quick-release (QR) axles front and rear, with a standard width (100mm front, 130mm rear). In recent years changes in axle configurations have introduced many new variables: disc brakes, 12mm thru axles, wider hub spacing, and many others. As a result, there is more work to ensure that any given bike can fit properly on the balance. The axle variations are accommodated by adapters installed on the strut. When testing one style of bike this is not a problem, but when testing competitor bikes it can have a big impact on setup time and complexity, because many different adapters must be switched in and out.

Once the front and rear axles have been secured in the struts, we make two types of adjustments. The first is to the wheelbase setting of the balance. The tires sit on rollers that are contained within the balance (Figure 29). The wheelbase (distance between the two wheels) of the balance rollers must be adjusted to match the bike being tested. This distance is different for almost every bike and so is a regular adjustment. Most tunnel balances allow relatively easy movement of one or both sets of struts.

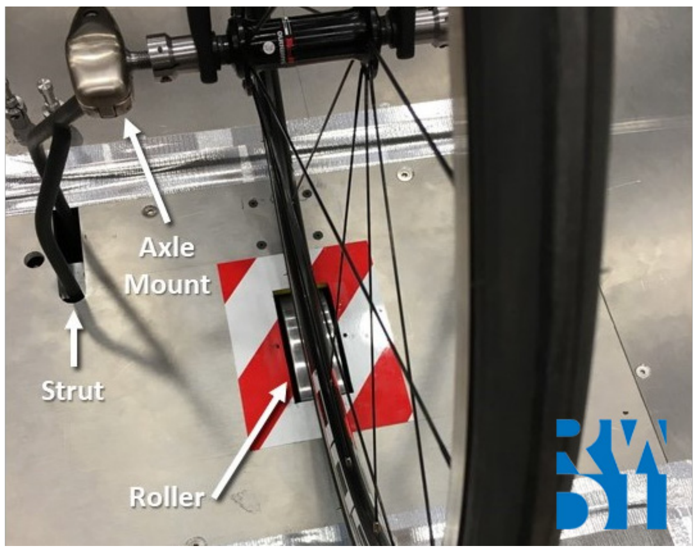

Figure 29. Strut and wheel roller detail (RWDI)

The second adjustment is to the height of the struts. This dimension is based on the tire diameter and changes with wheel style, tire, inflation pressure, and other factors. Because the rollers are powered and drive the tires, there must be enough normal force holding the tire to the roller to prevent slipping. The wheels rotate at the same speed as the air flow, which is normally 50 km/h. Thus, quite a high preload on the tires is required to stop them from slipping. Ensuring this preload is consistent and sufficient is done solely by adjusting the axle height; getting this adjustment right takes quite a bit of experience and patience. For this reason, we use consistent wheels and tires to avoid having to make this adjustment every run. Again, this adjustment can be an issue when testing competitor bikes or Out of Box spec.

When setting DZ up on the bike, we have to ensure that all of his joints are positioned properly. This adjustment is normally done by matching sets of marks at each joint. We also double-check his position on the bike by measuring his back angle at a specific location with a digital level. The back angle will change for different fits (i.e. road vs TT) but needs to be the same (within a small tolerance band) so that the drag results are consistent.

5.3.7. Beta Schedule (Yaw Angles)

The key variable in drag measurements is the yaw angle, as discussed in Section 2.2. When testing in the wind tunnel, the yaw angle is often referred to as the beta angle. (This term is a holdover from aerospace terminology of the past.) Because we're interested in how a bike will perform in both good and bad wind conditions, the question becomes "How many yaw angles should we measure in the tunnel?" It is generally accepted that the most important real-world yaw angles (those most often experienced by cyclists) occur in the range of ±15 or 20 degrees (see more discussion on this in Section 5.3.8). So, do we measure every degree? Every 5 degrees? It takes roughly a minute to record data for a single yaw angle (actual recording time plus stabilizing time). Thus, part of our decision process in developing the run list is to choose a number of yaw angles that will give us enough information to understand how the bike behaves without taking too much time in the tunnel.

We define our test points with what we term the beta schedule. An example is:

$$\text{B7-R1515}$$

This notation means that we will take seven data points over the yaw range of +15 to –15 degrees, equally spaced, for measurements at yaw angles of 15, 10, 5, 0, –5, –10, and –15 degrees. This is the typical range we use for most of our testing. It provides enough detail to understand the behaviour of the bikes but minimizes test time. For more specific tests or for research where we need more detail, we will use different beta schedules, such as

$$\text{B9-R2020} \quad (\pm 20^\circ \text{ every } 5^\circ)$$
$$\text{B15-R2020} \quad (\pm 20^\circ \text{ every } 2.5^\circ)$$

5.3.8. Presentation of Average Drag

Our next concern is how to present the data so that they are easy to interpret. Our practice is to produce "drag vs yaw" plots that show individual data points for drag at each yaw angle, with straight lines between the points. While we could do a more exact curve fit of the data, that would take more time and introduce more variables. The straight-line approximation between points shows trends easily and is consistent with our own previous work and with what most other manufacturers and tunnels present. Although these plots are good for showing trends, they can be hard to interpret. For broader purposes, it is far more useful to represent drag in a single number. Having one number makes it much easier for marketing staff, dealers, and consumers to compare models. We generally report an average drag across a specific yaw range and then compare these average values for different bikes.

But how should we calculate this average? For our standard B7–R1515 beta schedule, we simply mathematically average all the data points to give an average drag value. This is the top-level drag result that we report for our bikes. But condensing a drag vs yaw plot into one number inherently introduces approximations. In this standard average, all yaw angles are considered equally. How does this compare to the real world?

Much work has been presented by other manufacturers and in scientific journals regarding estimated distributions of yaw angle in real-world riding. Some researchers have performed measurements of specific courses at different times and with different riders to measure the actual distribution of yaw angles. However, these data are specific to those places and times. How can they be applied generally to other situations? Other authors have used a more random distribution of wind directions, which gives a different distribution of yaw angles. In all cases, some function is used to weight each yaw angle differently when calculating average drag. The bicycle industry, or at least some key players within it, seems to agree that narrower yaw angles should be weighted more highly than the wider angles, because the propulsion of forward motion will tend to correct wider yaw angles, biasing the distribution toward smaller angles.

At Cervélo, we are continuing to perform research in this area but have not yet determined to our own satisfaction what averaging method best represents the real world. As a result, until we have learned more, we still present our average values as a straight average, so that they are comparable to our previous test results. Keep in mind that this approximation does not affect our graphically presented data in any way, only the single average values reported. We still take into consideration the shape of the recorded drag curves and adjust our designs based on the expected use scenarios for different types of bikes.

5.4. Wind Tunnel Testing as a Development Tool

By David Killing and Richard Matthews

5.4.1. Why We Use the Wind Tunnel

Wind tunnel testing has always been a core part of Cervélo product development. From the very early days, the wind tunnel has been used as a tool for developing more aerodynamic bikes, not just as a marketing strategy.

Over time, we've added CFD and real-world track testing to our arsenal of development tools, but the wind tunnel still has a huge role in our development process because it holds some key advantages over other tools.

- It allows for accurate testing of real-world products (not just simplified CAD models, as used for CFD).
- A detailed part is not harder or more expensive to run than a simple part.
- Because we can test the bike "as built," it provides the best representation of the final aero performance for the customer.
- It is much easier to run benchmarks with competitor products in the wind tunnel than to build CAD models.
- It can be used to validate initial development done in CFD.

5.4.2. Wind Tunnel Testing in the Design Cycle

It is clear how the wind tunnel can be used to test final designs and competitors, but how does it fit into the design process? As described in Section 4, CFD is well suited to investigating design concepts and details, as they can be analyzed quickly from CAD data. However, to test a concept or idea in the wind tunnel, we need an actual physical part, and the part must be integrated into a bicycle in a realistic way. Even so, the wind tunnel is still a valuable tool for evaluating designs early in the development process. We often use the wind tunnel in conjunction with CFD to evaluate design concepts.

In a typical design evolution, we will first do many quick evaluations with CFD. If this exploration yields several possible designs but no clear differentiation in performance, or if there are major conceptual differences that we want to evaluate, we will then plan to wind tunnel test some of these ideas or designs. Wind tunnel testing gives real-world results in which we have very high confidence and which can be used to make clear decisions on design direction.

5.4.3. Wind Tunnel Experimental Techniques

To get clear direction for design, we have to ask the right questions and conduct the kind of tests that will answer those questions. This section will discuss how we go about testing new ideas in the wind tunnel. The key advantage of this work is that we're testing a real bike—except that most of the time, the bike can't be 100 percent real. It is not practical to 3D print a full new bike model for each little detail we want to investigate. Instead, we use the following techniques to get around this inconvenient limitation while keeping the test bike as realistic as possible. They are described in order of increasing fidelity to the final bike.

5.4.3.1. Add-ons to Real Parts

A quick method for testing specific areas or components is to add on 3D printed pieces to existing frame or fork parts (Figure 30). We use this method to get a very fast and accurate comparison of drag with and without the add-on piece. This test is quick and easy: Pieces can often be added with the bike still on the balance, making the testing much faster and more accurate. The downside is that it only works when the shape to be tested is larger than the underlying part.

Figure 30. Add-on part (seat tube variation) in tunnel (LSWT)

In a variation of the add-on technique, we sometimes create the add-on piece on the spot out of hand-molded clay (Figure 31). This method is perfect for quickly checking ideas while at the tunnel, but it can be difficult to replicate the results afterwards. In addition, it takes a person with good sculpting talent and aerodynamics experience. (Luckily, we have several people like this at Cervélo!)

Figure 31. Clay add-on in tunnel (LSWT)

5.4.3.2. Skeleton Frames ("Mules")

Early in design, we use skeleton frames, called mules, to try out different shapes for specific areas of a bike. In this method, a metal (usually steel) frame is covered in 3D-printed parts to "skin" it to the final bike shape (Figure 32). The metal frame is strong enough to support DZ and allows the 3D-printed skin pieces to be simpler and cheaper. This method also lets us try out different shapes in the same area. For example, we can investigate such details as how a specific radius or tube shape can be varied and how the drag changes as a result. A disadvantage of this method is that there are joints where the small skin pieces meet; they are usually taped with thin foil. There can be a lot of joints, which is undesirable because it takes longer to change out pieces and because joints may contribute effects that would not exist in the final frame. Another disadvantage is that the underlying metal frame cannot change, so variations can be tested only within that overall frame design.

Figure 32. P5 mule with 3D printed parts

5.4.3.3. Full Bike Models

A model of a complete frame (or fork, etc.) is usually the final stage of wind tunnel testing before production. We will often 3D print a full frame, including fork, bars, and so on (Figure 33). This model is very expensive (many thousands of dollars) because we have to use a very rigid and strong material that is sufficient to support our rider mannequin, DZ. A typical low-cost 3D printer cannot handle these types of materials. In addition, a full frame or fork is very large and needs the capacity of the biggest 3D printers. For these reasons we only make full models at critical decision points or to test final designs.

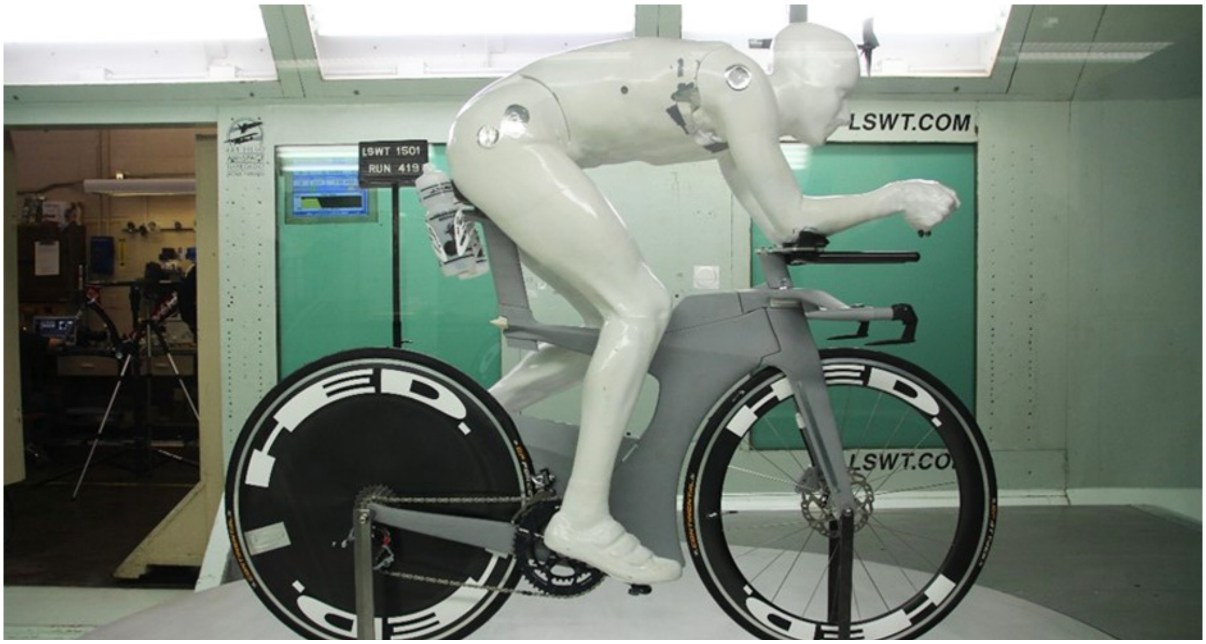

Figure 33. P5X 3D printed model in tunnel (LSWT)

5.4.3.4. Production/Prototype Bikes

In the final stage of development, we use actual production or prototype carbon parts (Figure 34). We do this test to verify final performance against preproduction results, because sometimes the move to production tooling introduces slight changes relative to the CAD shapes. This method is relatively easy, because the parts are real; however, sometimes ancillary pieces (derailleur mounts, etc.) can be prototypes. For competitor bikes we obviously test the production models, which poses the measurement challenges already discussed.

Figure 34. S5 prototype (carbon) in tunnel (RWDI)

6. REAL-WORLD VALIDATION

By Graham Shrive

After the techniques described in the previous sections have been applied to perfect the bike, one last step is required to translate those results into faster rides. This last step marries the drag from the frame with the drag from the rider, which is determined by position, clothing, helmet, and so on. The nature of the system is that rider drag dominates the overall drag equation. However, that doesn't mean the equipment is irrelevant; it just means the fit has to be precise. With the best possible fit, the rider can exploit the low-drag equipment designed using the various tools discussed earlier. To take this step into the real world, Cervélo validates product performance in tests by top athletes in a track environment.

6.1. Measuring Drag at a Track

Why do we measure at a track and not outdoors? In general, a track provides the closest thing to controlled conditions that a human can reasonably achieve on a bike in normal use. Inside a closed velodrome, the drag the rider encounters is from effectively still air, with uniformly low or no velocity. This condition translates into an effective wind velocity of zero for the rider and consequently an effective yaw of zero degrees. With these two factors removed, the track is an ideal environment in which to remove the remaining variables—other than the rider and bike—from the equation.

In real-world performance testing, the variable actually measured is the power output of the rider. Thus, to put track results in the same terms as wind tunnel results, we need to express drag in solely in terms of power output. Recall that the power output required from the rider is based on the following factors:

Eq. 6

$$P_{out} = P_{aero} + P_{RR} + P_{mech} + P_{elev}$$

where P_{out} is power output required from rider and P_{aero}, P_{RR}, P_{mech}, and P_{elev} are the quantities of power consumed to overcome the forces due to aerodynamics, rolling resistance, frictional losses in the mechanical components of the drivetrain, and changes in elevation, respectively.

Solving Eq. 6 to isolate the aerodynamic forces, we get:

Eq. 7

$$D_{aero} = \frac{P_{out} - P_{RR} - P_{mech} - P_{elev}}{V}$$

Given that the goal here is to isolate the aerodynamic drag forces acting on the rider, it is desirable to reduce the remaining variables the rider may encounter in the real world, including inconsistencies in road surface and differences in cornering effort on a real-world course. Let's consider these variables one by one.

The easiest to eliminate is rolling resistance. At a track, the surface condition is uniform above the blue line and will typically be free of paint (Figure 35). This uniformity removes—or at least normalizes—the effect of rolling resistance on power consumed.

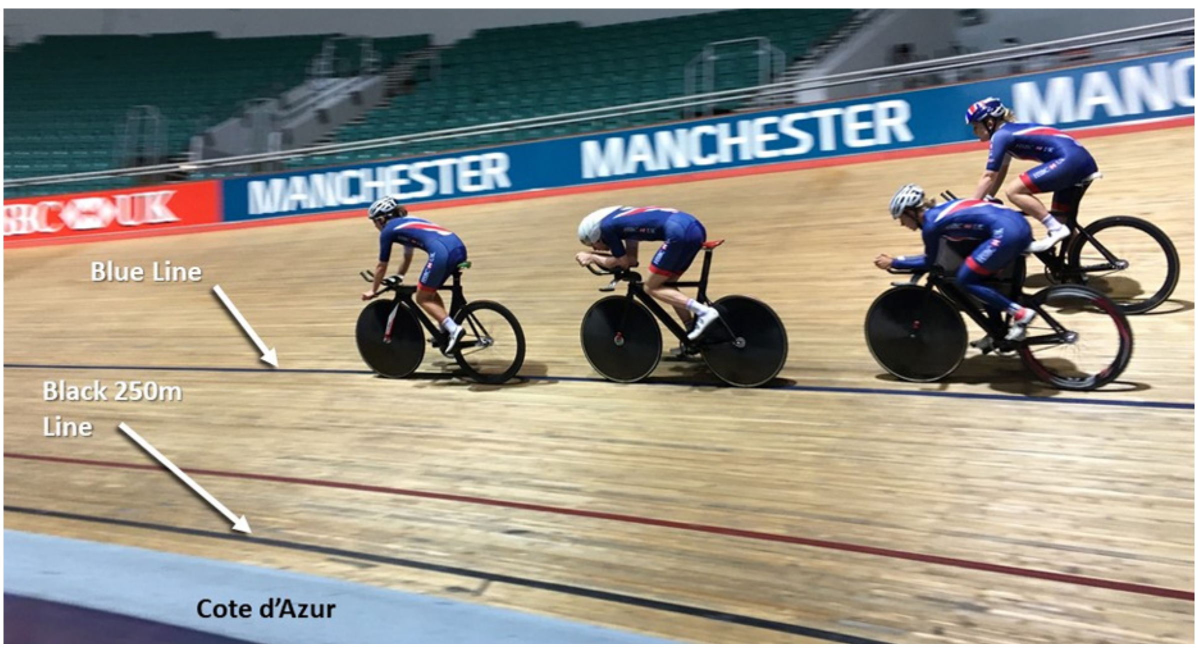

Figure 35. Regions of a velodrome track (NCC, Manchester UK)

The drivetrain losses on the bike are effectively the same across a fixed combination of power output and gearing. Thus, if we lock the rider into one gear and keep the cross-chaining to a minimum, we can disregard drivetrain losses, or just consider it a variable with an unchanging value. Based on the these arguments, we can assume that power consumption due to rolling resistance, mechanical losses, and changes in elevation is unchanging in the controlled test environment of the velodrome, and hence these quantities are replaced by constant C_1. Equation 7 can then be simplified as follows:

Eq. 8

$$D_{aero} = \frac{P_{out} - C_1}{V}$$

Finally, we can combine Eq. 8 and Eq. 1 and rearrange to obtain the following expression for C_DA:

Eq. 9

$$C_DA = \frac{P_{out} - C_1}{\frac{1}{2}\rho V^3}$$

This is where things get tricky—and where the more advanced systems for real-time aerodynamic measurement start to shine. The prevailing assumption is that a rider who maintains a generally straight disposition relative to the black 250 meter line at the track will effectively neither gain nor lose elevation, keeping that variable near constant. The challenge here is that as the rider traverses the banks in the track, there is actually some degree of work being done by the body (and thus power is consumed) as the body is rotated off of the normal axis. Furthermore, no one is perfect, and even pros vary up and down from the black line by small amounts from one run to another. However, we've worked with a few very advanced drag measurement systems that successfully use an algorithm to remove this variation in elevation. For such an algorithm to work, it needs two kinds of input: (1) a mapping of the specific track where tests are being done and (2) real-time feedback from a rider-carried module that measures body roll. The pre-eminent system using this methodology was created by Alphamantis (recently purchased by Garmin) and is employed by a number of fitting services around the globe.

Having removed the elevation from the equation, we can start drilling down into the rest of the problem. When we're comparing different setups, what we're really looking for is the smallest C_DA term in Eq. 9. (Recall that this term combines the drag coefficient, C_D, and the frontal area, A. To express C_DA solely in relation to power output in Eq. 8, we also have to remove the velocity variable in order to hold power output constant. Again, there is an imprecise way and a precise way to do this, by (respectively) asking the rider to hold the same velocity or compensating in real time for the changing velocity. The Alphamantis system also measures and compensates for velocity variation.

6.2. Perfecting Fit

Having removed the variables inherent to a live rider, we now have controlled conditions in which to evaluate both changes in equipment and changes in rider position. Generally, with professional riders, equipment is relatively fixed based on sponsorship agreements, but occasionally we can make small changes that can be evaluated relatively effectively at the track. For example, we can make suggestions about helmet styles or computer locations. Most importantly, the rider's position can be evaluated by trained fitters to balance power output, C_DA, and metabolic and biological measures (e.g. lactate clearing ability, comfort, and blood oxygen levels). At the track, different fits can be evaluated not simply for raw aerodynamic performance but also for how they affect the rider's ability to perform for extended periods. This kind of extended testing is not feasible in a wind tunnel because of the strain on the athlete and the extreme costs involved.

7. CONCLUSIONS

This paper has presented details of how aerodynamic design is pursued at Cervélo. The approaches described have been refined through more than 20 years of experience as the industry leader in aerodynamic bicycle design. We use all of them on a regular basis.

Our success rests on a triad of methods: computational fluid dynamics (CFD) simulations, wind tunnel testing, and real-world validation. Computational studies and wind tunnel testing have complementary strengths, with each providing information and advantages impossible to obtain with the other. Both have a role at every design stage. Both of them answer, however, to the ultimate verification: the performance of real athletes on a real track.

Taken together, the discussions in this paper reveal several themes that run through all our work.

- At every design stage, we seek to make choices that are informed by good science and hard data, for example by choosing appropriate mathematical methods and investing in detailed wind tunnel testing throughout design iteration.
- We constantly seek to optimize every technique in our repertoire, for example by refining simulation code, using automated parameter optimization, or developing a database of error bands to support go/no go decisions at the wind tunnel.
- Likewise, we constantly seek to improve the consistency of our methods and thus the repeatability of our results, as evidenced by the many ways we check and re-check bike setups before testing.
- We are committed to fairness and accuracy in how we study and report about bikes; this paper is a demonstration of that commitment.

Between the launch of the 1994 Barrachi and our newest offering, the 2019 S5, the industry has evolved spectacularly. Through all the changes, one idea is a constant at Cervélo: There is always something new to be learned—and then translated into speed.

CONTRIBUTORS

Like our design work, this document was a collaboration among many voices, each of whom brings a unique perspective to Cervélo's technical practice.

Howard Buckley specializes in CFD, aerodynamics, and wind tunnel testing of our bikes. He works closely with our entire design team to improve the aerodynamics of all of our products.

David Killing leads our design team in developing new products and draws on all of the skill sets available within Cervélo to do this.

Richard Matthews is responsible for coordinating Cervélo's R&D efforts, as well as maintaining all of our design tools and mentoring the design team.

Robert Pike oversees all of our wind tunnel testing and detailed design of mechanical parts for our products.

Graham Shrive directs the engineering team and interfaces with our professional race teams to ensure they are included in the product development cycle.

V1 2018-09-26

www.cervelo.com