SOLUTION BRIEF

Manufacturing
Modeling & Simulation

Winning Races through Faster Innovation

Predictive Simulation Helps Roush Yates Engines Design Better Engines in Less Time

There is always some degree of uncertainty when moving from engineering simulations to physical prototypes. The real world has more variables and finer granularity than our digital models. Yet recent advances in hardware and software have significantly reduced the gap, enabling engineers to build and run digital simulations faster and extract deeper and more reliable insights into the performance of complex physical systems.

There are few industries where fast, predictive capability is more critical to success than professional motorsports. Race cars are complex, tracks are all different, and the racing season is long and arduous. Yet the difference between winning and losing a 500-mile race can be a matter of inches. Even minor design optimizations can make the difference between defeat and victory.

Roush Yates Engines has risen fast in this uber-competitive world. Since 2004, cars sporting Roush Yates Engines* have won more than 300 races in some of the world's toughest competitions, averaging more than 20 wins per year across NASCAR, IMSA, and FIA Series. Speed on the track is only part of the story. Roush Yates Engines designs and builds more than 1,000 race engines per year and competes in more than 128 events. The company's engineering team must deliver world-class innovation at high volume and with fast turnaround times. Efficient digital simulations that accurately predict real-world performance are key to their success.

Deep Insight into Engine Dynamics

Roush Yates Engines uses digital simulations throughout its design cycle. According to Todd English, vice president of business development, "We scrutinize and analyze every component and optimize every engine for a specific race on a specific track. Since we can evaluate most components in a virtual environment quicker than we can manufacture and test them, simulation is one of the most valuable tools we have in this modern age of competitive motorsports."

Simulations take time and consume resources, so careful integration into design flows is important. The design team at Roush Yates Engines develops and evaluates digital models for key components and subsystems separately. These parts are then combined into full engine models to examine how they perform together, including how they perform in the vehicle and on the track under realistic race conditions. As ideas solidify, the team transitions gradually toward larger and more complex simulations and higher resolutions to fine-tune their designs.

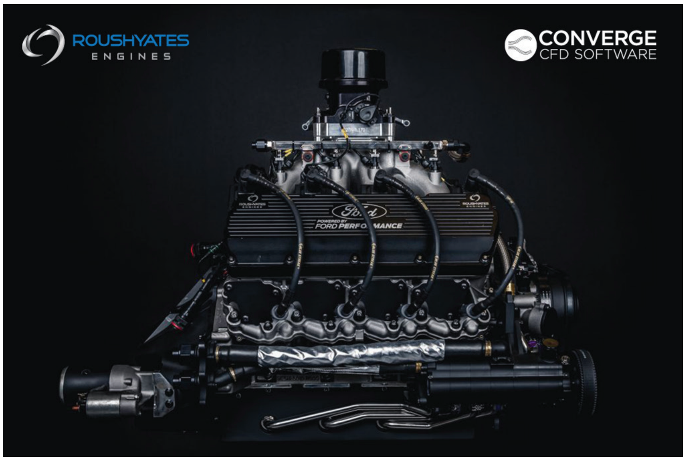

Figure 1. Roush Yates Engines designs and builds more than 1,000 racing engines per year. Fast, accurate engineering simulations are essential to enable world-class innovation at such high volumes.

Software Advances for Predictive Accuracy

Digital models illuminate engine dynamics in ways that are difficult or virtually impossible to explore on a firing engine. Yet the behavior of the models must correlate closely with physical reality to enable effective design decisions. The team at Roush Yates Engines uses CONVERGE* simulation software, which provides a range of modeling options for internal combustion engines. A number of features in the application help to improve predictive accuracy.

- **Automatic mesh generation** provides an optimized numerical model based on design characteristics and simulation requirements. This eliminates the time, effort, and errors associated with manual approaches.
- **Dynamic remeshing** allows the number and size of the cells in the mesh to be optimized automatically at each time step of a simulation. Resolution and accuracy can be increased in the areas of greatest interest, without driving up overall simulation runtimes.
- **Fully coupled chemistry** allows for greater accuracy and efficiency when analyzing how physical and chemical interactions affect each other and impact power delivery.

Together, these capabilities help the engineering team move from design, to simulation, to results more quickly. They also help them shine a brighter light onto the most critical phenomena, such as fuel flow, aeration, dispersion, and ignition.

Hardware Advances for Speed and Efficiency

Dynamic remeshing and fully coupled chemistry add to simulation workloads, potentially extending runtimes beyond acceptable limits. Says English, “When developing a component, we might have 15 designs that we want to evaluate. If each design case takes 12 hours to solve, it could take more than a week to run them all.”

To support increasing workloads without longer runtimes, Roush Yates Engines is introducing more parallelism into its computing platforms using Intel® Xeon® Scalable processors (formerly code-named Skylake) and Intel® Xeon Phi™ processors. These processors offer more cores and threads than their predecessors. Several additional features help to optimize parallel throughput, including support for extra-wide vectors (up to 512 bits) and high-speed on-die memory. These processors also support Intel® Omni-Path Architecture (Intel® OPA) high-performance fabric. Intel® Xeon® processors and Intel OPA are both part of the Intel® Scalable System Framework, which helps to simplify the design of balanced, high-performing clusters.

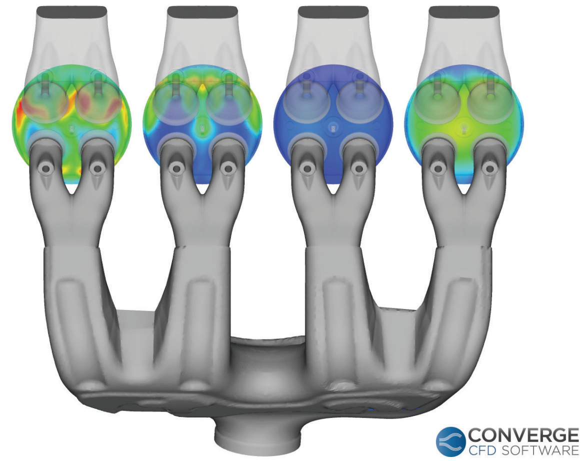

Figure 2. Recent advances in simulation software are helping engineering teams move from design, to simulation, to results more quickly, while shining a brighter light on the most critical phenomena.

Firing on All Cylinders

Although increasing compute parallelism can be an effective strategy for many high-performance computing applications, it is not appropriate for all workloads. Importantly, the chemistry and flow solvers in CONVERGE parallelize independently, so they are well suited to a highly parallel execution environment. Convergent Science worked with Intel to optimize their code for the new platform.

Recent CONVERGE benchmarks show that the new Intel Xeon Scalable processor architecture can provide up to 1.3X higher performance than the previous-generation Intel® Xeon® processor v4 family using the same processor speed.[1] The performance boost offers significant time-to-solution advantages for developing winning engine designs. Optimizations for Intel Xeon Scalable processors also tend to improve performance for other Intel® processors, so the benefits extend across multiple compute platforms. For Roush Yates Engines, this means that any simulation can be run efficiently on the fastest available computing resource.

Taking It to the Track

Roush Yates Engines’ work doesn’t end after an engine is built. The company provides track-side support to help customers get the best performance in the only place where it really matters. Drivers and pit crews make split-second adjustments throughout a race to optimize traction, fuel efficiency, power delivery, and more. Any one of these decisions can mean the difference between success and failure. According to English, “We benefit from these engagements as much as our customers do. We get to monitor engine operations before, after, and in some cases during each race. We use that data to compare real-world performance with our modeling, so we can continue to improve the predictive value of our simulations.”

Accelerating into the Next Era of Reliable Speed

Roush Yates Engines isn't resting on its many laurels. Every engine design must be better than the last, which means taking advantage of technology advances at every level. From materials and physical tools, to simulation software and computing platforms, new innovations must be integrated into the design flow early enough to deliver competitive advantage yet late enough to avoid the costs and risks that are often absorbed by the earliest adopters.

It's not an easy balance to maintain. However, as the industry continues the march toward fully predictive simulations, Todd English believes it's imperative to stay near the front of the pack. "Roush Yates Engines builds world-class engines, which requires world-class computing and design tools. We work with the best to deliver the best. And we have no intention of slowing down."

[1] Testing conducted by Intel using CONVERGE* 2.4.8 (custom build for Intel® Advanced Vector Extensions 2 and Intel® Advanced Vector Extensions 512) and Flame-D* workload. Baseline server configuration: 2 x Intel® Xeon® processor E5-2697 v4 (2.3 GHz, 18-core), 128 GB memory (8 x 16 GB DDR4 @ 2400 MT/s), Red Hat Enterprise Linux* 7.2, 2 x Intel® SSD S3500 Series. New server configuration: 2 x Intel® Xeon® Gold 6148 processor (2.4 GHz, 20-core), 192 GB memory (8 x 16 GB DDR4 @ 2400 MHz), Red Hat Enterprise Linux 7.3, 800 GB Intel® SSD Data Center S3500 Series.