# Design and Simulate in Parallel

PERFORM SIMULTANEOUS CAE OR STRUCTURAL/FLUID ANALYSIS ON THE SAME SYSTEM BEING USED FOR DESIGN.

*Produced by the editors of* Desktop Engineering

Desktop Engineering®

*Sponsored by:*

NVIDIA

hp

# CONTENTS



# APPENDIX

# EXECUTIVE SUMMARY

Best-in-class engineering departments are replacing traditional workflows with more efficient practices that streamline the product development process. These new workflows are based on replacing a heavy reliance on physical testing with more reliance on simulation.

However, a modern workflow intended to speed product development by relying on simulation instead of extensive physical testing is often crippled by improperly equipped workstations that struggle under the load of complex simulation jobs. Idled engineers resort to scheduling simulations after hours, hoping they're set up correctly so they don't fail in the middle of the night. Others turn to scheduling time on data center resources. Some even have to send their designs to outside service providers to have the simulations processed.

All of these coping mechanisms add time to the design process. Less time equals fewer and less-complex simulations that might miss a critical flaw. As engineers struggle with inefficient workflows, new design ideas may not even be pursued.

To take advantage of the increased use of simulation, engineers need the right equipment and workflow that allows them to design and simulate simultaneously. This white paper explains how to unlock the potential of such a workflow while achieving more than 3x faster simulation runs.

The benefits of such a simultaneous approach to design engineering are illustrated in this paper by three companies' real-world experiences: Astrobotic, which is competing for the $30 million Google Lunar X prize to land a robot on the moon; Liquid Robotics, an unmanned maritime vehicle developer; and Briggs & Stratton, a manufacturer of engines and outdoor power equipment. Each explains how they are able to work with higher fidelity simulations, review more design candidates in less time, and reduce simulation processing times.

# 1 | DESIGN OPTIMIZATION CHALLENGES

Design engineers face incredible pressure to create higher quality products faster and more affordably. A competitive global market, uncertain economy and disruptive technologies have compounded product development challenges.

Many engineers have turned to computer-aided engineering (CAE) software to keep up. Unfortunately, they often find CAE software slows their workstations to a crawl, making them useless for any other task while a simulation is running. For engineers using workstations that aren't properly equipped for advanced CAE, simulating a single event could take hours, days, or even weeks, depending on the complexity of the model. Repeating this process for multiple design candidates could add weeks or months to the development cycle.

## Settling for a Separated Design and Simulation Process

To cope with a design process crippled by overloaded workstations, many companies are separating simulation from the design process, primarily by resorting to batch simulation runs, or investing in data center resources and outside service providers.

Sitting idly by during the busy workday while a workstation chugs away at a simulation for hours on end prompts many engineers to run simulations in batches overnight or on weekends. If the simulations are set up correctly, a batch simulation process will yield results the next workday. Too often, however, the simulations fail to complete because of a missed variable. Already a compromise to maximum efficiency, waiting to run a simulation at night only to have it fail wastes even more time.

The lost time and opportunities are enough for many companies to consider investing in data center resources dedicated to simulation. Using such resources means simulations can be run during work hours without tying up workstations. If a simulation fails, at least it can be caught and re-run during the course of the workday. Unfortunately, data centers are not in the budget for many small- and medium-sized businesses (SMBs). Even those companies who are able to make the investment in hardware and maintenance may find that this approach often saves less time than anticipated. For many engineers, the popularity of offloading simulations to dedicated high-performance computing (HPC) resources leads to scheduling their jobs and then waiting for their turn in a long queue.

> **Already a compromise to maximum efficiency, waiting to run a batch simulation at night only to have it fail because of a setup error wastes even more time.**

Faced with the imperfect solutions of simulation via batch and dedicated simulation processing, some engineers reach out for help. They'll outsource their simulation work to contractors who specialize in it and have the necessary computing resources. For example, in the case studies that appear later in this white paper, one company explains how it was forced to rely on its simulation software vendor to run its simulations for them. Collecting and sending files, then waiting for simulation results to come doesn't maximize efficiency, but many SMBs simply can't afford to lose the use of their workstations to run a simulation, or make the upfront investment in data center resources.

## The High Price of Concessions

These coping mechanisms—batch runs of simulations, and offloading simulation work to data center resources or a contractor—are better than a sequential design process that relies on expensive and time-consuming physical testing instead of simulation. However, they still result in a disjointed design approach that prevents companies from fully realizing the high productivity and design optimization promise inherent in a simultaneous design approach.

**The heavy cost of simulation runs, measured in hours and days, imposes a frugal attitude that limits innovation: Explore as few ideas as possible.**

Engineers too often feel forced to make painful cuts in model fidelity and number of variations they evaluate in order to meet development milestones. Thus, the heavy cost of simulation runs, measured in hours and days, imposes a frugal attitude that limits innovation: Explore as few ideas as possible.

There is a better way. To break away from sequential design and truly create an efficient design cycle, engineers need to be able to design and simulate simultaneously on the same machine.

## 2 | SPEED PRODUCT DEVELOPMENT WITH PARALLEL PROCESSING

Simultaneously designing and simulating on the same machine allows engineers to use their time more efficiently. Their increased productivity lets them evaluate a greater number of designs in the same amount of time, ultimately leading to the development of an optimal product.

When a workstation's resources remain available while the CAE software is engaged, it allows engineers to refine designs while they verify them in the background. This gives engineers more control over a simulation job. They can periodically check in with the job in progress, see how far along it is, and, if necessary, interrupt it and start anew as soon as failures are identified.

Access to the CAE environment on their workstation also encourages more exploration and experimentation, leading to innovative thinking. Without the need to submit jobs to an HPC queue, without the system sluggishness usually associated with running simulations on a standard workstation, designers and engineers have greater freedom to rigorously evaluate new ideas, increasing the odds of creating breakthrough innovations.

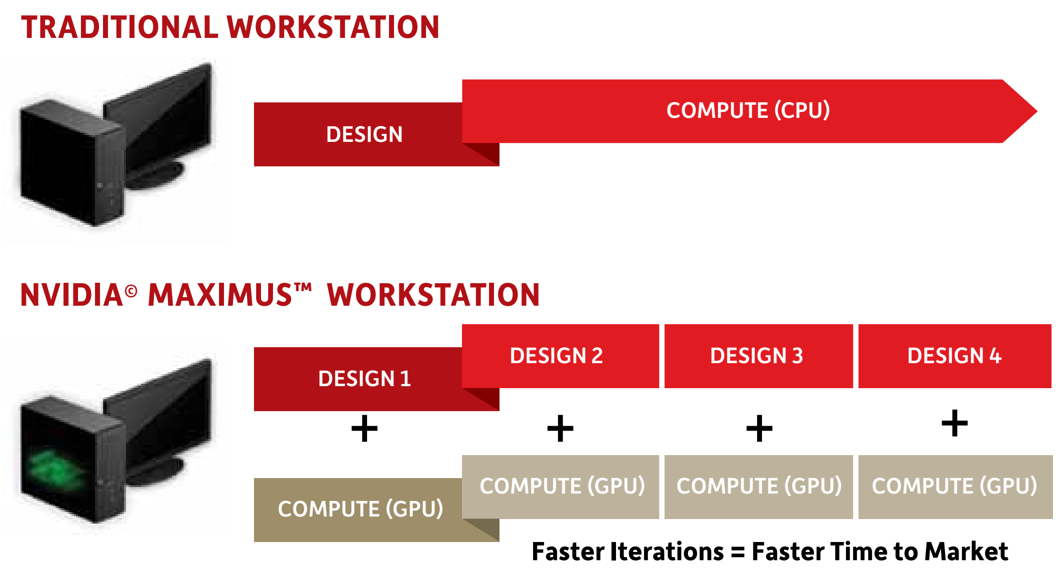


**Designers and engineers can now perform compute-intensive tasks like CAE, rendering, or structural/fluid analysis on the same system they are simultaneously using for design work, without having an impact on their applications' speed and interactivity.**

*Source: NVIDIA*

To carry out simultaneous design and simulation on a single workstation, both the hardware and software must be designed to support parallel processing, which involves splitting operations into parts that execute simultaneously on different processors in the same computer.

The hardware must provide sufficient processing power to satisfy the needs of both the CAD and CAE programs—*not one at a time, but both at the same time*. To facilitate a simultaneous workflow, a multicore system must be properly configured to produce the best accelerated computing, visualization, and simulation.

Hardware is only half the equation. To unlock the full benefit of simultaneous design requires using design and simulation software that is written to take advantage of parallel processing. Most major CAE software vendors, including ANSYS, MathWorks, MSC and Dassault Systèmes, have products that take advantage of multiple cores from CPUs and GPUs.

When the hardware cannot support parallel processing to the degree required by an intensive CAE task, the simultaneous workflow model falls apart. As Jason Calaiaro, the chief information officer for Astrobotic, the subject of one of the case studies in this white paper, notes: "When the CPU cores are fully occupied with analysis jobs, you cannot do anything else—not even something as trivial as browsing the web or writing an email. Essentially, the workstation has no available resources, so it becomes a dead node."

NVIDIA and HP recognize that Astrobotic and other businesses like it need to carry out simultaneous design and simulation on a single workstation without sacrificing design speed or interactivity. To respond to the need for simultaneous design, HP has created workstations with NVIDIA Maximus™ technology. Maximus turbocharges parallel processing, allowing workstations to run design and simulation software simultaneously.

NVIDIA's Maximus technology enables a simultaneous design and simulation workflow by distributing tasks to two GPUs: an NVIDIA Quadro® graphics card and an NVIDIA Tesla® companion processor, which in turn frees up CPU resources for the work they are best suited to do. The Quadro family is specifically developed for graphics-intensive applications like CAD, while Tesla is designed for the parallel-processing jobs commonly required in digital testing and simulation. Maximus transparently

## CUDA Compatibility

Because Maximus-equipped workstations rely on the additional cores in NVIDIA GPUs to take on the lion's share of the workload, they are suited for CAE programs compatible with NVIDIA's Compute Unified Device Architecture (CUDA) GPU computing specification. Recognized CAE software developers—ANSYS and Dassault Systèmes SIMULIA, to name but two—have refined their software suites so they can fully leverage the additional power available in Maximus-equipped workstations.

**NVIDIA's research shows that Maximus can more than triple the speed of SIMULIA Abaqus jobs over a base two-core workstation license.**

and automatically assigns visualization and simulation or rendering work to the right processor, enabling engineers to simulate on the Tesla GPU without affecting their ability to design using the Quadro GPU.

For example, NVIDIA's research shows that Maximus can more than triple the speed of a typical SIMULIA Abaqus job over a base two-core workstation license, all while engineers are still working in other applications, including CAD, with no loss in interactivity. See the table below for more details.

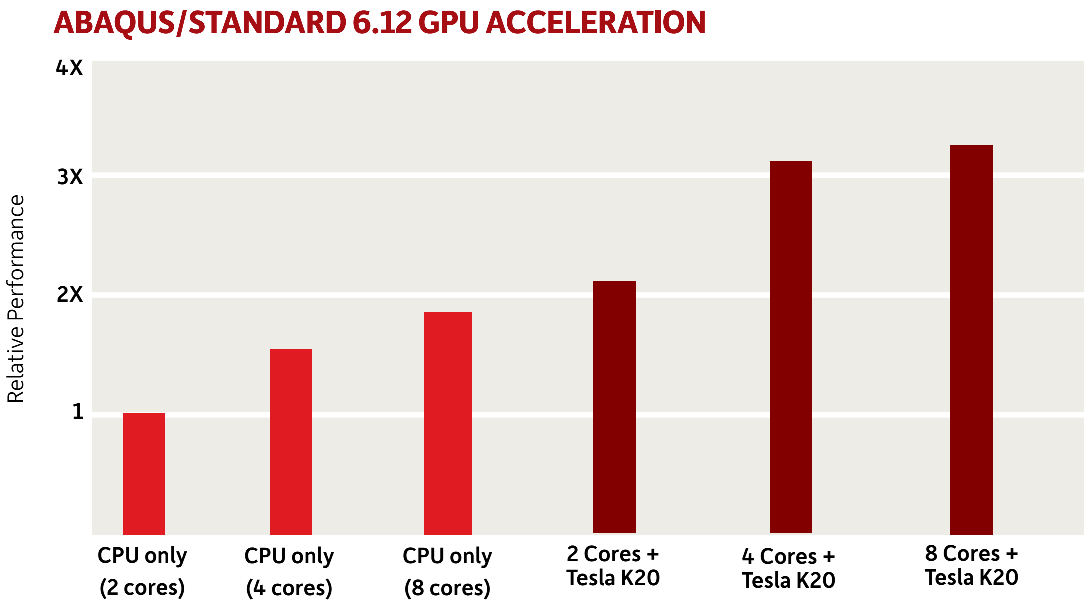


**The effect of adding an NVIDIA Tesla K20 to a 2-, 4- and 8-core workstation shows a marked improvement in Abaqus model processing performance.**

*Benchmarks performed using S4B model, engine geometry, 3.2M degrees of freedom, static nonlinear, direct sparse. System: HP Z820 workstation with Sandy Bridge CPUs (2x E5-2643, 3.3GHz, 48GB memory), Tesla K20, and Windows 7. Data provided by NVIDIA.*

## 3 | CASE STUDIES

On paper, workstations equipped with Maximus graphics technology look great. The multi-core Quadro and Tesla GPUs enhance the CPU by taking on some of the graphics and CAE processing, enabling engineers to design and simulate on the same workstation at the same time. But, just as real-world performance is the ultimate test of a design, the proof of Maximus' worth is how it performs under real-world conditions.

Among adopters of NVIDIA's Maximus technology are Astrobotic Technology, Inc., a space robotic engineering company; Liquid Robotics, an unmanned maritime vehicle developer; and Briggs & Stratton, a household name in engines and outdoor power equipment. The two robotics companies are well suited to take advantage of NVIDIA Maximus technology because they were bogged down by a slow, sequential workflow. They relied on CAD and CAE to develop, evaluate, and prove their concepts, and they were ready to find an alternative to reduce their design cycle. Briggs & Stratton took a different approach by dedicating its Maximus-equipped workstation to running simulation jobs to free up its engineers' workstations for design work.

### CASE STUDY: Astrobotic

**GOAL:** Land and control a robot on the surface of the moon.

**CHALLENGE:** Simulations needed to analyze aspects of the launch and landing sequence occupied a workstation for up to 10 hours, and sometimes failed.

**SOFTWARE:** SolidWorks, MathWorks MATLAB, ANSYS

**SOLUTION:** An HP Z800 workstation, equipped with NVIDIA Maximus' dual-GPU architecture, allowed higher-fidelity simulations to be run in less time, without taking the workstation away from simultaneous design work.

### ASTROBOTIC

Astrobotic, a spinoff of Carnegie Mellon University's Robotics Institute, is one of 26 engineering teams currently competing for the $30 million Google Lunar X prize, set aside for the first privately funded group to safely land a robot on the surface of the Moon, have that robot travel 500 meters over the lunar surface, and send video, images and data back to the Earth. The teams have until the end of 2015 to get to the Moon, meet the prize objectives, and win the prize purses. The Pittsburgh-based company envisions driving a new era of space exploration, science, tourism, resource utilization, and mining with its technology, which includes a lander and an attached rover.

The company's plans literally aim for the moon, but it's not a large business. Astrobotic consists of about 20 people who are either on its staff or collaborate from Carnegie Mellon.

## High-Fidelity Simulation No Longer a Bottleneck

While remotely controlling a robot on the moon as it takes videos and snaps pictures is no small feat, the real test for Astrobotic's engineers is to stick the landing. The company can physically test its lander's mechanics and remote controls, but physically testing the landing sequence is not practical for the small company.

"That's why we invested a lot to develop a comprehensive, sophisticated simulation environment that allows us to recreate an entire launch-to-landing simulation inside a computer," said Jason Calaiaro, Astrobotic's chief information officer. While Calaiaro didn't divulge specific investment costs, the company's software environment consists of SolidWorks for mechanical design, and MathWorks' MATLAB and ANSYS for digital simulation.

The CAD and CAE software was only part of the solution. Calaiaro admitted (rather sheepishly) that Astrobotic's pre-Maximus hardware consisted of three-year-old workstations. Running an ANSYS analysis job on an older machine, he revealed, could take up to 10 hours because of the level of precision needed.

At Astrobotic, the complexity of the simulation is measured in degrees of freedom (representing the possible directions of deformation for each tiny geometric mesh that must be calculated). Knowing the limitation of its hardware and length of time each analysis exercise took, engineers were often forced to simplify the jobs to around half a million degrees of freedom just to get them to run faster.

### KEY FEATURES: Astrobotic's Griffin Lander

- Operating Environment: Cruise, orbit, and surface operations at any lunar destination.
- Lander Mass: 525 kg.
- Rover and Payload Mass Capability: 260 kg.
- Average Power: 200W (600W peak)
- Landing Precision: 100 m of target coordinates
- Dimensions (LxWxH): 3 x 3 x 2 m
- Features: Image Registration for Precision Landing, Laser Scanner for Hazard Detection and Safe Landing, Bipropellant Thrusters for Orbit Capture, Attitude Control, and Landing, Passive Thermal Control, Direct-to-Earth Communication

*— Source: Astrobotic*

"There's a lot of special analysis that needs to be done … A [lunar lander] is most likely to fail at the last minute before it lands, just because there is so much that has to happen autonomously," Calaiaro said. "The robot needs to understand its altitude, what orientation it's in, where it's going to land, identify obstacles correctly, select a clear landing site, and guide itself down. These all happen in such a quick succession that you cannot intervene."

Simplification of the simulation means less accuracy, which is critical to the landing sequence. But even the trade-off in fidelity to speed up simulations wasn't enough to keep Astrobotic engineers working while the CAE software crunched numbers.

"When the CPU cores are fully occupied [with analysis jobs], you cannot do anything else—not even something as trivial as browsing the web or writing an email," he said. "Essentially, the workstation had no available resources, so it became a dead node."

As a workaround, Astrobotic coordinated overnight analysis jobs into its development cycle. Returning to work the next morning to find an analysis job had crashed or failed because of errors in setup was, Calaiaro said, "more common than anyone would like to admit."

To meet its deadlines, the company asked for help. ANSYS, a sponsor of the robotics company, ran many of Astrobotic's calculations for them.

While he appreciated the help, it was obvious to Calaiaro that Astrobotic needed to update its design resources.

### Refine and Test Models More Completely and in Less Time

To help rectify its sluggish sequential design process, the company acquired an HP Z800 workstation, equipped with multicore Intel Xeon CPUs with 12 cores and NVIDIA Maximus' dual-GPU architecture. The parallel processing capabilities of Maximus had an immediate impact on Astrobotic's simulation quality and speed.

"Maximus shines when you get above about 1.5 million degrees of freedom," said Calaiaro. "Now we can do complete analyses on our lander that runs 2 million to 3 million degrees of freedom, which means we can refine and test our models more completely and in less time."

The Maximus-equipped HP Z800 allows Astrobotic's design engineers to create more iterations of higher quality models in less time. They can design and simulate simultaneously, and are able to check in and recalculate simulations on the fly if something doesn't look right.

> **"The NVIDIA Maximus-powered system is like getting three people's worth of use on a single machine. This system is a beast."**
>
> *—Jason Calaiaro, Astrobotic's Director of Information Systems*

“The NVIDIA Maximus-powered system is like getting three people’s worth of use on a single machine,” said Calaiaro. “This system is a beast. We haven’t yet found anything it can’t handle—even simultaneous CAD, analysis, and additional number crunching in remote rendering jobs.”

## Additional Investments Planned

After experiencing the simultaneous workflow made possible by the HP workstation equipped with NVIDIA Maximus, Astrobotic has decided to invest in an additional Tesla GPU, bringing the total number of GPUs in the system to three.

With the additional Tesla unit’s parallel processing power, Calaiaro hopes to run the complete lunar mission on the single workstation. Such a simulation requires so much fidelity in physics computation that assigning it to one workstation would have been unthinkable in the past.

“It’s unprecedented for what we’ll be able to do on a desktop,” Calaiaro observed. “Without Maximus technology, we could not have developed the landing software used for the lunar lander robot.”

## LIQUID ROBOTICS

Liquid Robotics ' history can be traced all the way back to the day its original founder, a venture capitalist named Joe Rizzi, got the idea to digitally capture and broadcast the songs of humpback whales through the Internet. Work on his pet project led him to invest in developing commercial products involving ways to autonomously collect data on ocean currents, temperatures and waves. By 2010, Liquid Robotics was a global ocean robot data services company.

PHOTO: BEGUERY, CNRS-DT INSU

### CASE STUDY: Liquid Robotics

**GOAL:** Create an unmanned marine robot that can autonomously navigate the ocean while collecting data.

**CHALLENGE:** A complex, one-minute simulation took five hours of processing, and consumed all of the engineer's workstation resources.

**SOFTWARE:** SolidWorks, MathWorks MATLAB, ANSYS

**SOLUTION:** HP Z800 workstations, equipped with NVIDIA Maximus' dual-GPU architecture allow multiple design candidates to be simulated—saving the cost and time of seaborne testing.

The company's flagship product is the Wave Glider, an unmanned marine robot that is powered by the ocean's endless supply of wave energy and solar power. That means no manpower, no emissions and no refueling as the surfboard-sized robot moves across the ocean's surface, dangling a payload of sensors beneath it.

### Slow Design Process Stifled Innovation

The ocean is a harsh, unpredictable environment that makes physical testing difficult. Tim Ong, vice president of mechanical engineering at Liquid Robotics, explained that field testing his firm's products involves the cost of a full-scale Wave Glider unit, plus the cost of deployment via a shipborne operation. Any catastrophic failure in the field could lead to loss of the vehicle itself.

There is also a huge investment in time to physically test a product that is designed to operate autonomously for an entire year. For example, physically testing just one critical component, a spring system that controls the angle of attack of the WaveGlider's wings, took at least five months to equal a year's worth of real-world wing adjustments.
At this rate, even if the design process were restricted to just three iterations, it would have taken the company a year-and-a-half just to validate one component of a new product—a serious impediment in design development. The inevitable conclusion, as summed up by Ong: "It's a whole lot faster to study the fatigue of the system and whether it will survive if we test it digitally."

Liquid Robotics began using ANSYS and MathWorks' MATLAB to simulate the vehicle's mechanisms that it designed with SolidWorks. The simulations were far and away faster than the physical tests, of course, but they slowed to a relative crawl on the company's workstations, which weren't properly configured for parallel CAE tasks.

"In certain simulations, it took us 300 seconds to compute every second of simulation," recalled Ong. "For a one-minute simulation, we had to put in five hours processing … for some time-consuming scenarios, you end up deciding whether to go down that road or not based on how long it would take to simulate it."

**"Since we had to wait till the end of the day to run (simulation), we basically had one shot. If it failed, everything you set up and ran the previous night was wasted."**

*—Tim Ong, Vice President of Mechanical Engineering, Liquid Robotics*

In addition to the time sink restricting engineers' ability to test what-if scenarios that might have led to product innovations, it also crippled their productivity. When their workstation's resources were completely consumed by running a simulation, the engineers were essentially left without a computer.

"If you wanted to do anything else while running a simulation or modeling, you were out of luck," said Ong.

Like Astrobotic, Liquid Robotics engineers began running simulations on evenings and weekends so they could still use their workstations during the workday. But they were also frustrated by failed simulation runs that required inputting dozens of numeric parameters just to set up.

"Structural problems and fluid problems [required in simulating the Wave Glider] have to be set up just right," said Ong. "Since we had to wait till the end of the day to run them, we basically had one shot. If it failed, everything you set up and ran the previous night was wasted."

### The Results of a Simultaneous Workflow

Parallel processing power has transformed Liquid Robotics' workflow. Instead of passing design work to the engineers whose computers weren't bogged down running a simulation, NVIDIA Maximus graphics technology allows engineers to simultaneously design and simulate.

The company's engineers now use HP Z800 workstations with two CPUs and Maximus to speed their simulations. Six of the CPU cores plus an NVIDIA Tesla companion processor run CAE software, leaving the workstation's other six CPU cores, along with its Quadro GPU, to run SolidWorks and other design and office programs.

In addition to increasing productivity by getting engineers back in front of their workstations, the simulation process is also streamlined by Maximus.

"In a lot of the analysis jobs, you're looking for convergence," said Ong. "If an engineer is using the same machine to run analysis while he's working on CAD design, he can just switch over to the analysis job, and say, 'No, that's not what I want. Stop. Let's try something else.'"

Maximus has made simulation runs faster and its engineers more productive. In a word, Maximus has given Liquid Robotics *time*.

### KEY FEATURES: Liquid Robotics' WaveGlider

- Float Dimensions: 208 cm x 60 cm (82 x 24 in.)
- Sub Dimensions: 40 cm x 191 cm (16 x 75 in.)
- Wing Dimensions: 107 cm wide (42 in.)
- Mass: 90 kg (200 lb.)
- Displacement: 150 kg (330 lb.)
- Water Speed: 0.4 to 2.0 kts.
- Base Payload: water speed sensor, AIS receiver
- Propulsion: Mechanical conversion of wave energy into forward propulsion
- Battery: 665 Watt-hours, lithium-ion rechargeable
- Solar Power: 80 Watts (peak) for battery charging, onboard electronics and payloads

*— Source: Liquid Robotics*

## BRIGGS & STRATTON CO.

While Astrobotic and Liquid Robotics are fairly young companies developing and perfecting new products, Briggs & Stratton has been manufacturing engines for more than 100 years. Its 3,000 employees build more than 10 million small engines that 80% of lawnmower brands use in their lawn tractors and push mowers, according to the company.

Despite the differences in size and scope, Briggs & Stratton had the same problem as Astrobotic and Liquid Robotics: lost productivity while workstation resources were consumed with simulation.

### Temperature Management Challenges

A critical design element of the company's outdoor power equipment—which includes snow blowers and portable generators as well as lawn mowers—is temperature management. Based on the anticipated operating conditions of the machines, they have to be designed to withstand a certain temperature range and heat load.

For example, a push lawn mower might typically be used for an hour at a time, but within that hour, the thermal load of the engine will gradually climb to a peak. A lawn tractor is typically used by consumers with more land, so it may operate longer on average. A power generator could run for an entire day. To ensure the different products function smoothly, Briggs & Stratton engineers need to simulate their heat load and mechanical operations, and study the heat's effects on the equipment's structure.

Briggs & Stratton uses PTC's Pro/ENGINEER Wildfire for equipment design. For studying the effects of the temperature on the product's mechanical components, engineers use Dassault Systèmes' SIMULIA Abaqus. And for analyzing the airflow, cooling, and combustion activities within the engine, they use ANSYS Fluent.

Dave Torres, a senior analysis engineer who determines hardware configurations for simulation in the company, said engineers learned to do without the use of their

**CASE STUDY:**
**Briggs & Stratton**

**GOAL:** Simulate complex heat loads and mechanical operations of engines in various environments.

**CHALLENGE:** Simulations, even on dedicated workstations, took too long to complete.

**SOFTWARE:** PTC Pro/ENGINEER Wildfire, Dassault Systèmes SIMULIA Abaqus, ANSYS Fluent

**SOLUTION:** Five HP Z820 workstations, equipped with NVIDIA Maximus' dual-GPU architecture, enabled a 60% reduction in processing time on some jobs, and up to a 30% reduction in simulation processing times on average.

CAD program when their workstations were tied up with SIMULIA Abaqus and Fluent operations.

"You could partition a machine so you can be running Pro/E while running simulation," said Torres, "but then you're slowing down the simulation solve time, so your wait time will be longer."

> **"In some cases, it leads us to explore different ideas we may not have had in the past because we didn't have time. We might give another idea a chance, because we now have the time to do it."**
>
> *— Dave Torres, Senior Analysis Engineer, Briggs & Stratton Co.*

Briggs & Stratton's approach was to divide and conquer. It invested in eight single-processor workstations and dedicated them to CAD, and eight dual-processor workstations it dedicated to simulation. Thus, the design work was separated from the resource-hogging demands of structural, thermal, and fluid analyses.

Even with this divided setup, Torres sought a way to further decrease the time it took to run complex simulations without over-simplifying them. He first considered creating an HPC cluster that would pool computing resources for simulation runs. Even though Briggs & Stratton is a large, established company, Torres did his homework and investigated whether it had the necessary infrastructure to support such an environment.

He is wary of the level of IT support that would be needed to install and maintain an HPC cluster. "We're still looking into that as a possible investment," Torres says. "But the appeal of Maximus is, it gives us many of the advantages of an HPC cluster without having to set up a cluster."

### KEY FEATURES: Briggs & Stratton's Professional Series (V-Twin) RIding Mower Engine

- Advanced Debris Management System: Proactively deflects dirt and debris from the engine; any remaining debris is trapped in the high-capacity, dual-sealing air cleaner cartridge
- Dura-Bore Cast Iron Sleeve: Withstands wear and abuse while providing improved oil control
- Dynamically Balanced Crankshaft (V-Twins): Minimizes engine vibration, resulting in smoother running and longer engine life
- Full Pressure Lube System: Automotive-type oil system for cooler operation and longer life
- Overhead Valve (OHV): Design enables the engine to run cooler and cleaner, while delivering more power, longer engine life and improved fuel economy

*— Source: Briggs & Stratton*

### Simulation Processing Times Drastically Reduced

Torres initially had some reservations about the performance gain promised by Maximus technology, as he assumed it was developed for rendering and visualization. Torres wasn't sure how the company's CAE software would perform on the new hardware. "The wait time was cut down from about 24 hours to 14-15 hours" on SIMULIA Abaqus and ANSYS Fluent jobs, Torres said. On average, Briggs & Stratton began to see up to a 30% reduction in simulation processing time.

"It also turned out, [studying the effects of heat on structures] in Abaqus was the type of simulation best suited for the GPU, so it was a good fit," said Torres. He noted that the performance gain varies depending on the nature of the job—not just whether the software supports parallel processing, but whether the job itself is well suited for parallel processing.

The results were more than enough to convince Torres that his group needed more Maximus-equipped workstations. Briggs & Stratton acquired five HP Z820 workstations with Maximus to replace some of the workstations dedicated to simulation.

According to Torres, with Maximus-equipped workstations, Briggs & Stratton engineers are less concerned with model complexity and size. They do not feel pressured to simplify the simulation problem by reducing the mesh count to cut down on the time it takes the simulation to run. Instead, they run complex models as they are, which gives them a more accurate analysis of the products' behavior.

## 4 | HARDWARE AND SOFTWARE CONSIDERATIONS

To determine whether to invest in Maximus technology, you must first decide if it's right for you or your organization. It's not a matter of size. You may be part of a small company trying to do something that's never been done before, like Astrobotic or Liquid Robotics, or part of a large, established company like Briggs & Stratton that is trying to do something it's been doing for 100 years better than ever.

It's the outcome that matters. A slow product development cycle—whether it's caused by inefficiencies related to excessive physical testing, or waiting for simulations to run through an HPC server queue or workstation—determines the quality and time to market of many companies' end products.

If you need to evaluate more design candidates via high-fidelity simulations to design optimal products, but are frustrated by inadequate computing resources that limit your engineering productivity and creativity, consider Maximus technology.

### Upgrade or Replace the Workstation?

If your workstations are relatively new, you may be able to upgrade them with Maximus technology. A Maximus-enabled platform is a combination of an NVIDIA Quadro card, Tesla card, and software driver technology. To host this combination of GPUs, you'll need a capable workstation chassis from a qualified original equipment manufacturer that has the necessary physical space, connections, power, and cooling capabilities.

#### KEY FEATURES: NVIDIA Tesla K20

- SMX streaming multiprocessor technology for up to a 3x performance per watt advantage compared to NVIDIA Fermi GPU architecture
- Dynamic Parallelism and Hyper-Q GPU technologies for simplified parallel programming and faster performance

#### KEY FEATURES: NVIDIA Quadro K5000 GPU

- Bindless Textures that give users the ability to reference more than 1 million textures directly in memory while reducing CPU overhead
- FXAA/TXAA film-style anti-aliasing technologies for image quality
- An all-new display engine capable of driving up to four displays simultaneously with a single K5000
- Display Port 1.2 support for resolutions up to 3840x2160 @60Hz

*— Source: NVIDIA*

Current HP workstations capable of supporting Maximus configurations include the Z420, Z620 and Z820. Older HP workstations that may be upgradeable include the Z400 and Z800, if they have the larger power supply configured. It is important to note that Maximus configurations need to use Tesla and Quadro pairings that use the same architecture. For example, a Fermi-based Quadro card cannot be paired with a Kepler-based Tesla card. See the table on page 24 for allowed pairings.

For many people with underperforming workstations, acquiring a workstation already configured with Maximus technology is a better option. In addition to all the benefits normally associated with the latest workstation technology, it ensures your workstation is properly configured for maximum parallel processing productivity.

"A new HP Z-workstation would be recommended for any Kepler-based Maximus solution," said Louis Gaiot, global commercial solutions at HP. "The new NVIDIA Kepler architecture GPUs are fully qualified on the HP Z420, Z620 and Z820 based systems, and appropriate Maximus configurations that cover a wide range of GPU computing capabilities have been carefully matched with the right CPUs and system memory to provide a great experience. In fact, HP offers the widest range of workstation-based Maximus solutions in the industry."

Your individual workload will determine the best workstation configuration. For example, a design-heavy workload with an occasional need for CAE simulations will require different capabilities than a simulation-heavy workload with a constant need for complex CAE simulations. To get more info on HP Workstations with Maximus, click on the link in the Appendix on page 24.

## RECOMMENDED CONFIGURATIONS

| | MODEL | PROCESSORS | MEMORY | POWER |
|---|---|---|---|---|
| **Occasional use of CAE** | HP Z420 convertible mini tower | NVIDIA Maximus: Quadro K2000 + Tesla K20 Single Xeon X86 CPU (4-6 cores) | 16-32 GB RAM | 600W 90% efficient, wide ranging, active Power Factor Correction |
| **Moderate use of CAE** | HP Z620 rackable tower | NVIDIA Maximus: Quadro K4000+ Tesla K20 Single Xeon X86 CPU (6-8 cores) | 48-64 GB RAM | 800W 90% efficient power supply |
| **Intensive use of CAE** | HP Z820 rackable tower | NVIDIA Maximus: Quadro K4000 + 2x Tesla K20s Dual Xeon X86 CPUs (8+ cores) | 64-128 GB RAM | 1125W 90% Efficient wide-ranging, active Power Factor Correction |

*Source HP and NVIDIA.*

### Software Considerations

Once you've determined the appropriate hardware, you'll need to consider the software required to enable the simultaneous workflow that's right for you. Not all CAE software versions are Maximus-enabled. At present, SIMULIA Abaqus 6.12-X or newer and ANSYS Mechanical 13.0 SP2 or newer are designed to make use of Maximus technology. MathWorks recommends the MATLAB R2011b and newer releases (together with the Parallel Computing Toolbox) to make the best use of Maximus.

If you're already running software that is capable of taking advantage of GPUs, but are frustrated by your hardware's speed, the licensing costs when upgrading to a Maximus-equipped workstation will be relatively minor. Different CAE vendors have different pricing models for accessing multiple cores that you will need to take into account.

**ABAQUS USERS:**
**Calculate Your Speed Boost**

If you're a current SIMULIA Abaqus user, you can see how GPUs would increase your performance at AccelerateAbaqusOnGPU.com. The online tool provides a performance gain breakdown based on your specific Abaqus model running on various GPU-enabled hardware scenarios.

If your current workflow relies heavily on physical testing and you are planning to roll out multiple new CAE seats in order to begin a simultaneous workflow, you must take into account the cost of each seat. Software is almost always the most significant cost driver of a new workflow. Training to use the new software and the investment in time away from production that training takes must also be added in. On the plus side, the return on investment (ROI) you'll see in terms of time savings and improved product quality will be even more dramatic with the new simulation-led workflow.

### Calculate ROI

Once you've added up the costs for hardware and software, you'll see it's dwarfed by how much a fast, simultaneous design workflow will pay off. How can you help your management team put that number into context? There are two key outcomes to value when estimating your ROI: productivity increases and innovation gains.

It's fairly easy to estimate the productivity gains that will be realized in an optimized workflow using NVIDIA Maximus graphics technology. It's a matter of comparing the costs vs. the savings over time that result from increased engineering productivity. Begin with the salary (including cost of benefits) of an engineer on your team to calculate the dollar figure associated with the percentage of time that engineer is actively engaged on his current workstation doing design and simulation tasks. This is the amount you are currently spending on design and simulation for that engineer. Now project the productivity gains of a Maximus-equipped workstation.

While each situation is different, having an HP Maximus-equipped workstation "is almost like having twice as many engineers," according to Astrobotic's Calaiaro. "The gains in productivity are substantial and tangible. Not only do our engineers get things done now more effectively, but they have a better time doing it."

Even if you take a very conservative estimate of productivity gains, the ROI case for Maximus is measured in months, just based on salary savings alone. But the real ROI comes from the ability to get better products to market faster. What is the value of the extra time a Maximus-based workflow provides?

For example: What might you have saved by finding a low-cost composite material that exceeded your quality standards for weight and strength? How many more errors could you have found by simulating hundreds of design candidates, rather than testing a few? How much could your warranty costs decrease by finding and correcting errors early in the design process? How much might your revenues grow because you identified a design feature among the increased number of iterations that your competitors don't have and your customers value?

**"The Maximus-powered system lets us take on more projects, which boosts our revenues and eventually will allow us to get to the moon faster."**

*—Jason Calaiaro, Astrobotic's Director of Information Systems*

As Liquid Robotics' Ong puts it: "Time is money. Getting the new product out is important. Potential sales that didn't happen because we couldn't test a new feature … that could be a big number."

Astrobotic's Calaiaro looks at it another way. "The Maximus-powered system lets us take on more projects, which boosts our revenues and eventually will allow us to get to the moon faster."

The value of discovery was not lost on Briggs & Stratton's Torres, either. "It allows us to get through more simulation iterations," he said. "In some cases, it leads us to explore different ideas we may not have in the past because we didn't have time. We might give another idea a chance, because we now have the time to do it."

And if that idea is the one that beats the competition, it's priceless.

## 5 | LEAD YOUR COMPETITION

The slow pace of a sequential design workflow—caused by its reliance on extensive physical testing or improperly equipped hardware that strains under the load of advanced simulations—is an albatross around the neck of many design engineering teams. It not only slows their time to market, it also limits the number of design candidates they can perform in the pursuit of the best design.

**Companies can choose to follow their competitors just to keep up, or they can choose to lead the competition into a simultaneous workflow, using it as competitive advantage to gain market share.**

Contrast that with best-in-class organizations that have invested in a robust parallel processing workflow that allows them to design and simulate simultaneously. A simultaneous design workflow doesn't just reduce their time to market, which is hard enough to compete against, it allows engineers to design better products in less time via multiple iterations. That's nearly impossible to compete against.

It's not a question of if engineering teams will move toward the most productive workflow, it's a matter of when. They can choose to follow their competitors just to keep up, or they can choose to lead the competition into a simultaneous workflow, using it as a competitive advantage to gain market share. Delaying the choice means falling further behind.

It's not difficult to correlate a more productive design workflow and a healthier bottom line. But adopting a simultaneous design workflow goes beyond the simple "time equals money" equation. When given to design engineers, time also equals product innovation. Time equals the next big thing that could transform your business.

The switch from sequential workflow to simultaneous workflow is the shift from a timid approach toward product design to a bold, creative one that yields productivity and quality improvements as a matter of course, but also provides the opportunity to disrupt the market with a truly innovative design.

# APPENDIX

## Resources

**NVIDIA Maximus Page:**
http://www.nvidia.com/object/maximus.html

**NVIDIA Maximus System-Builder's Guide:**
http://www.nvidia.com/content/quadro/maximus/maximus-system-builders-guide-win-7-64.pdf

**Test-drive your SIMULIA Abaqus/Standard jobs with GPUs for free:**
http://www.AccelerateAbaqusonGPU.com

**NVIDIA Maximus Technology for ANSYS Mechanical:**
http://www.nvidia.com/content/quadro/maximus/maximus-technology-for-ansys-mechanical.pdf

### VALID PAIRINGS FOR MAXIMUS

| KEPLER-BASED CONFIGURATIONS | | | | |
|---|---|---|---|---|
| NVIDIA Tesla | NVIDIA Quadro | HP Workstations | | |
| | | Z420 | Z620 | Z820 |
| K20 | K600 | Yes | Yes | Yes* |
| K20 | K2000 | Yes | Yes | Yes* |
| K20 | K4000 | ^ | Yes | Yes* |
| K20 | K5000 | – | ^ | Yes*** |
| K20 | K6000 | – | – | Yes*** |
| **FERMI-BASED CONFIGURATIONS** | | | | |
| NVIDIA Tesla | NVIDIA Quadro | HP Workstations | | |
| | | Z420 | Z620 | Z820 |
| C2075 | 600 | Yes | Yes | Yes** |
| C2075 | 2000 | Yes | Yes | Yes** |
| C2075 | 4000 | – | Yes | Yes*** |
| C2075 | 5000 | – | – | Yes*** |
| C2075 | 6000 | – | – | Yes*** |

^ Special configuration needed. Consult an HP representative for more information.
* Z820 accepts up to two Tesla K20 GPUs with Quadro K600, K2000 and K4000.
** Z820 accepts up to two C2075 GPUs with Quadro 600 and 2000.
***For multiple Tesla GPUs with Quadro 4000, K5000, 5000, K6000 and 6000, consult an HP representative for special configurations.

**Cards based on NVIDIA Fermi architecture on cannot be mixed with cards based on NVIDIA Kepler architecture.**

*Source HP and NVIDIA.*

## ACKNOWLEDGEMENTS

This white paper is based on interviews with companies using NVIDIA Maximus technology and HP workstations. It was prepared by *Desktop Engineering* as an educational resource for the engineering community. It was sponsored by HP and NVIDIA.