BENCHMARKING STUDY

Productivity Study: ANSYS Mechanical

Simulation run times are up to **4.8 times faster** when using a current workstation equipped with a powerful graphics processing unit, compared to an older system with an older GPU.

SPONSORED BY:

DELL NVIDIA

PRODUCED BY

DE Digital Engineering

Executive Summary

Simulation-led design offers many benefits to design engineering teams: Time and cost savings via less physical product testing, the ability to quickly determine the best initial designs to develop and the freedom to digitally experiment with innovative "what-if" scenarios are chief among them. In today's complex product design and development environment that might include smart products, autonomous vehicles, electric powertrains and digital twins, moving simulation further forward in the design process can give engineering teams a competitive advantage. However, hardware, software and cultural challenges have kept many companies from reaping the rewards of simulation-led design. Three recent developments explored in this paper promise to change that.

Affordable, Advanced Computing

As the availability of affordable computing power increases and simulation software is optimized to take advantage of those hardware improvements, more design engineering teams are able to realize the incredible benefits of a simulation-led workflow. We put this notion to the test by benchmarking modern ANSYS Mechanical 2019R1 software on a three-year old Dell workstation and comparing it to the same software running on its comparable present-day counterpart. We discovered that updating the computer enabled completion of the same simulations as much as 71% faster.

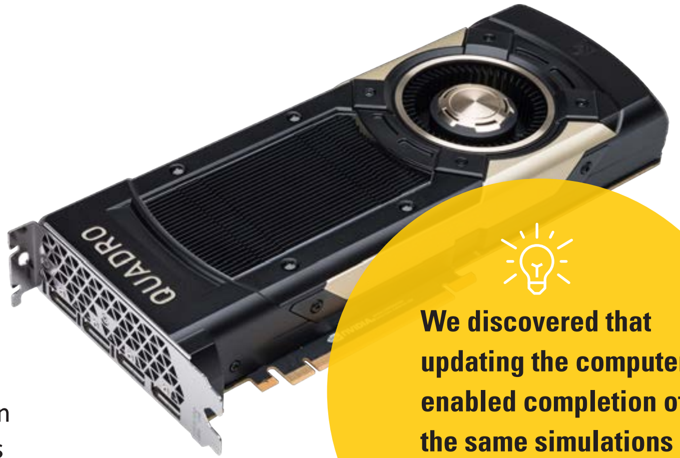

We discovered that updating the computer enabled completion of the same simulations as much as 71% faster.

Advances in GPU Technology and Adoption

When it comes to core count, GPUs have obvious benefits. The NVIDIA Quadro GV100 GPU in the modern workstation features 5,120 CUDA cores. Because ANSYS Mechanical is optimized to take advantage of distributing the workload across those cores on many types of simulation, the modern GPU further reduced the amount of time required to complete each analysis, in some cases by as much as 59%.

Democratization of Simulation

Simulation has long been the purview of professional analysts, but modern simulation platforms like ANSYS 2019R1 make simulation workflows more accessible to non-analysts. The company has steadily improved its user experience by rolling out a new user interface to speed user adoption and reduce learning times. The improved interface provides new capabilities around customization, ease-of-use features and instant search and launch tools. For design engineers accustomed to working in computer-aided design packages, the company recently introduced ANSYS Discovery Live, which also powers PTC Creo Simulation Live. The products provide instantaneous 3D simulation, tightly coupled with direct geometry modeling, to enable interactive design exploration and rapid product innovation.

This paper reports the detailed results of our benchmarking study and provides more information on the hardware, software and cultural changes that are making simulation-led design a reality for companies large and small. The study was sponsored by Dell and NVIDIA, and conducted by *Digital Engineering*'s editors.

The Productivity Study

To determine the speed boosts possible by updating an engineering workstation and GPU, *Digital Engineering*'s editors were provided with loans of workstations from Dell, GPUs from NVIDIA and software from ANSYS. The study environment consisted of a modern Dell Precision 7920 Tower and a three-year-old Dell Precision 7810 Tower workstation. They were configured with a current NVIDIA Quadro GV100 GPU for the modern workstation and an NVIDIA Quadro K6000 GPU for the older workstation. ANSYS Mechanical 2019R1 simulation software benchmarks were conducted for nine different models on each of the workstations.

The goals of the study were three-fold:

1. To compare the performance of the same set of simulations using current software on a modern workstation equipped with a modern GPU to what would have been a common engineering workstation configuration and GPU three years ago.

2. To show the effect of using a discrete GPU in a modern workstation vs. running simulations without a discrete GPU in that same workstation.

3. To show the effect of using a discrete GPU in a modern workstation vs. running simulations without a discrete GPU in an older workstation.

THE HARDWARE	OLDER WORKSTATION Dell Precision 7810 Tower	NEWER WORKSTATION Dell Precision 7920 Tower
		
CPU	Dual Intel Xeon E5-2620 V3,, 2.4GHz	Dual Intel Xeon Gold 6146, 3.2GHz
CPU Cores	6 each (12 total)	12 each (24 total)
RAM	64GB 2133MHz	192GB 2666MHz
Hard Disk System	1TB SATA HD	1TB M.2 PCIe NVMe, no RAID
GPU	NVIDIA Quadro K6000 12GB GDDR5 2,880 CUDA cores	NVIDIA Quadro GV100 32GB HBM2 5,120 CUDA cores 640 Tensor cores

As a baseline, we'll focus on using 4 CPU cores. We decided on 4 CPU cores because at the time of testing, ANSYS Mechanical licensing allows use of 4 cores per user in its base license. To use more cores, users can purchase HPC Packs. ANSYS defines "core" as either a single logical CPU core or a single GPU video card (regardless of the number of CUDA cores). CUDA stands for Compute Unified Device Architecture. It is a programming language that can leverage NVIDIA GPUs to perform tasks with greater performance. Each GPU can contain hundreds to thousands of CUDA cores, as the NVIDIA Quadro K6000 (2,880 CUDA cores) and NVIDIA Quadro GV100 (5,120 CUDA cores) used in our tests both do.

One key observation that is backed-up in this study is that users face diminishing returns in terms of the cost/performance ratio when increasing CPU cores – increasing the number of cores does not necessarily produce a linear increase in speed (see chart below). Under the ANSYS licensing model, one HPC license can unlock the entire GPU, regardless of the number of CUDA GPU cores. It may be more cost-effective to purchase the same number of CPU cores in combination with a large GPU, because users can experience increased performance with a lower licensing investment.

The NVIDIA Quadro GV100

The NVIDIA Quadro GV100 was designed to meet the demands of AI-enhanced design and visualization workflows for workstation users. In addition to working with higher fidelity CAE models, it is well-suited to rendering, features AI-accelerated denoising, and allows users to work with complex, photoreal datasets in virtual reality. Users can also accelerate AI training/inferencing thanks to the Tensor Cores in the Quadro GV100 and NVLink, which allow multiple GPUs to be used together. All Quadro cards are certified with a range of professional applications.

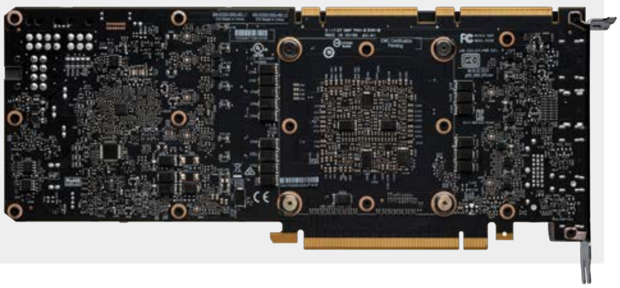

Effect of Increasing CPU Cores Compared to Normalized Solve Time (1T)

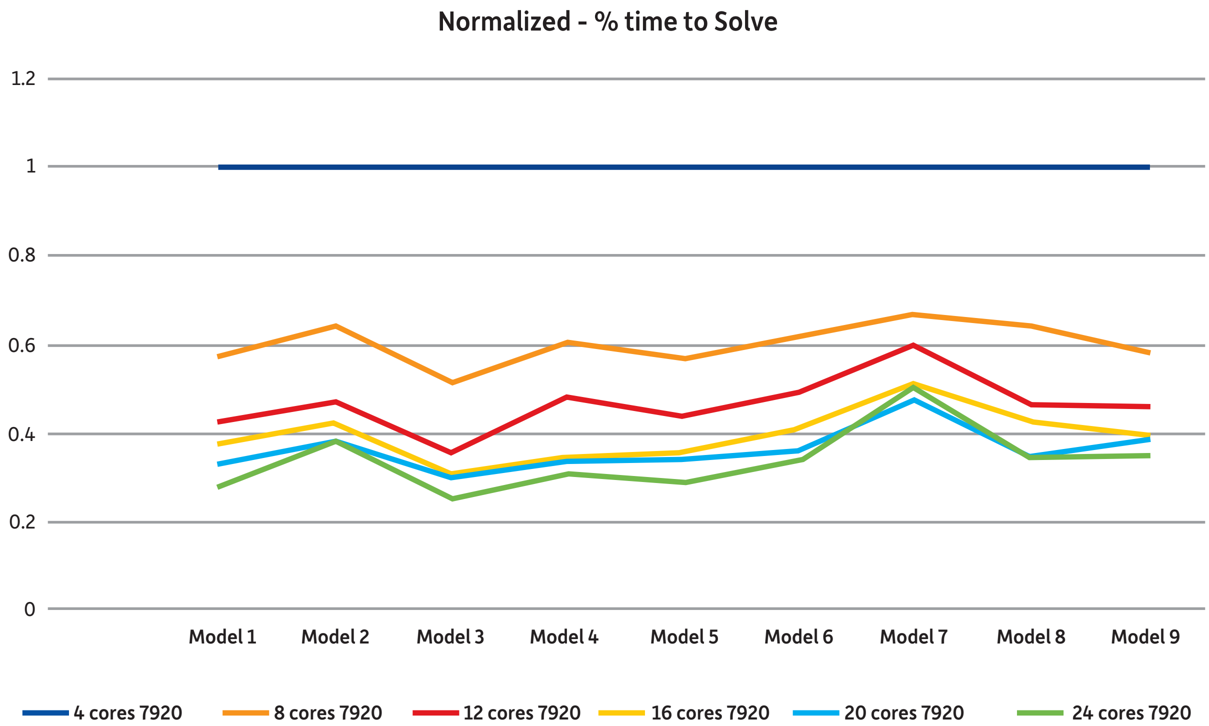

The older Dell Precision 7810 Tower was equipped with 64GB of RAM, compared to 192GB of memory in the Dell Precision 7920 Tower. To make the benchmarking more comparable, we reduced the RAM available to the 7920 to 64GB for our baseline when comparing the older and newer systems.

However, we will also highlight results from the full range of benchmarks we performed. For example, because the dual processors in the older Dell Precision 7810 Tower provided 6 CPU cores each, tests run on that workstation were limited to a maximum of 12 cores. But with its dual 12-core Intel Xeon processors, we were able to utilize 24 cores for each analysis in the modern Dell Precision 7920 Tower. Likewise, we will also call out the effects of using all 192GB of RAM that the Dell Precision 7920 Tower had to offer.

For all nine models we tested against, turning off hyperthreading led to the fastest results in all test cases. We would recommend you test the effect of turning off hyperthreading with your own data during deployment to ensure optimal results.

The Dell Precision 7920 Tower

The Precision 7920 Tower we benchmarked showed impressive performance, and it still had a lot of room to scale up. For example, it can be configured with up to 28 cores per processor or a total of 56 cores when customized with Dual processors; up to 3TB of 2666 MHz RDIMM/LRDIMM expandable memory; and up to 900W of graphics power with as many as 3 double-width graphics cards. Its FlexBay design holds up to 10 2.5- or 3.5-inch SATA/SAS drives or up to 4 M.2 or U.2 PCIe NVMe solid-state drives (SSDs). With the hot-swap feature on the M.2 and U.2 PCIe NVMe SSDs, you can remove drives without shutting down the workstation.

AI Can Optimize Your Workstation

Think about all the different potential hardware-software combinations possible, and you'll get an idea of the challenge faced by workstation manufacturers like Dell. Every manufacturer wants to provide the best computing experience, but how is that possible when there are so many different use cases? Dell uses artificial intelligence to optimize performance in its Precision family of workstations.

"We used to only have a static profile-based approach," says Dell Senior Product Manager Niraj Shah. "Our performance engineers would use industry standard benchmarks and work with partners to create a profile. We had good results with that, but the approach was not scalable because people were using different kinds of applications in different ways than the industry benchmarks."

In October 2018, Dell began factory installing the Dell Precision Optimizer Standard version 5.x on all Precision workstations. In the Premium version of the software, which is available as a 30-day free trial, the Dell Precision Optimizer software collects data as applications are running to characterize the individual user's workflow. Data such as CPU utilization, the number of threads being engaged, the top threads, memory utilization by process and cache by process all go through a machine learning model trained to optimize it.

Shah says Dell Precision Optimizer Premium has been able to increase the performance of common design engineering applications significantly, citing a 500% increase in PTC Creo and a 224% increase in Siemens NX software according to Dell internal testing.

"The idea is that now you have a tool that can use your application behavior, in the way you're using it, to come up with the right optimization for you,"

"The idea is that now you have a tool that can use your application behavior, in the way you're using it, to come up with the right optimization for you," Shah says. "Whether it's off-the-shelf or homegrown software, just point the file to Dell Precision Optimizer Premium and start using it as you normally would."

The Benchmarking Models

Model 1 (V19cg-1)

POWER SUPPLY MODULE

Degrees of Freedom: 5,300,000

Analysis Type: Static, linear, thermal

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	27%
7920 (no GPU) vs. 7920 (w/GPU)	59%
7810 (no GPU) vs. 7920 (w/GPU)	63%

7810 (w/GPU) vs. 7920 (w/GPU)

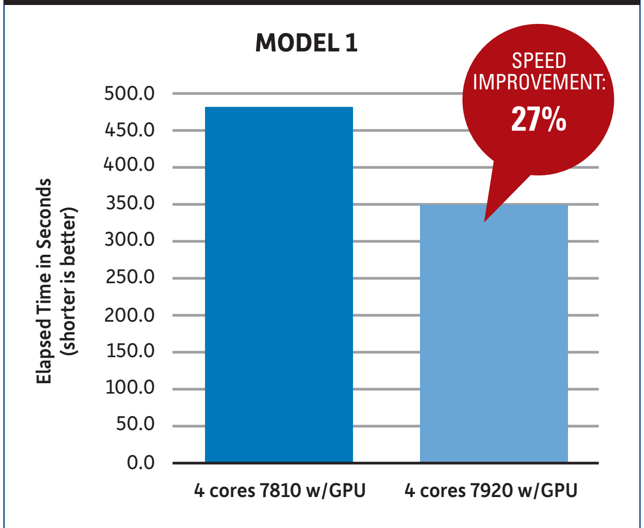

7810 (no GPU) vs. 7920 (w/GPU)

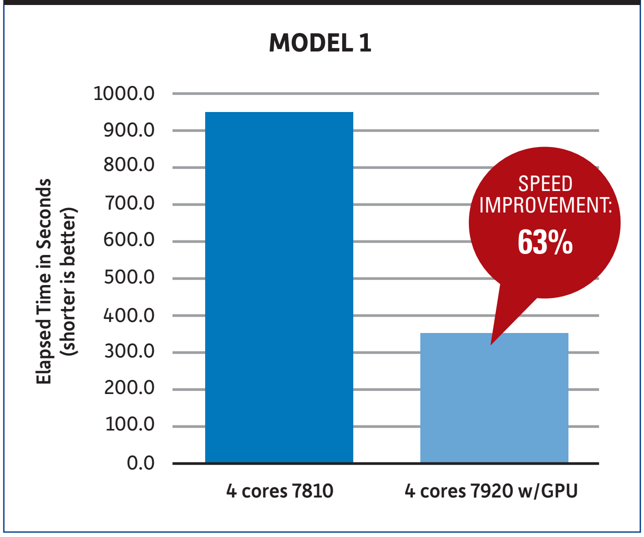

Jacobi preconditioned conjugate gradient (JCG) solver, ummetric matrix, static, linear, thermal analysis. This is a medium-sized job for iterative solvers and a good test of memory bandwidth.

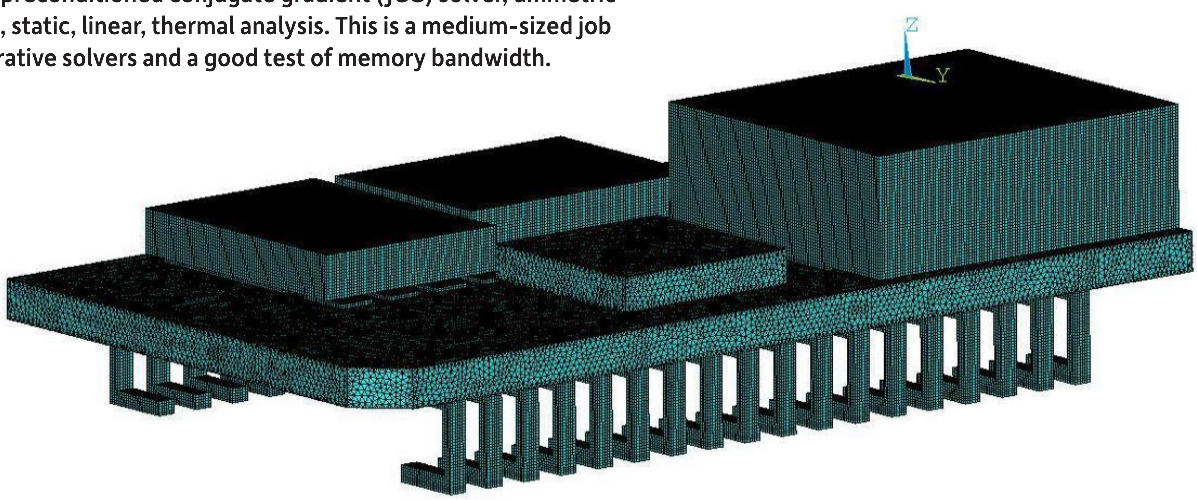

The Benchmarking Models

Model 2 (V19cg-2)

TRACTOR REAL AXLE

Degrees of Freedom: 12,300,000

Analysis Type: Static, linear, structural

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	45%
7920 (no GPU) vs. 7920 (w/GPU)	26%
7810 (no GPU) vs. 7920 (w/GPU)	40%

7810 (w/GPU) vs. 7920 (w/GPU)

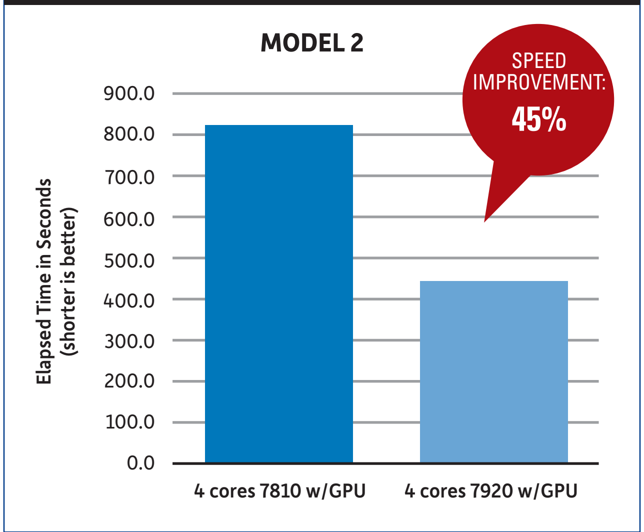

7810 (no GPU) vs. 7920 (w/GPU)

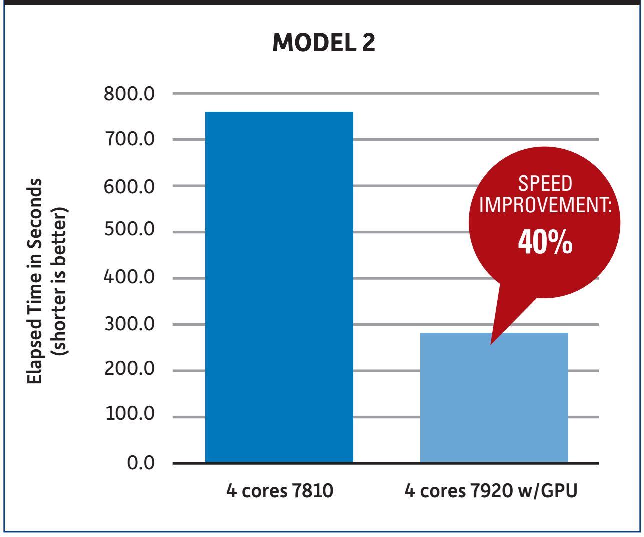

Preconditioned conjugate gradient (PCG) solver, symmetric matrix, static, linear, structural analysis. Large-sized job for iterative solvers, good test of memory bandwidth.

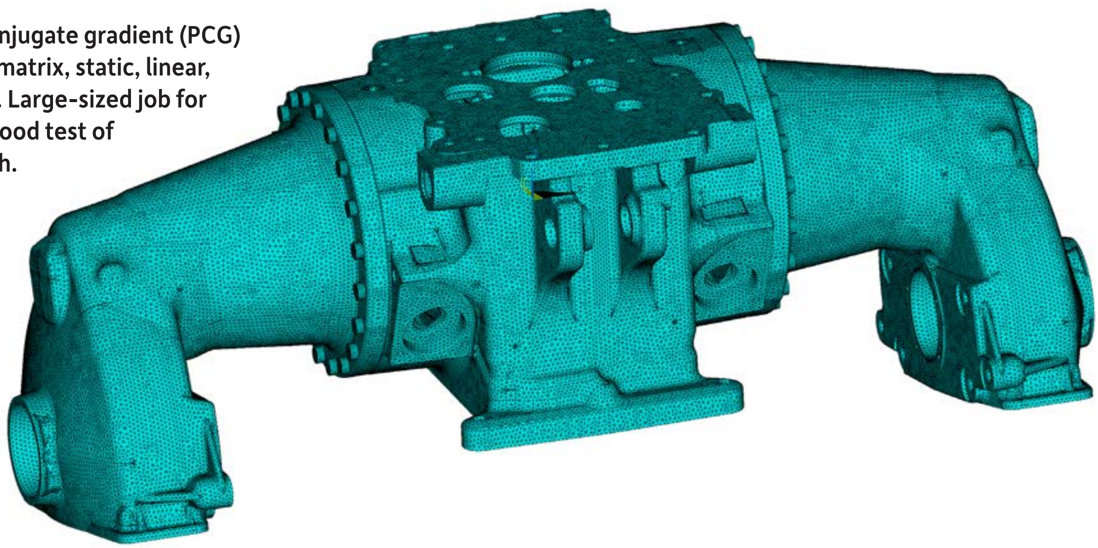

The Benchmarking Models

Model 3 (V19ln-1)

GEAR BOX

Degrees of Freedom: 7,700,000

Analysis Type: Modal, linear, structural (10 modes)

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	0%*
7920 (no GPU) vs. 7920 (w/GPU)	12%
7810 (no GPU) vs. 7920 (w/GPU)	32%

** Model 3 was a more CPU-intensive model, and showed no improvement when both the 7810 and 7920 were equipped with GPUs. However, when the Precision 7920 was equipped both with a GPU and more RAM (192GB), there was an 8% performance improvement.*

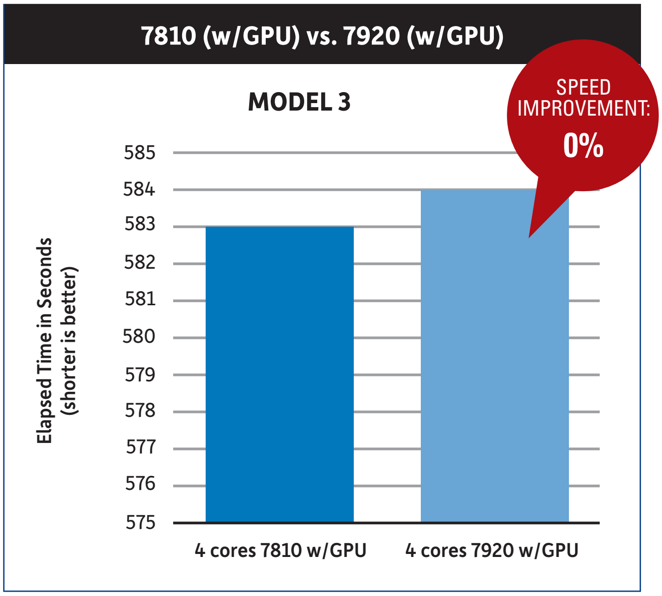

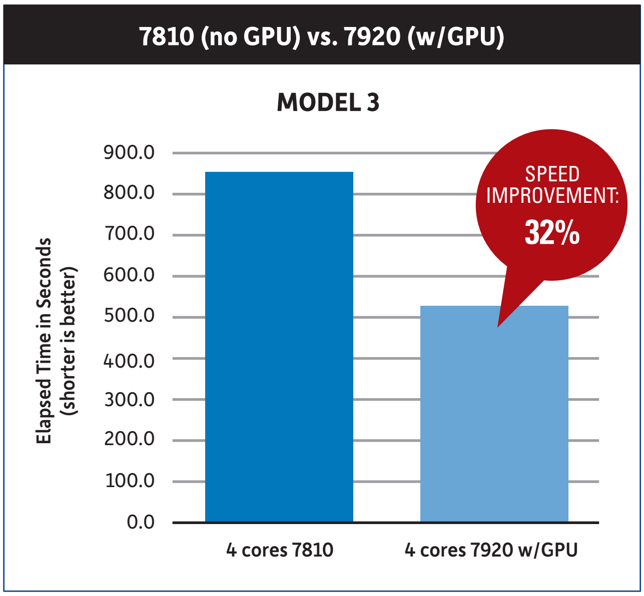

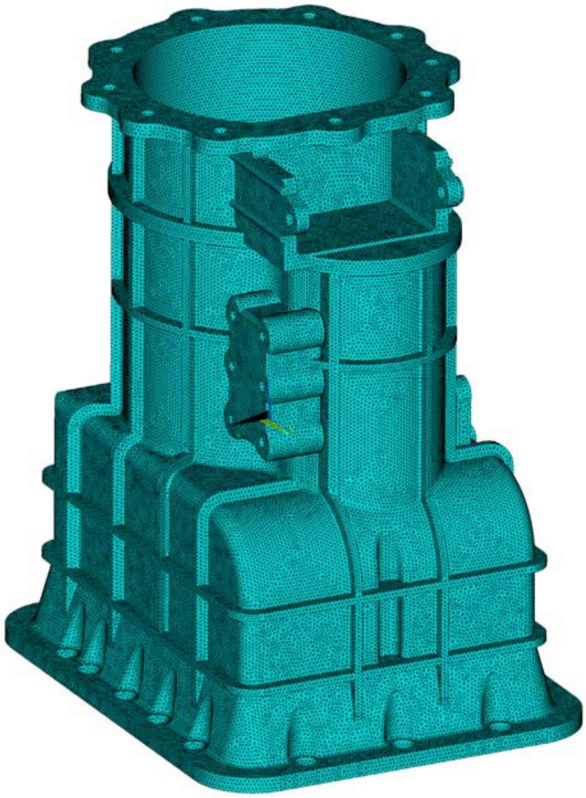

PCG Lanczos eignensolver, symmetric matrix, model, linear, structural analysis requesting 10 modes. Medium-sized job for iterative solvers, good test of memory bandwidth.

The Benchmarking Models

Model 4 (V19ln-2)

RADIAL IMPELLER

Degrees of Freedom: 2,000,000

Analysis Type: Modal, cyclic symmetry, linear, structural (50 modes)

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	36%
7920 (no GPU) vs. 7920 (w/GPU)	39%
7810 (no GPU) vs. 7920 (w/GPU)	60%

7810 (w/GPU) vs. 7920 (w/GPU)

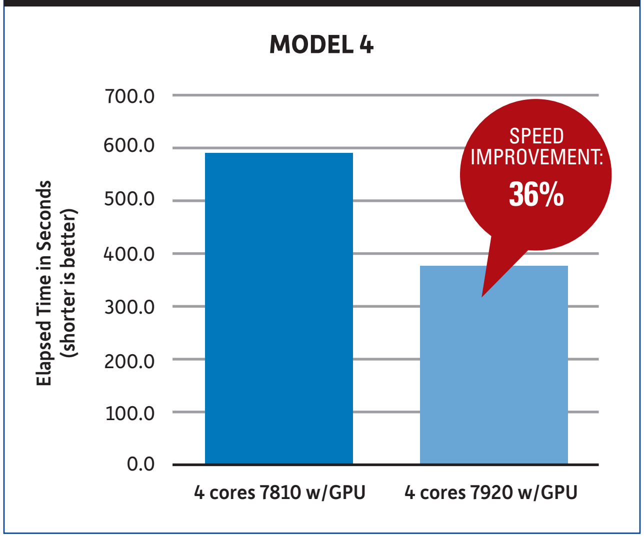

7810 (no GPU) vs. 7920 (w/GPU)

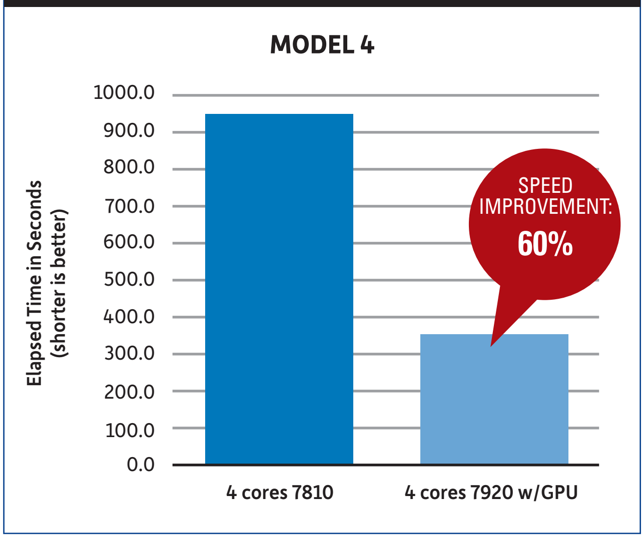

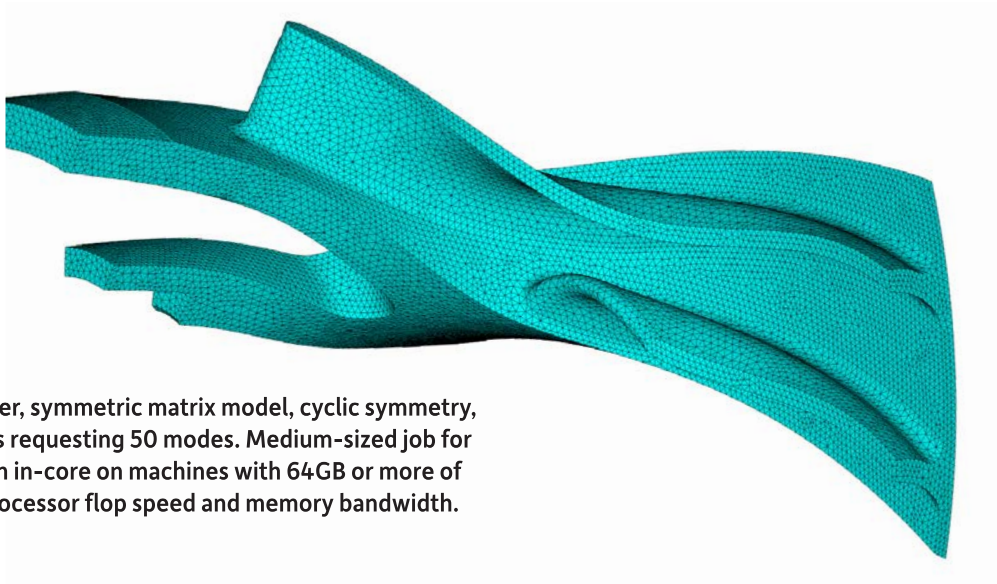

Block Lanczos eigensolver, symmetric matrix model, cyclic symmetry, linear, structural analysis requesting 50 modes. Medium-sized job for direct solvers, should run in-core on machines with 64GB or more of memory, good test of processor flop speed and memory bandwidth.

The Benchmarking Models

Model 5 (V19sp-1)

PELTIER COOLING BLOCK

Degrees of Freedom: 650,000

Analysis Type: Static, nonlinear, thermal-electric coupled-field

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	36%
7920 (no GPU) vs. 7920 (w/GPU)	24%
7810 (no GPU) vs. 7920 (w/GPU)	51%

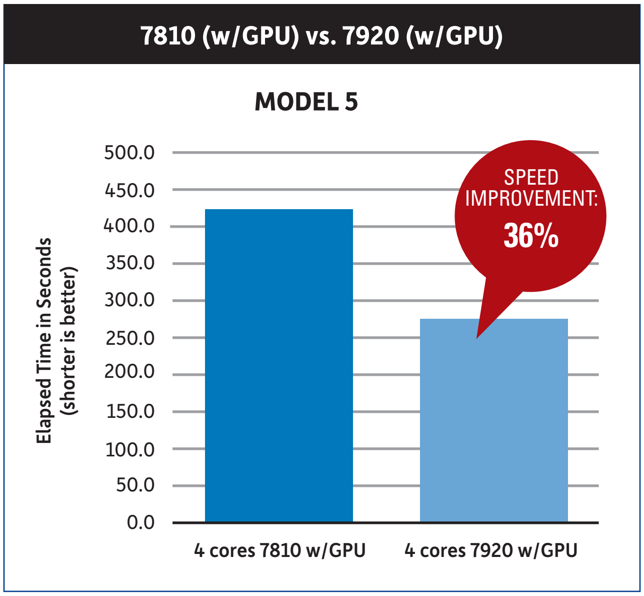

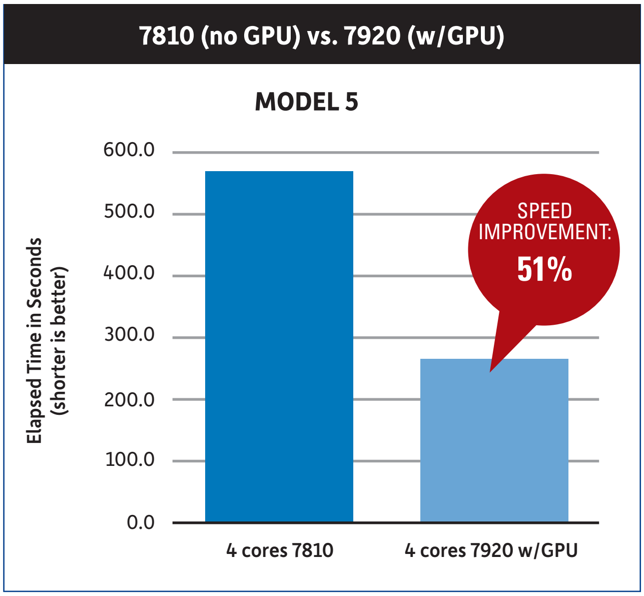

Sparse solver, non-symmetric matrix, static, nonlinear, thermal-electric coupled field analysis. Medium-sized job for direct solvers, should run in-core on machines with 32GB or more of memory, good test of processor flop speed if running in-core and I/O if running out-of-core.

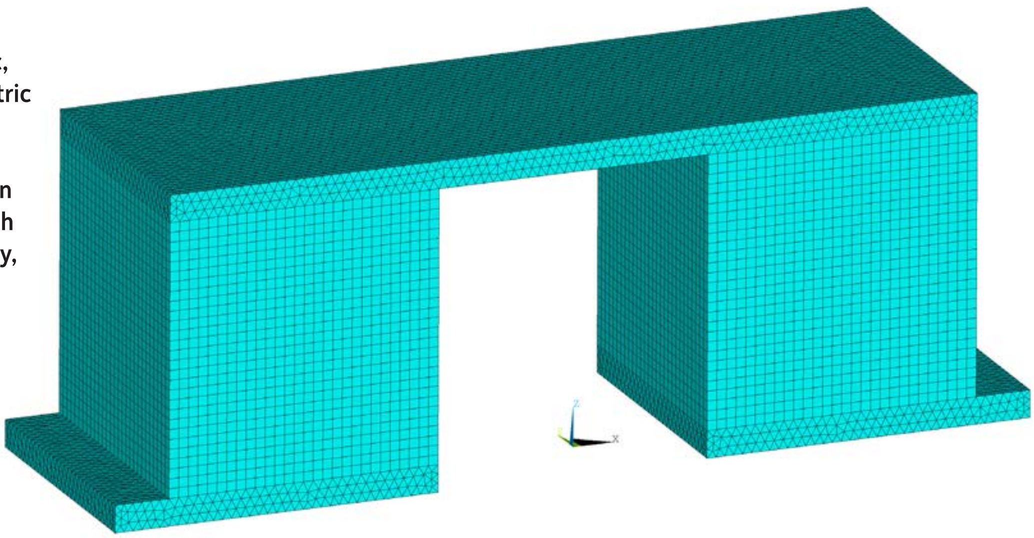

The Benchmarking Models

Model 6 (V19sp-2)

SEMI-SUBMERSIBLE

Degrees of Freedom: 4,700,000

Analysis Type: Transient, nonlinear, structural

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	33%
7920 (no GPU) vs. 7920 (w/GPU)	16%
7810 (no GPU) vs. 7920 (w/GPU)	44%

7810 (w/GPU) vs. 7920 (w/GPU)

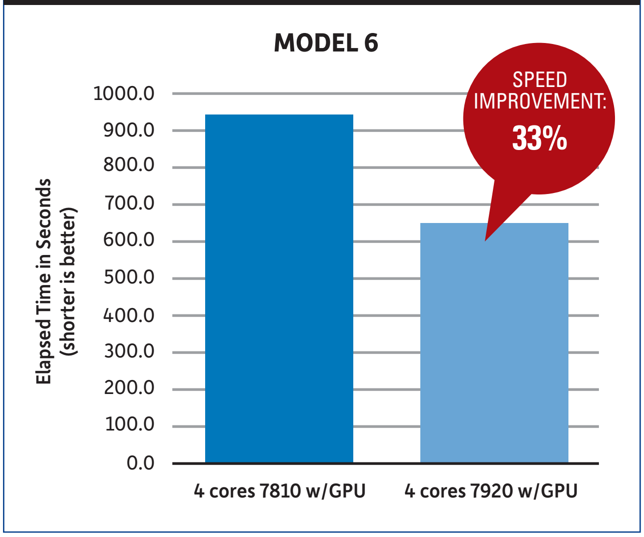

7810 (no GPU) vs. 7920 (w/GPU)

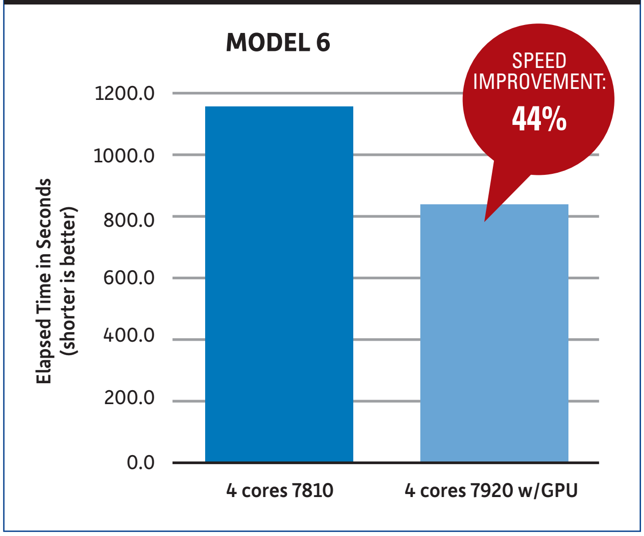

Sparse solver, symmetric matrix, transient, nonlinear, structural analysis. Medium-sized job for direct solvers, should run in-core on machines with 48GB or more of memory, good test of processor flop speed if running in-core and I/O if running out-of-core.

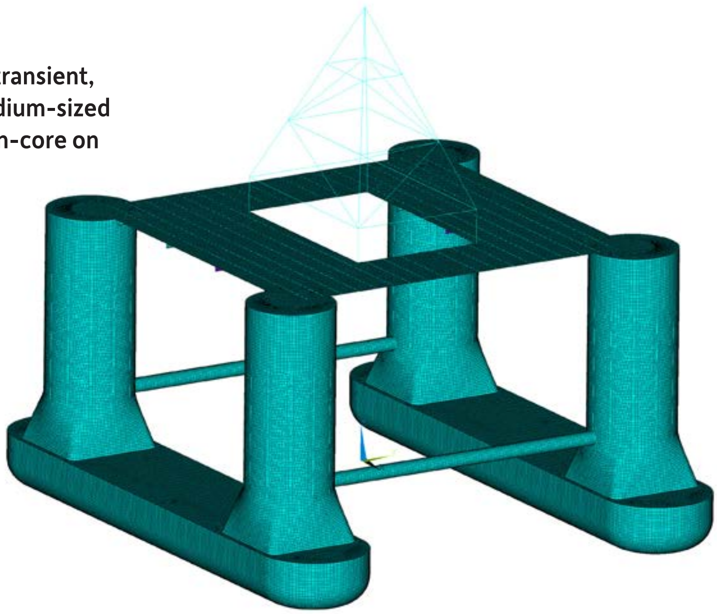

The Benchmarking Models

Model 7 (V19sp-3)

SPEAKER

Degrees of Freedom: 1,700,000

Analysis Type: Harmonic, linear, structural (1 frequency)

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	51%
7920 (no GPU) vs. 7920 (w/GPU)	46%
7810 (no GPU) vs. 7920 (w/GPU)	69%

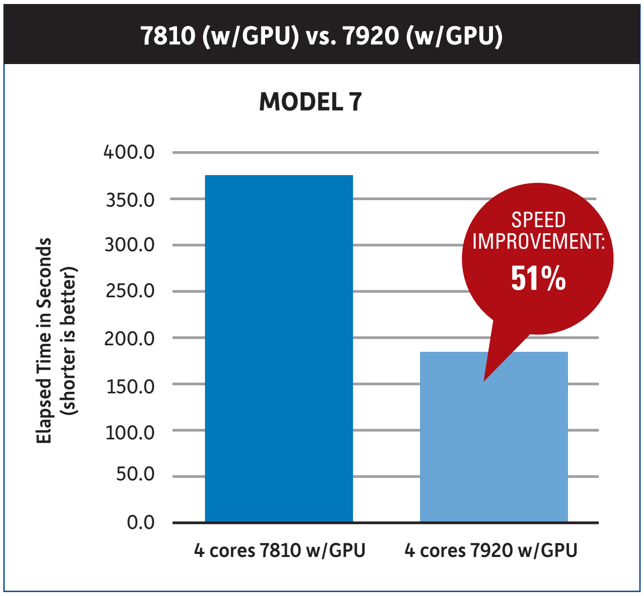

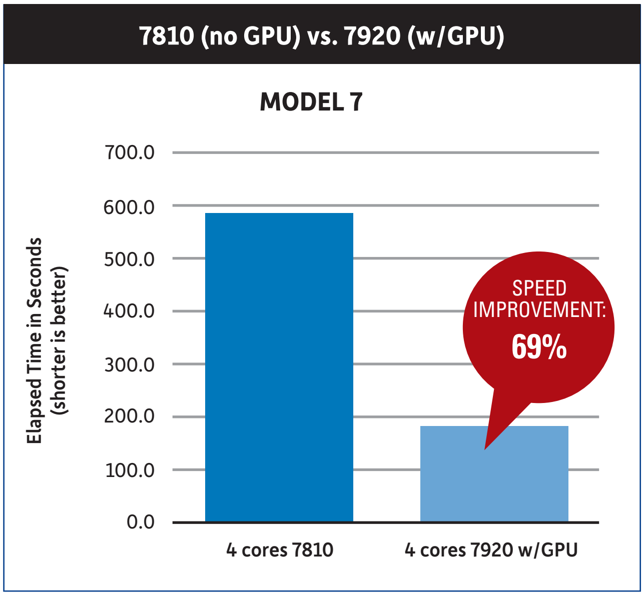

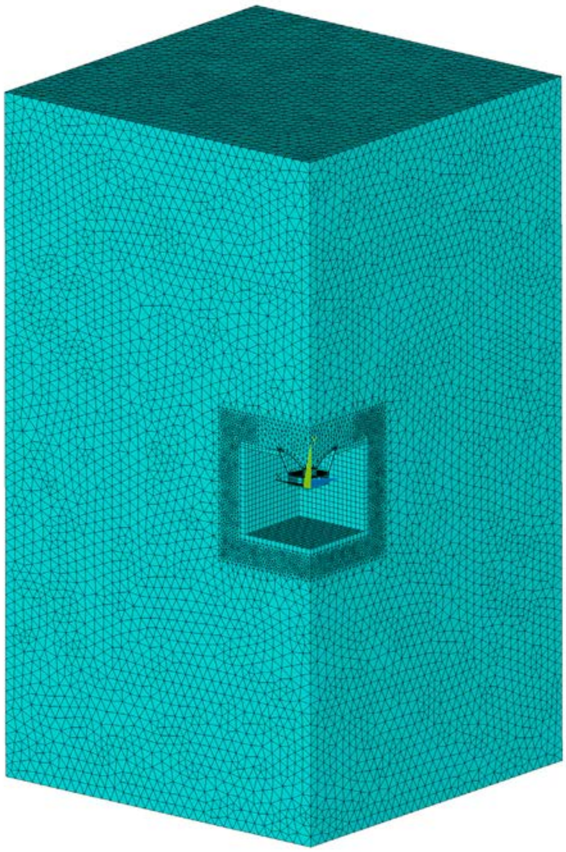

Sparse solver, symmetric matrix, harmonic, linear, structural analysis requesting 1 frequency. Medium-sized job for direct solvers, should run in-core on machines with 64GB or more of memory, good test of processor flop speed if running in-core and I/O if running out-of-core.

The Benchmarking Models

Model 8 (V19sp-4)

TURBINE

Degrees of Freedom: 3,200,000

Analysis Type: Static, nonlinear, structural (1 cumulative iteration)

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	50%
7920 (no GPU) vs. 7920 (w/GPU)	57%
7810 (no GPU) vs. 7920 (w/GPU)	74%

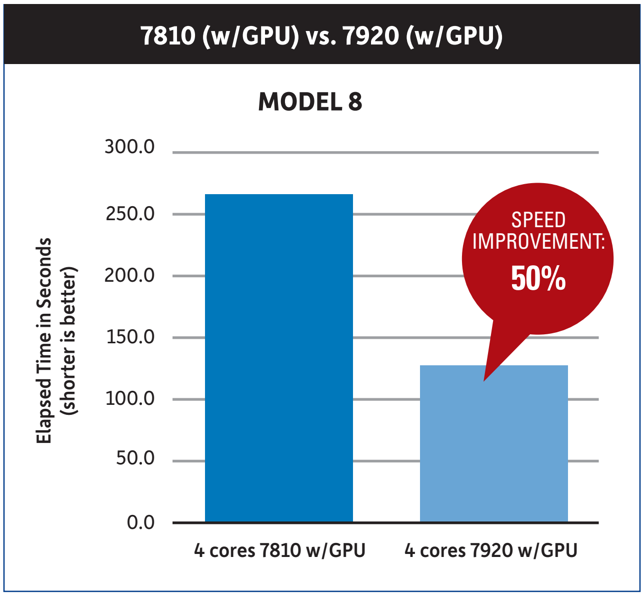

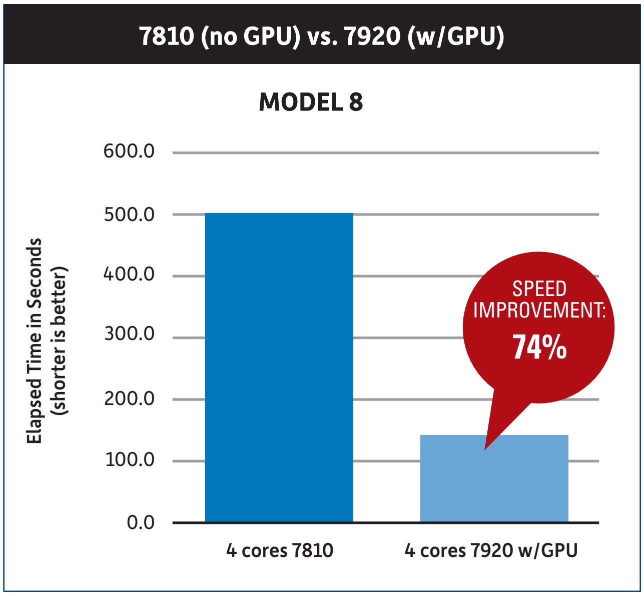

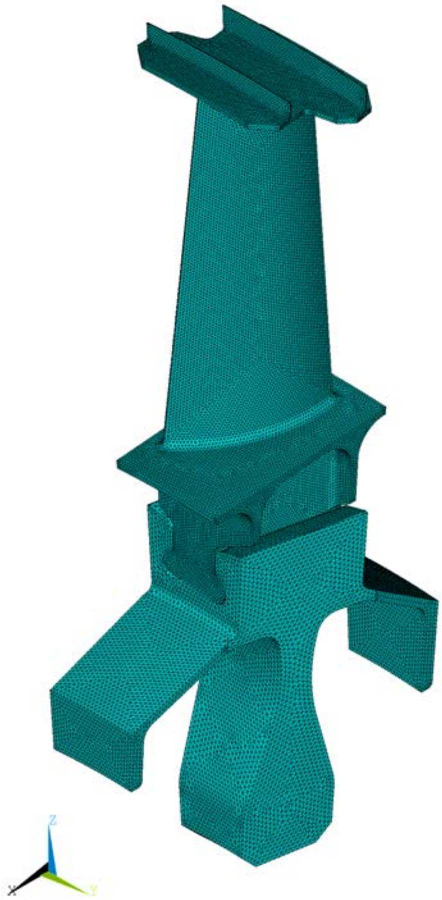

Sparse solver, symmetric matrix, 3200K DOFs, static, nonlinear, structural analysis with 1 iteration. Large-sized job for direct solvers, should run in-core on machines with 96GB or more of memory, good test of processor flop speed if running in-core and I/O if running out-of-core.

The Benchmarking Models

Model 9 (V19sp-5)

ELECTRONIC BALL GRID ARRAY (BGA)

Degrees of Freedom: 6,000,000

Analysis Type: Transient, nonlinear, structural (1 cumulative iteration)

Configuration: 4 cores, 64GB of RAM, Hyperthreading off, no DPO

SPEED IMPROVEMENTS	
7810 vs. 7920 (both w/GPUs)	53%
7920 (no GPU) vs. 7920 (w/GPU)	33%
7810 (no GPU) vs. 7920 (w/GPU)	67%

7810 (w/GPU) vs. 7920 (w/GPU)

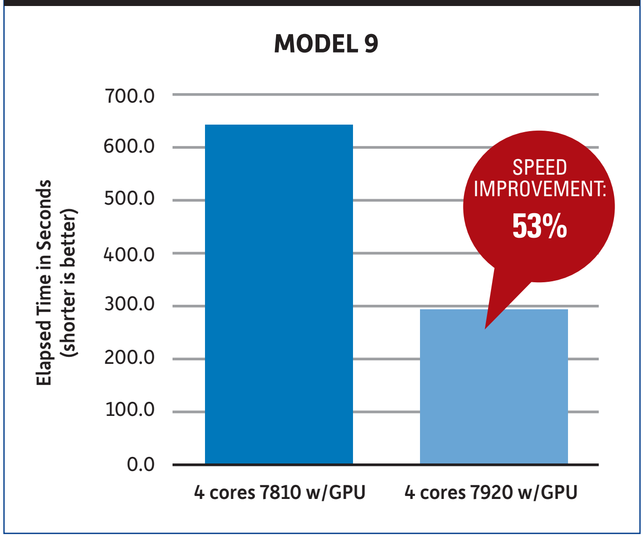

7810 (no GPU) vs. 7920 (w/GPU)

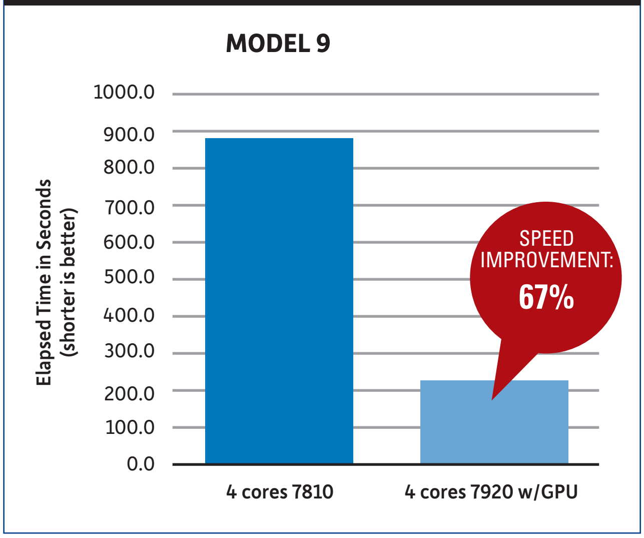

Sparse solver, symmetric matrix, 6000K DOFs, transient, nonlinear, structural analysis with 1 iteration. Large-sized job for direct solvers, should run in-core on machines with 128GB or more of memory, good test of processor flop speed if running in-core and I/O if running out-of-core.

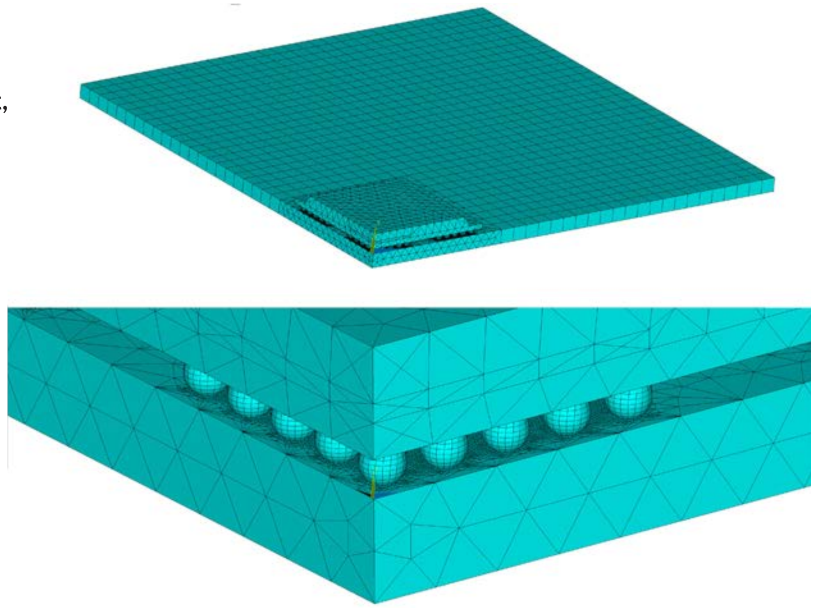

Benchmark Results

When benchmarking the simulation of a rear axle, for example, the newer GPU was responsible for 27% of the 45% improvement gains realized by the newer hardware.

Overall, the productivity study shows that a workstation upgrade delivers a significant performance improvement for ANSYS Mechanical 2019R1 users, dramatically reducing the time required to complete each simulation. On average, the benchmarks show that upgrading from an older workstation resulted in a 42% productivity improvement. The most significant improvement was seen while benchmarking the simulation of an electronic ball grid array, which completed 53% faster on the newer workstation than the older workstation when using 4 cores and 64GB of RAM on each. That same model completed 78% faster when allowing the newer workstation to use all of its 24 cores and 192GB of RAM.

In each model but one, the use of GPU computing contributed to reducing the amount of time required to complete each analysis. When benchmarking the simulation of a rear axle, for example, the newer GPU was responsible for 27% of the 45% improvement gains realized by the newer hardware.

When isolating the benefits of the GPU, we saw improvements up to 61% when benchmarking the newer workstation with and without the GPU, and gains up to 74% when comparing the older workstation without a GPU to the newer workstation with the NVIDIA Quadro GV100 GPU—keeping the core count at 4 and the RAM at 64GB. That improvement jumps to 81% when employing 12 of the Dell Precision 7920 Tower's 24 cores and all of its 192GB of RAM.

ANSYS has done considerable work to optimize for GPU computing within ANSYS Mechanical. In our testing, adding an NVIDIA Quadro GV100 (which has 5,120 CUDA cores) and holding all other hardware configurations in a steady state led to a 10% to 60% time savings (an average of roughly 30%) in elapsed time to solve the models. That can mean tremendous savings with larger jobs. For example, a week-long run could be reduced by as much three days just by selecting the correct GPU and enabling it with a single HPC license (see chart below).

Dell Precision 7920, With and Without GPU

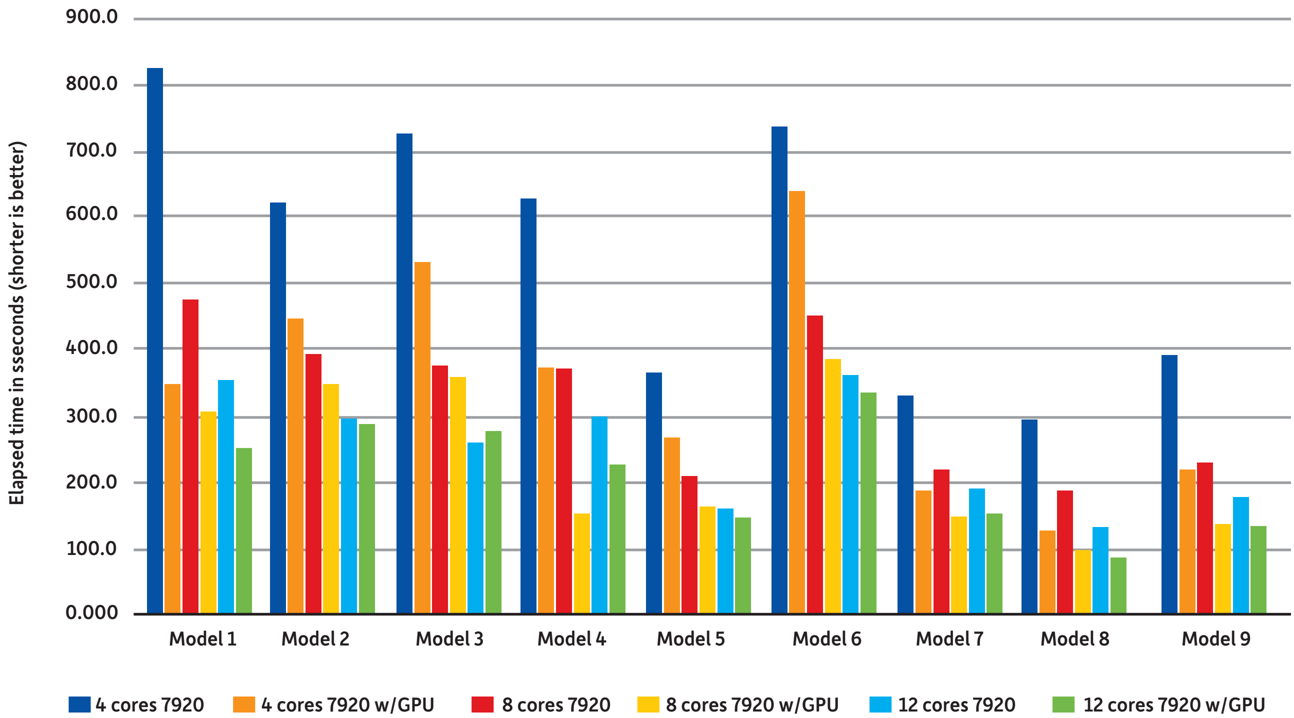

Model size will dictate memory requirements, so it is critical for users to ensure they have enough RAM to avoid caching to the hard drive. In our tests, 64GB of RAM was enough for most of the models, but in several cases where the dataset was larger, the overall time to solve was doubled if there was not enough available memory. In these use cases, the value of having 192GB versus 64GB of RAM was abundantly clear.

Leading Design

Design engineering teams looking for a product design and development edge in the age of digital transformation can make the case for the latest hardware and software. Workstations, GPUs and engineering software continue to advance in lock step, with one leapfrogging the other to the benefit of design engineers and simulation analysts who are taking on more challenging projects with tighter deadlines.

Being able to complete simulations almost 5 times faster by using newer hardware allows CAE experts to provide input earlier in the product development cycle, at a time when it has the most potential to optimize designs. The time saved can be used to iterate and virtually test more hypotheses. Likewise, making use of new software tools, such as ANSYS Discovery Live and Creo Simulation Live, helps design engineers better study their own models at the concept stage, helping to alleviate the bottleneck of analysts simulating small, early-stage design changes.

For firms using workstations that are not currently equipped with a GPU, the average solve time improvement of 30% (and as high as 60% for some models) possible using a new Dell Precision workstation with the NVIDIA GPU will make engineers more efficient and result in a return on investment in the new hardware in as little as a few months.

The benefits are not just limited to ANSYS Mechanical, as ANSYS offers a full suite of GPU-optimized applications. The investment in new hardware can be justified not just through the speed and performance increase when it comes to mechanical simulation, but via improvements across a range of tools and applications from ANSYS and its partners.

By empowering CAE analysts to do faster and more advanced simulation, and design engineers to take advantage of simulation, companies are making a sound investment in their future. Today's complex problems demand a thoughtful approach to product design, and boosting simulation productivity will not only accelerate the development cycle, it will result in innovation that delivers a competitive edge.

As the results of the benchmarking study show, combining current design and simulation software with modern hardware helps meet many technical challenges being faced by manufacturers today. Simulation-led design allows the design engineering team to explore more innovative solutions faster without compromising.

ANSYS Discovery Live Enables Real-Time Digital Exploration

ANSYS Discovery Live is expanding what the company calls Pervasive Engineering Simulation. The conceptual simulation software allows engineers to pose what-if questions upfront in the design process where most of the product costs are locked in, to explore thousands of design options and to receive immediate feedback. Discovery Live provides instant simulation tightly coupled with direct geometry modeling to deliver interactive exploration and rapid product innovation. ANSYS brings Discovery Live together with ANSYS AIM and ANSYS SpaceClaim into one product family geared toward design engineers.

From rapid initial concept exploration and 3D design to more detailed and comprehensive validation, the Discovery product family enables engineers at every stage. It also empowers design engineers to collaborate with simulation experts using ANSYS flagship products, for more comprehensive and detailed simulation of complex phenomena.Discovery Live is powered by NVIDIA graphics processing units (GPUs) and CUDA parallel computing that provide supercomputing capabilities to deliver results faster than more traditional methods.