Ansys

How to

Accelerate Simulation Software with Turnkey Hardware Systems

Discover the hardware solutions that can accelerate simulation tasks from 4X - 1,600X, which can supercharge engineering productivity and speed time to market.

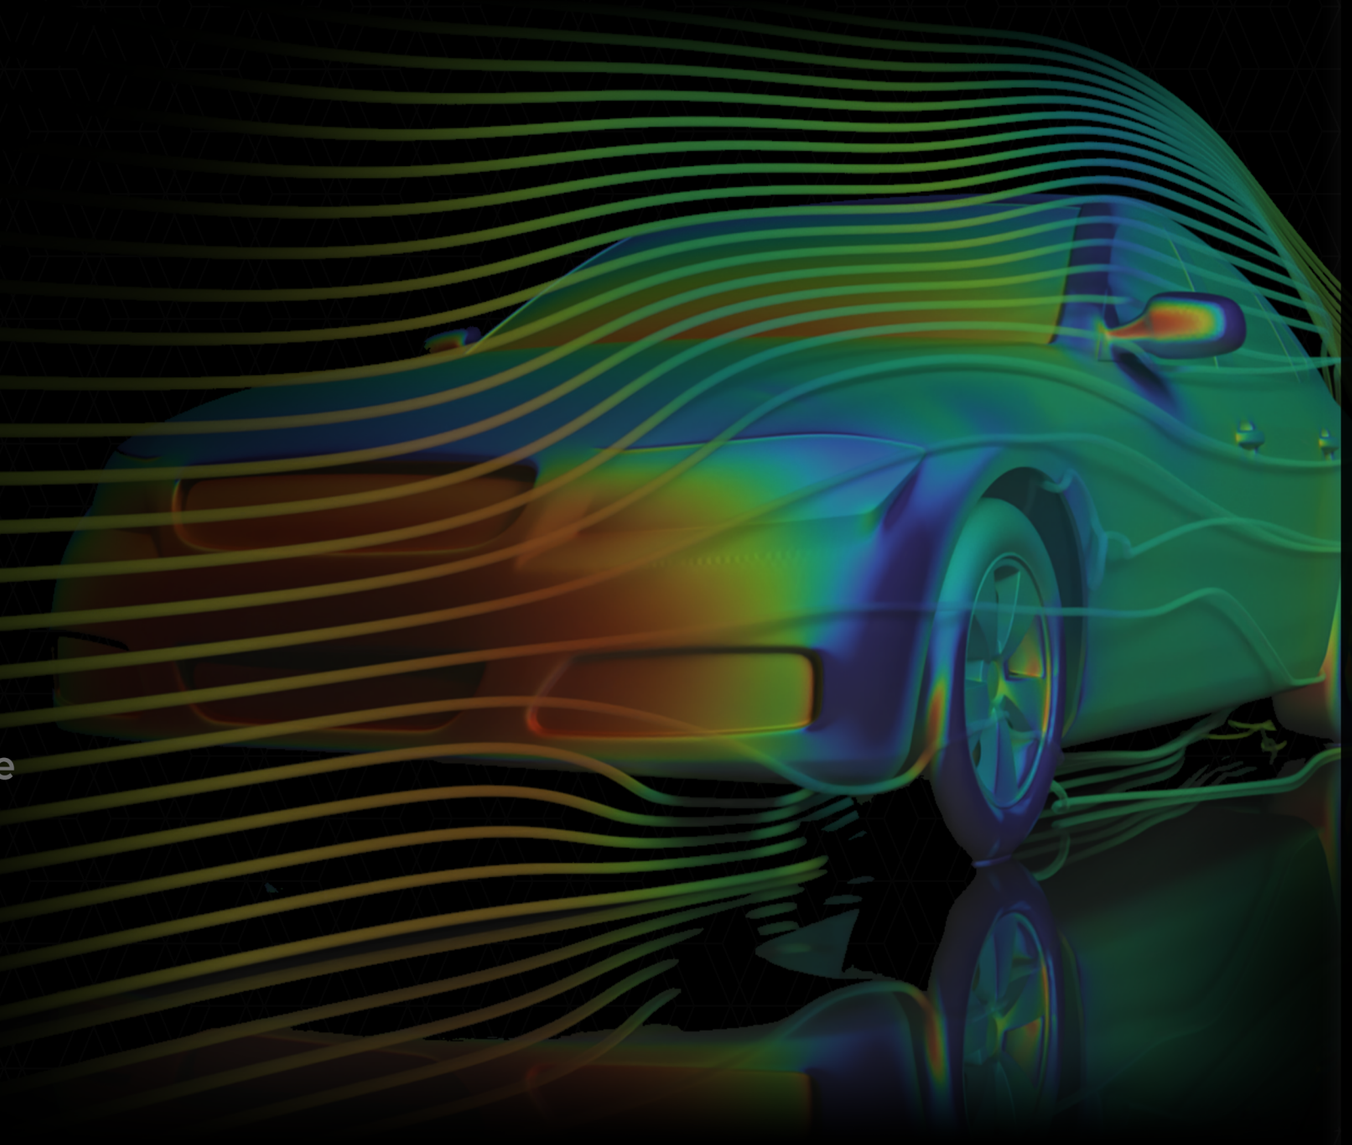

Executive Summary

Complicated problems take longer to solve. In the world of engineering, it's a given that simulation drastically accelerates the development cycle, enabling engineers to digitally test ideas and optimize products exponentially faster than physical testing. However, today's complex products and systems that require simulation of many different physics – say the thermal, structural, and electromagnetic forces at play in an electric vehicle or smart phone – can tax yesterday's computing environments. When multiphysics simulation software is designed to take full advantage of the latest computing hardware, companies can innovate even faster.

Solve Complex Simulations Faster

How much faster? In a series of third-party benchmarking studies conducted with Supermicro and NVIDIA hardware solutions, Ansys simulation software solved complex multiphysics simulations between 4X and 53X faster on optimized hardware, and 1,600X faster when using Ansys advancements in artificial intelligence to boost simulation speed. That's an almost unbelievable rate of acceleration, but it's possible with the right combination of hardware and software.

Make the Right Computing Choice

The computing landscape is always evolving, and Ansys consistently enhances its multiphysics simulation software. This constant innovation can make it challenging to determine where to invest resources for the best return. In this e-book, we break down the hardware and software environments that delivered these results – covering everything from specific processors, memory, and hard drives to the simulation model sizes and physics. But we don't just showcase the top-performing environments to impress you with maximum speeds; Instead, we present a range of configurations that offer varying levels of speed improvements, helping you find the optimal balance that suits your business needs.

Empower Innovation

Speed isn't everything in product design and development, but what organizations do with that saved time can make all the difference. When it comes to engineering, saved time can be invested to enable engineers to optimize existing products, simulate full systems to improve accuracy, and – perhaps best of all – try out more ideas that could lead to the next big thing. Beating the competition to market and optimizing your engineers' productivity are attractive perks as well.

BY THE NUMBERS:

USING NVIDIA-POWERED SUPERMICRO SYSTEMS, ANSYS SOFTWARE RUNS

UP TO 53X FASTER

AND SPEEDS UP AI TRAINING FOR SIMULATION

UP TO 1,600X FASTER

Our Strategic Partners

Ansys is collaborating with Supermicro and NVIDIA to deliver turnkey hardware, enabling unmatched acceleration for Ansys multiphysics simulation solutions. By leveraging the combined power of industry-leading hardware and software, Ansys customers can accelerate the solve time of larger, more complex models and further accelerate them through AI models to speed up simulations up to 1,600x faster. When running on Supermicro and NVIDIA technology, Ansys solutions reduce time-to-market and facilitate more robust design exploration for a wide range of applications, including automotive crash testing and external aerodynamics, aerospace gas turbine engines, 5G/6G antennas, and biopharmaceutical development.

Since its founding in 1993, NVIDIA has been a pioneer in accelerated computing. The company's invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined computer graphics, ignited the era of modern AI, and is fueling industrial digitalization across markets. NVIDIA is now a full-stack computing company with data-center-scale offerings that are reshaping industry.

Learn more: *nvidia.com*

Supermicro is a global technology leader committed to delivering first-to-market innovation for Enterprise, Cloud, AI, Metaverse, and 5G Telco/Edge IT Infrastructure. They are a Rack-Scale Total IT Solutions provider that designs and builds an environmentally-friendly and energy-saving portfolio of servers, storage systems, switches, software, along with global support services.

Learn more: *supermicro.com*

Showing Our Work

To help determine the optimal computing environment to speed up simulation tasks, such as large, complicated simulations and the use of simulation to train artificial intelligence algorithms, Ansys participated in three benchmarking studies in collaboration with NVIDIA, Supermicro, and MVConcept.

1. **Testing speed and scalability of four Ansys applications in the same GPU computing environment**
2. **Testing the speed and scalability of Ansys LS-DYNA® nonlinear dynamics structural simulation software in an NVIDIA Grace CPU Superchip environment**
3. **Testing the speed and scalability of Ansys Perceive EM™ radio frequency channel and radar signature simulation software in a GPU environment**

In addition to identifying the optimal hardware configuration, each study also analyzed the speed gains users could achieve by scaling the number of processors up or down.

Four Simulations, One Take-home Message

In the first study, Ansys collaborated with its HPC partner, MVConcept, a specialist in optimizing computer hardware and software environments. Working with experts from Ansys, NVIDIA, and Supermicro, four different simulation scenarios were tested using a Supermicro AS-4125GS-TNRT server configured with eight NVIDIA Hopper architecture-based H100 Tensor Core GPUs, as shown in this table.

The Software	The Simulation Models
Ansys Fluent® 2024 R2 computational fluid dynamic simulation software	A 250-million-cell, large-eddy simulation (LES) turbulence model, transient simulation of automotive external aerodynamics.
Ansys Mechanical™ 2024 R1 structural finite element analysis software	An 11-million element (50 million degrees of freedom), non-linear static simulation of a low-pressure turbine.
Ansys optiSLang® AI+ 2024 R1 process integration and design optimization software	Generating AI models of a parametric simulation trade study optimization using data from an Ansys HFSS™ 5G antenna module simulation.
Ansys Rocky™ 2024 R1 particle dynamics simulation software	A meshless, 16-million polyhedron particle discrete element method simulation of a rotating drum.

The Hardware: *Supermicro AS-4125GS-TNRT*

- GPU: 8x NVIDIA Hopper architecture-based H100 Tensor Core GPUs
- CPU: AMD 4th-Gen EPYC Genoa 9554 with 112 cores
- Memory: 24x 64GB (1,536 GB) DDR5 4800-speed RDIMM memory modules
- Solid-state drives (SSDs): 2x PCIe 4.0 NVMe 1.4 solid-state drives with 3,200GB and 7,680GB capacities
- Bios: Default settings for version 1.6a
- Operating system: Red Hat Enterprise Linux 9.4 (Plow)

In each of the benchmarks – which span a variety of physics and applications – Ansys multiphysics simulation software provided answers to complex engineering challenges significantly faster when taking advantage of modern hardware innovations. However, it's not just a matter of throwing more computing cores at the problem. Some hardware configurations provide only modest gains, but the benchmarks show that others provided 4X, 6X, 17X, 24X, 53X speed ups and more.

See what the researchers found to help you determine the right configuration of hardware and software to squeeze every bit of performance from your simulations in this e-book.

When you're ready, request your own free benchmark on a custom Supermicro system like the ones tested in this e-book to see how you can accelerate your Ansys simulations.

Ansys Fluent Results

Ansys Fluent software is used to model and analyze a variety of fluids phenomena, such as thermal management of electric vehicle batteries and e-motors, aircraft and automotive aerodynamics, wind turbine optimization, and many other fluid flows.

Fluent software is known for its efficient high-performance computing (HPC) capabilities, scaling to thousands of CPU cores. However, a new era has launched with the advent of GPU computing. GPUs offer the promise of significant increases in throughput for computational fluid dynamics (CFD) simulations. Ansys is at the forefront of this revolution with its GPU-resident implementation of the Ansys Fluent CFD solver, and this benchmarking study showcases the benefits of that. The table below shows the sizable improvement in speed when using just 4 NVIDIA H100 GPUs when compared to 512 CPU cores, a 12.6X speed up to be exact. When scaling up to 8 GPUs, there was an almost linear trend in solve time reductions.

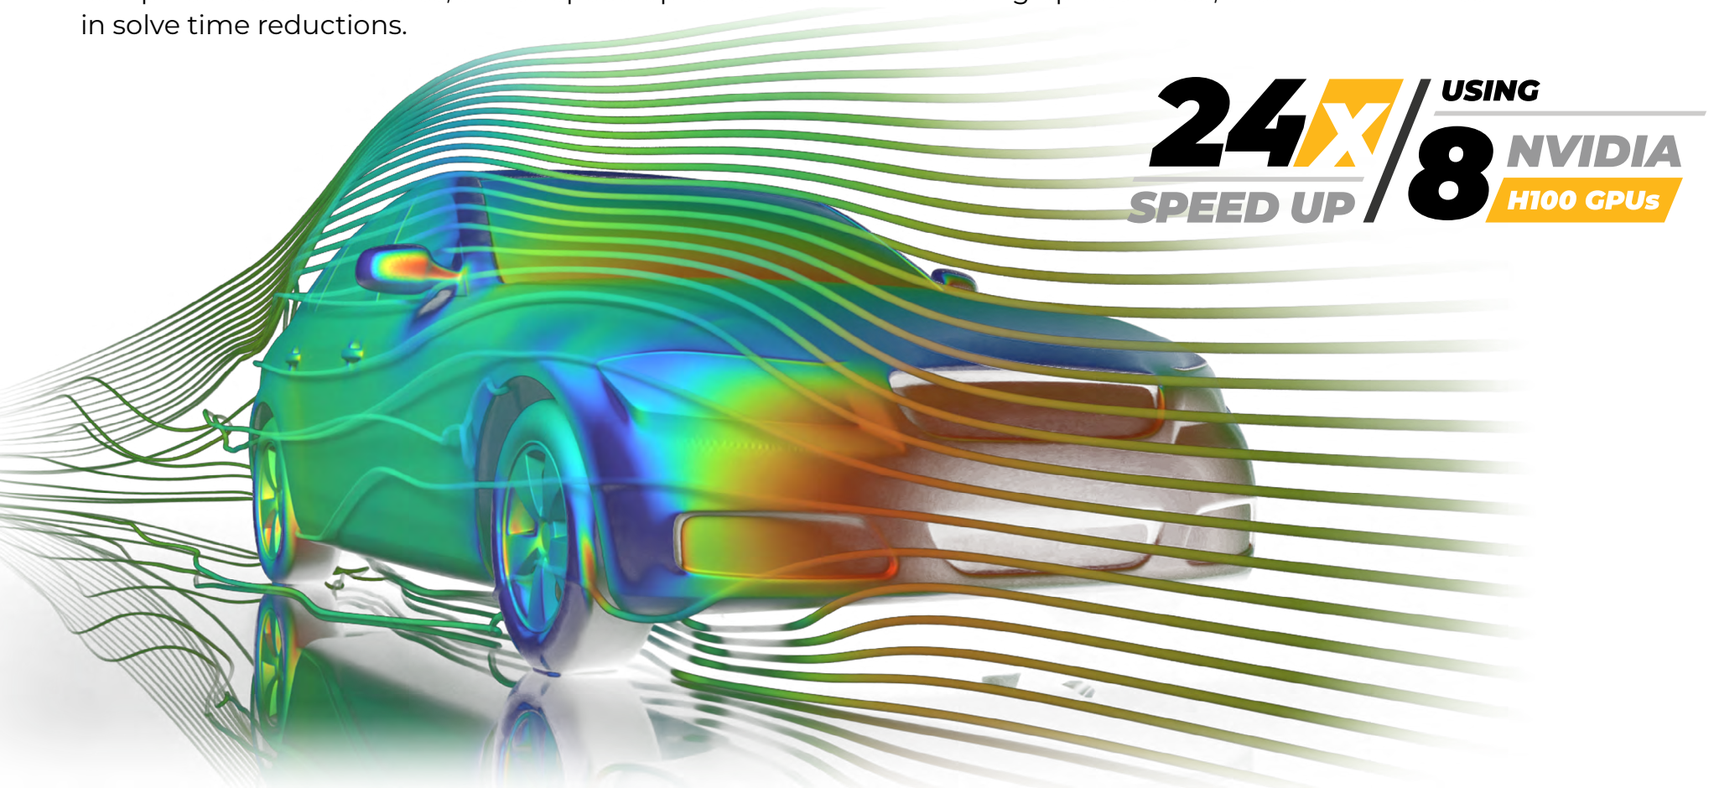

An Ansys Fluent simulation of automotive aerodynamics

ANSYS FLUENT SPEED UP

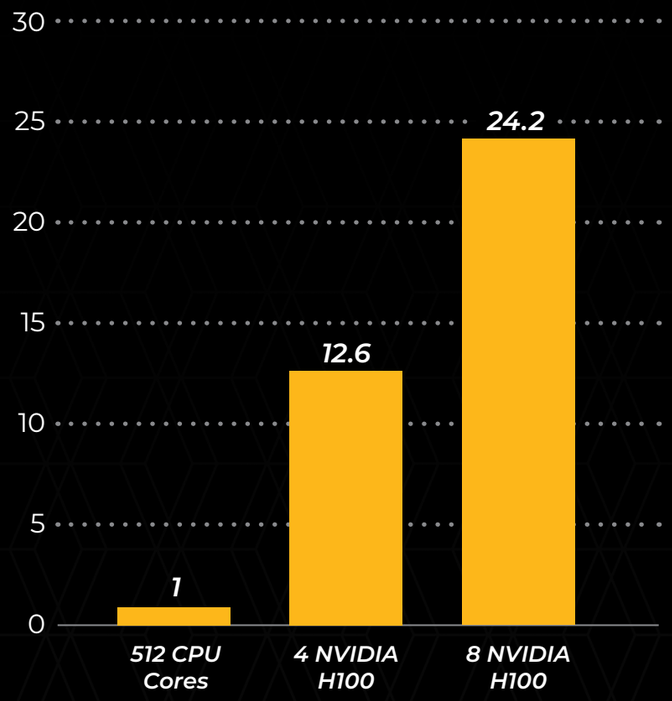

Linear scaling means that Fluent software's 12.6X speed up on four NVIDIA H100 GPUs nearly doubles (24.2X) when using eight NVIDIA H100 GPUs. The model used in the Fluent benchmark study is a 250-million-cell, large-eddy simulation of automotive external aerodynamics. The results imply that one GPU is equivalent to 1,500 CPU cores (24.2 speed up/8 GPUs=3.025 per GPU; 3.025*512 CPU cores=1,548).

Ansys Mechanical Results

Ansys Mechanical software predicts structural engineering issues such as stress, shock, vibration, deformation and more, including the interaction between mechanical and thermal or fluid physics. For decades, engineers seeking to understand how machine components will withstand wear and tear have trusted Ansys Mechanical software. Significant enhancements over the past few years, including optimized algorithms and new parallel processing technologies, are enabling Mechanical software users to solve models more efficiently on CPU and GPU hardware.

To showcase these enhancements, a non-linear static simulation of a low-pressure turbine (LPT) was simulated under a variety of computing environments. This is a common, but highly complex simulation involving nonlinearity and many degrees of freedom (DoF). The LPT portion of the gas turbine faces extremely high temperatures and plays a pivotal role in the overall performance of the engine.

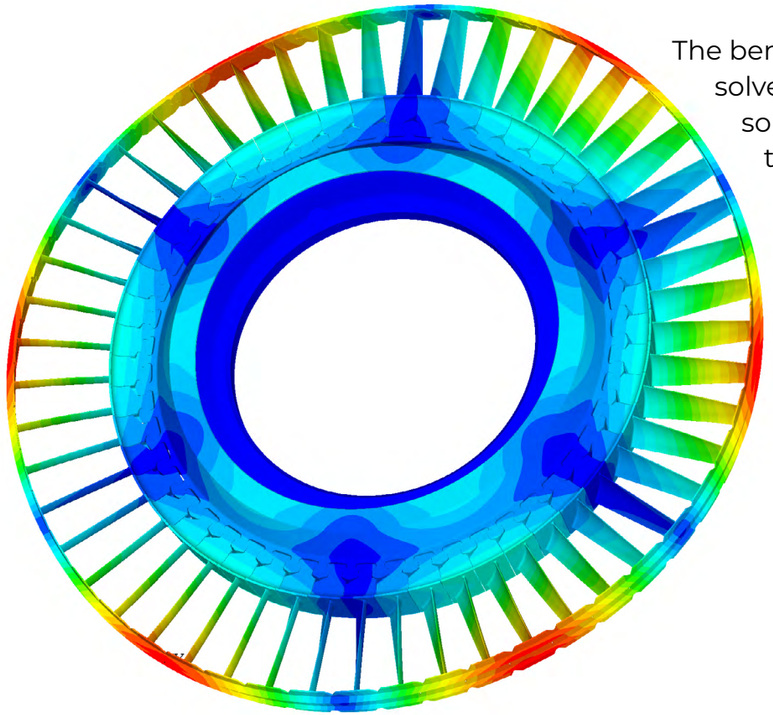

An Ansys Mechanical simulation of a low-pressure turbine commonly used in gas turbine engines.

The benchmark testing highlights the drastic reduction in simulation solve time when adding just a single GPU compared to a CPU-only solve. Six different benchmarks were performed, including three that were run on four, eight, and 16 CPU cores. These three CPU-only solves were then repeated with the same quantity of CPU cores but with the addition of a single NVIDIA H100 GPU.

6X SPEED UP / **ADDING 1 NVIDIA H100 GPU**

ANSYS MECHANICAL SPEED UP

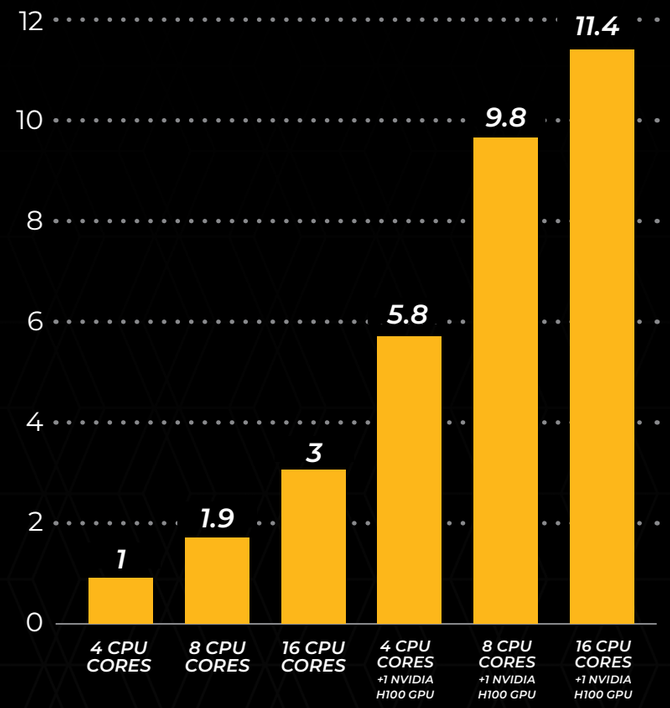

Results show that adding just a single NVIDIA H100 GPU to four CPU cores yields a 5.8X speed up when compared to running on just four CPU cores. Comparable speed up results can be seen for the other tests of eight and 16 CPU cores when adding a single GPU as well. These results highlight the value GPUs provide for Mechanical software users, enabling them to increase complexity of models, move from components to whole assemblies, include more nonlinearity, and evaluate and optimize more design scenarios.

Ansys optiSLang Results

Ansys optiSLang software is used to connect and automate toolchains when conducting parametric design studies. A parametric design study enables engineers to understand how a model should be shaped to meet different design goals, requirements, objectives, and more. Those parameters, or variables, can be virtually tested via algorithms in optiSLang to generate data for an optimal design. The Ansys optiSLang AI+ module uses artificial intelligence to speed up those optimizations. By training optiSLang software with simulation data derived from small parametric design studies that may take days to create, multiple AI models using different variables can then be run in seconds or minutes.

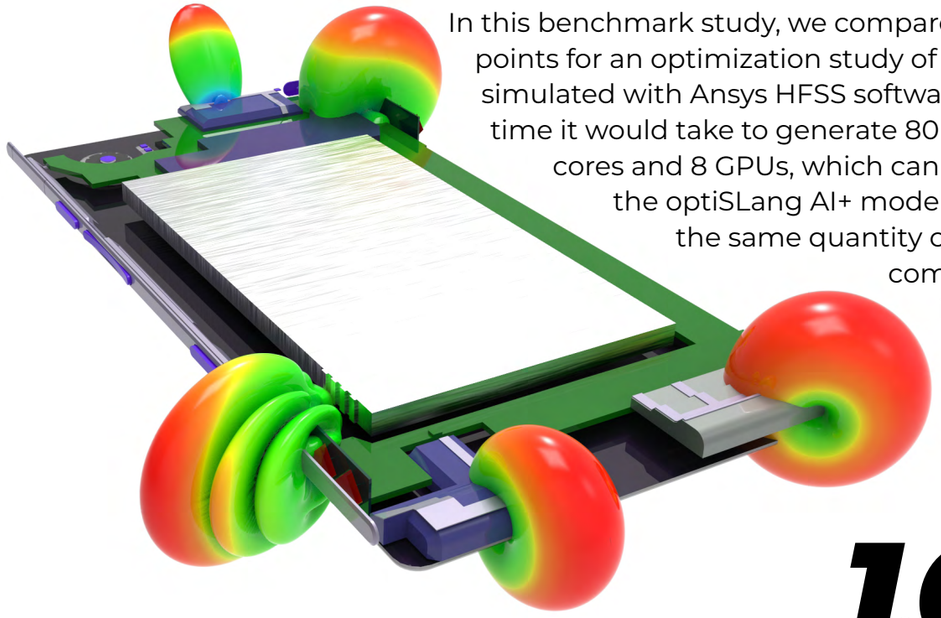

Ansys optiSLang software can be used to mathematically determine the best shape for product requirements.

In this benchmark study, we compared the time it would take to generate 80 design points for an optimization study of two 5G mmWave antenna modules in a cell phone simulated with Ansys HFSS software. The baseline result was calculated by finding the time it would take to generate 80 design points solving this HFSS model on 12 CPU cores and 8 GPUs, which can be seen in the HFSS portion above. We then used the optiSLang AI+ model to calculate how long it would take to generate the same quantity of design points when using the power of AI. It may come as no surprise that AI-based algorithms will result in significant speed ups, but this benchmark testing showed an astounding 1,600X speed increase using optiSLang AI+ software when compared to the non-AI based, HPC-accelerated approach.

1,600X SPEED UP USING OPTISLANG Ansys AI+ MODEL

ANSYS OPTISLANG SPEED UP

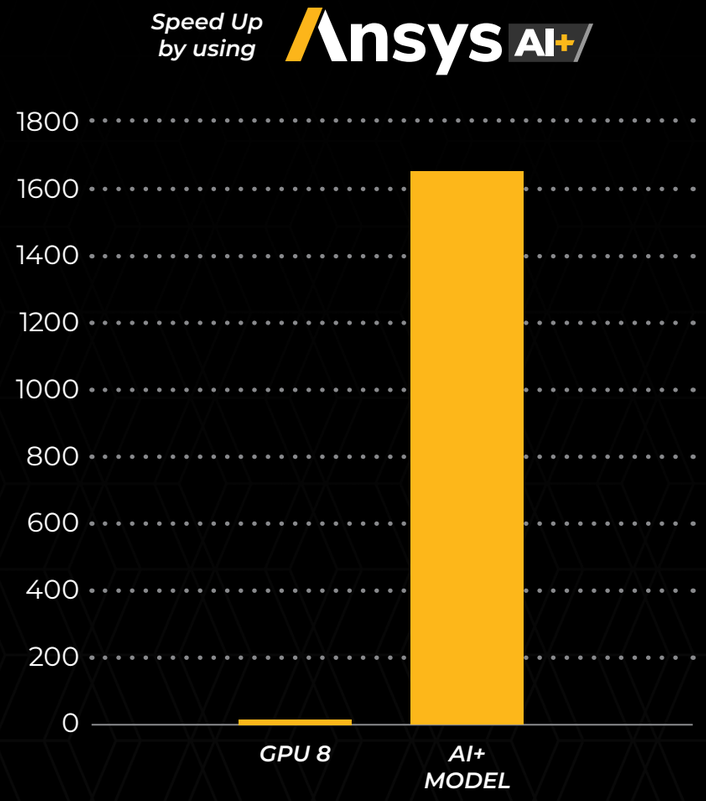

The dramatic AI+ speed increase makes it possible to run complex, multiphysics optimizations that were impossible before, resulting in new possibilities for customers to deliver lower cost, higher-performance products that were previously impossible.

Ansys Rocky Results

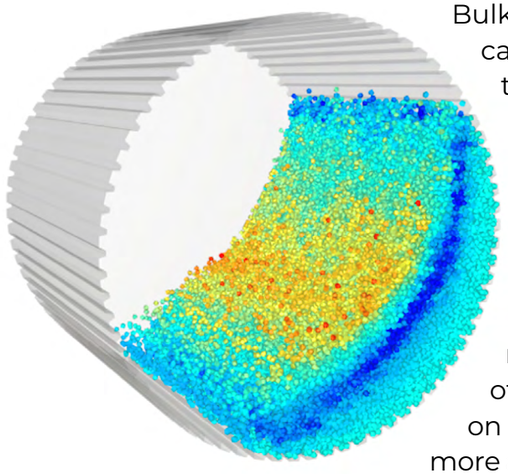

Ansys Rocky software simulation of a rotating drum of 16-million polyhedrons.

Bulk materials – from uniform shapes like pills to variable shapes like gravel and more – can be simulated using Ansys Rocky particle dynamics simulation software. The need to simulate how thousands or millions of particles flow and interact with each other, or objects around them, is something that crosses many industry sectors: food processing, drug manufacturing, construction, mining, and more.

The rotating drum simulation used in the benchmarking study is often used in the biopharmaceutical industry to apply coatings to tablets to achieve desired properties. The particles in the drum are modeled as polyhedrons, showcasing Rocky's unique ability to accurately model realistic, complex particle shapes. Ansys Rocky software can model the quantity, size distributions, and actual shape, rather than a representation of it, to accurately predict and optimize the coating process. Running these simulations on multiple GPUs enables customers to drastically accelerate compute times and perform more accuratge simulations by modeling the real shape of the particle and a larger quantity of them, which reduces the need for expensive scale-up experiments.

ANSYS ROCKY SPEED UP

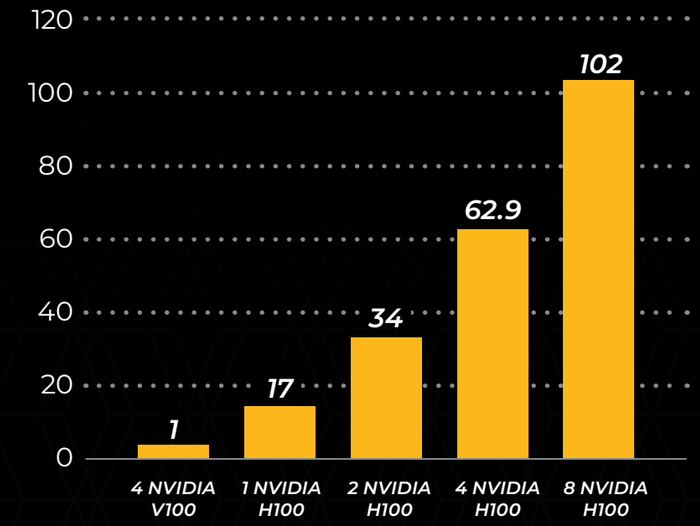

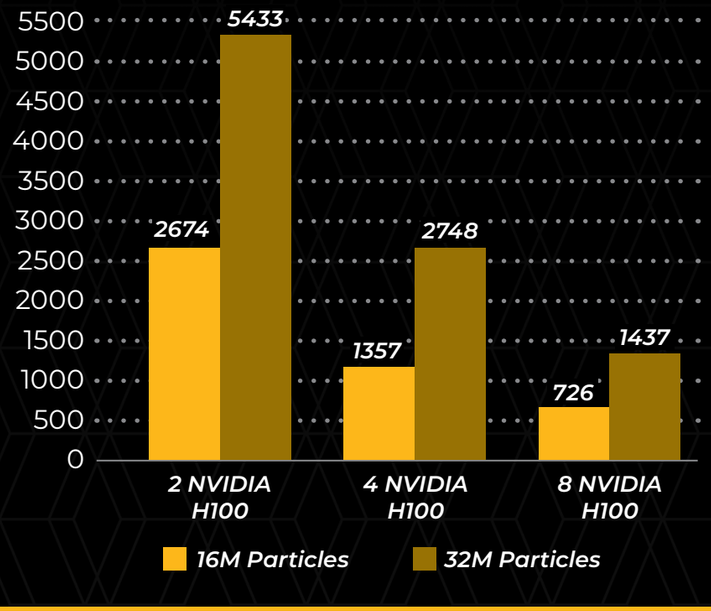

In the benchmarking study, 16-million polyhedral particles were simulated on a variety of NVIDIA H100 GPU configurations to see the performance speed ups when compared to this same model run on four NVIDIA V100 Tensor Core GPUs from a previous benchmark study. Results showed that using just one NVIDIA H100 GPU provided a 17X speed up when compared to the four NVIDIA V100 GPUs.

This model continued to see extremely strong speed ups as we scaled from one NVIDIA H100 GPU up to eight NVIDIA H100 GPUs.

17X SPEED UP / USING 1 NVIDIA H100 GPU

The benchmarking study also considered how different numbers of particles would affect Rocky simulation times. When simulating 16- vs. 32-million polyhedron-shaped particles in a rotating drum, the addition of additional GPUs is nearly linear compared to elapsed time. In other words, doubling GPUs from four to eight cuts simulation time roughly in half – from about 23 minutes for 16-million polyhedrons on four GPUs to about 12 minutes on eight GPUs; likewise from about 46 minutes for 32-million polyhedrons on four GPUs to about 24 minutes on eight GPUs. Outside the benchmarking lab, a model with more particles is often used to increase the accuracy of the simulation, so the near-linear acceleration will pay off for a number of use cases.

Ansys LS-DYNA Results

A separate benchmark study investigated computing environments that would accelerate Ansys LS-DYNA nonlinear dynamics structural simulation software. LS-DYNA software is an explicit simulation application – meaning it is mainly used to analyze sudden impacts, such as a car crash, aircraft bird strike, a helmet-to-helmet hit in football, or even the effects of an earthquake on a concrete dam. Engineers who want to simulate the response of materials to short periods of severe loading turn to LS-DYNA software.

For the benchmarking study, a virtual car crash simulation, known as ODB-10M, is recreated. It simulates a frontal crash scenario where the car hits an offset deformable barrier (ODB) to assess the frontal impact of vehicles to comply with, or exceed, safety requirements. The simulation was conducted in LS-DYNA software running in a Supermicro ARS-121L-DNR computing environment, equipped with the NVIDIA Grace CPU Superchip.

4X SPEED UP / **USING 1 NVIDIA GRACE SUPERCHIP**

The Hardware: *Supermicro ARS-121L-DNR*

- Processor: 2-node NVIDIA Grace CPU Superchip, 2x144 ARM Neoverse v2 cores
- Memory: 2x960GB LPDDR5X memory
- Solid-state drives (SSDs): PCIe Gen 5
- NVMe support
- Bios: AMI 32MB SPI Flash EEPROM

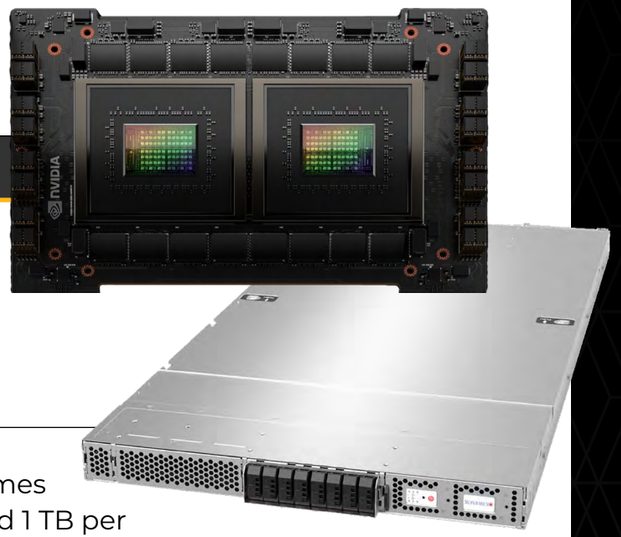

The NVIDIA Grace CPU Superchip is the industry's first ARM-based chip that consumes lower energy; resulting in power consumption savings when using the 144 cores and 1 TB per second of memory bandwidth. Using Ansys LS-DYNA software to solve these simulations on an ARM architecture provided a 4X speed up compared to 36 CPU cores. With the NVIDIA Grace CPU Superchip, engineers using LS-DYNA software can drastically reduce solve time to evaluate more collision scenarios for improved safety and reduce the time and expenses of physical testing.

ANSYS LS-DYNA SPEED UP

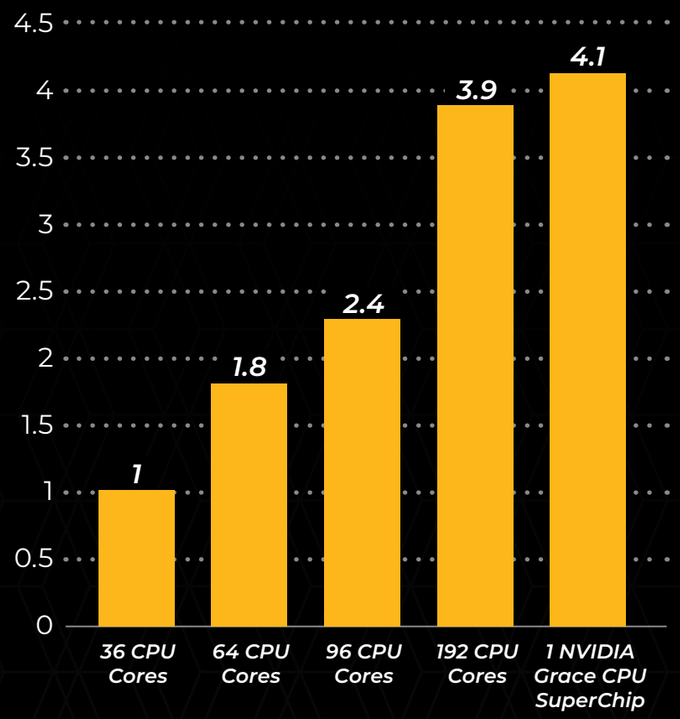

Ansys Perceive EM Results

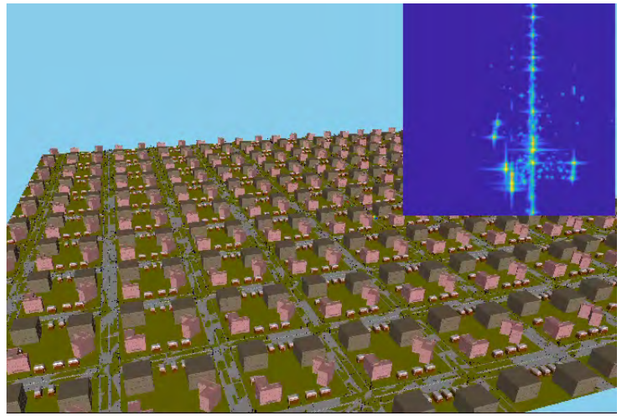

The Ansys Perceive EM benchmark computed eight radars (768 total channels) in a large city with street-level detail.

53X SPEED UP / USING 8 NVIDIA L40S GPUs

A separate benchmarking study examined the GPU performance of Ansys Perceive EM radio frequency wireless channel and radar signature simulation software. Perceive EM software is a high performance physical optics (PO)-based shooting and bouncing rays (SBR) technology delivered via a lightweight API that seamlessly integrates with any digital twin platform. Its rapid simulation performance is powered by advanced algorithms and a highly parallelized implementation on NVIDIA GPUs. This technology enables large-scale electromagnetic (EM) simulations, including automotive radar and RF communications scenarios, such as 5G and 6G.

Radar and wireless communications developers face signal propagation challenges. While artificial intelligence offers solutions, it also demands massive amounts of data for training. Ansys Perceive EM offers physics-accurate wireless channel modeling for on-demand synthetic data generation.

For the benchmarking study, eight radars across a large city were computed. Each radar had 96 channels, for a total of 768 radar channels that were computed in near real-time.

The Hardware: *Supermicro GPU SuperServer SYS-521GE-TNRT*

- GPU: 8x NVIDIA PCIe L40S
- CPU: 2x Intel 4th Gen Intel Xeon Platinum 8462Y+ with 32 Cores
- Memory: 32x 64GB DDR5 4800 ECC REG memory modules
- Solid-state drives (SSDs): 2x Micron 960GB M.2 22x80mm NVMe PCIe 4.0 TLC solid state drives
- BIOS: AMI 32MB SPI Flash EEPROM

ANSYS PERCEIVE EM SPEED UP

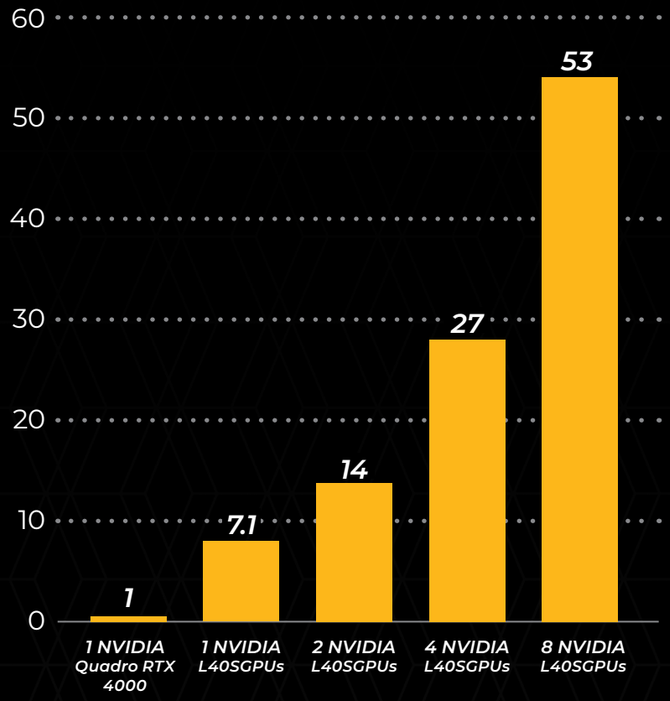

Ansys Perceive EM software running on one NVIDIA GPU achieves the same performance as 1,000,000+ CPU cores. The benchmarking study showed a 53X speed up running on eight NVIDIA Ada L40S GPUs in a Supermicro GPU SuperServer SYS-521GE-TNRY environment, compared to one NVIDIA Quadro RTX 4000 GPU.

Benchmark Your Models

The benchmarking results presented in this e-book show what is possible when advanced simulation software from Ansys is paired with modern computing environments. While our benchmarking examples cover a wide range of applications and physics, the most crucial benchmark is your own. That's why we've created a free performance benchmark program specifically for you.

Instead of demonstrating evidence of the benefits of HPC and GPU computing on standard benchmark models, we want to show you the time savings and scalability that is possible for your own model. Stop reducing the size and details of your simulations to reduce run times, and get all the information you need from high-fidelity, multiphysics simulations.

The Benchmarks Defined

CPU, central processing unit: A computer component that processes instructions. It consists of at least one processor core that performs calculations, but may include multiple processor cores in a single chip. A computer may have multiple CPUs, each of which may include multiple cores.

GPU, graphics processing unit: A computer component originally designed to process graphics instructions, but now also commonly used for general-purpose computations to improve overall performance. A computer may have multiple GPUs.

MIUPS, Million Iteration Updates Per Second: A benchmarking method used to standardized testing across products.

Process: In this e-book, "process" is used to define the number of compute tasks that are asked by the user. The number of processes is not always equal to the number of physical cores and logical cores used on the CPUs and/or GPUs.

Efficiency: In this e-book, "efficiency" is the speedup normalized by the number of processors used, presented as a percentage. It indicates the overall utilization of the CPUs during a parallel calculation. The remaining time is spent waiting for other functions, such as parallel communication or work on other processors, to complete.

Speed up: In this e-book, "speed up" is the ratio of wall-clock time required to complete a given calculation using a single processor to that of the equivalent calculation performed on a concurrent machine. Its value ranges from 1 to the number of processors used for the parallel run. When the speedup is equal to the number of processors used, the speedup is called perfect or linear. When the speedup exceeds the number of processors, this is referred to as super-linear, and is often caused by the availability and use of larger amounts of "fast" memory (e.g. cache or local memory) compared to the single processor run.

Rating: In this e-book, "rating" is the primary metric used to report performance results of the benchmarks. It is defined as the number of benchmarks that can be run on a given machine (in sequence) in a 24-hour period. It is computed by dividing the number of seconds in a day (86,400 seconds) by the number of seconds required to run the benchmark. A higher rating means faster performance. The ratings in our reports are based on Total Wall Clock Time.

REQUEST YOUR FREE BENCHMARK

Discover how to accelerate your Ansys simulations, improve the productivity of your engineering team, and reap the full benefits of simulating more complex systems faster than ever.

REQUEST YOUR FREE BENCHMARK

Powering Innovation That Drives Human Advancement

ANSYS, Inc.
Southpointe
2600 Ansys Drive
Canonsburg, PA 15317
U.S.A.
724.746.3304

Visit **www.ansys.com** for more information.