HPC brings major benefits to almost every industry and application. In turbomachinery design, small performance increases can add up to significant cost benefits. To achieve this, large models are required for accuracy and to model complete systems, and many design iterations are conducted. HPC is a key enabler of both high-fidelity simulation and design exploration.

Bigger, Better, Faster: HPC Technology Leadership

Intense focus on HPC software development enables breakthrough productivity on current and emerging hardware solutions.

By Ray Browell, Lead Product Manager, and Barbara Hutchings, Director, Strategic Partnerships, ANSYS, Inc.

Today, ANSYS customers in a variety of industries are running simulations in high-performance computing (HPC) environments, using multiple processors — and often large clusters of computers — to solve numerically intense problems using parallel processing. These customers conduct larger, more detailed and more accurate simulations, study entire systems and the interaction between components, and gain greater confidence that their work will predict the actual performance of physical products subjected to real-world forces. Such productivity comes as a result of significant software development efforts, which focus on delivering optimized performance and scaling on the latest — and ever-changing — computer hardware.

In fact, HPC has become a software development imperative. As processor speeds have leveled off due to thermal constraints, hardware speed improvements are now delivered through increased numbers of computing cores. For ANSYS software to effectively leverage today's hardware, efficient execution on multiple cores is essential. As core counts continue to increase, and with the advent of many-core processors and availability of graphics processing units (GPUs), ongoing software architecture changes must be made to maintain productive use of high-performance computing. Software development to build and maintain parallel processing efficiency is therefore a major ongoing focus at ANSYS.

Performance and Scaling

The latest HPC enhancements in ANSYS software deliver breakthrough performance milestones. For structural analysis, Distributed ANSYS solvers now show outstanding scaling for large problems out to dozens — and even hundreds — of cores. Significant scaling improvements have been achieved for the direct sparse solver by applying parallelism to the equation re-ordering scheme. For the iterative PCG solver, parallel treatment of the pre-conditioner significantly extends scaling. These core solver improvements remove scaling bottlenecks, pushing performance out to

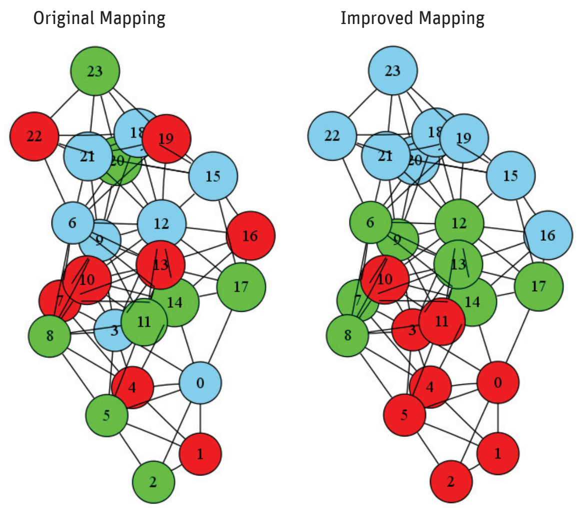

Network-aware partitioning in ANSYS Fluent enables faster solutions by minimizing network traffic.

much higher core counts and enabling efficient processing of ever-larger, more detailed structural simulations.

Scaling of ANSYS fluids simulations also continues to reach new heights. Scaling improvement in the recent version of ANSYS Fluent (13.0) was enabled by a new solution methodology called hybrid parallelism. This is a significant algorithmic change, enabling the core linear solver to exploit fast shared memory parallel communications within a machine, in combination with distributed memory parallel communications based on message passing interface (MPI) software between machines. The result is much faster execution on multicore systems, leading to outstanding scaling. Hybrid parallelism is also being applied to specific physical models, including particle tracking and radiation ray tracing, yielding significant speedup.

Improved scaling continues in Fluent 14.0, in which scaling using a pre-release version has been demonstrated out to nearly 4,000 cores for problem sizes on the order of 100M cells. Of course, such extreme scaling is important for very large problems, but it also demonstrates the ability to scale simulation models using partition sizes as small as 10,000 cells per core — or even less. This means that smaller models can be accelerated using more cores than ever before. Fluent also incorporates network-aware mapping of partitions: Those requiring high levels of communication are logically placed to minimize network traffic. With the adjacent mapping technique, communication overhead is minimized, leading to faster overall solution time, especially on slower networks.

In ANSYS HFSS 14.0, the distributed memory parallel technique, domain decomposition, extends to efficient electromagnetic analysis of structures with highly repetitive geometries, such as antenna arrays. A 256-element waveguide array has been analyzed in under two hours with less than 1GB of total memory across 16 compute engines. This analysis provides a comprehensive solution to the finite array, including all coupling matrices and edge effects.

Performance gains have been realized throughout release 14.0 via adoption of the latest Intel® compilers and math kernel libraries. For mechanical analysis, this has

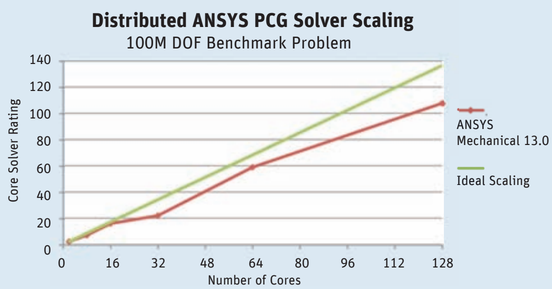

For structural analysis, the Distributed ANSYS solver offers near-linear scaling for the solution of large problems. Problem sizes as large as 100 million degrees of freedom can be analyzed quickly and cost-effectively.

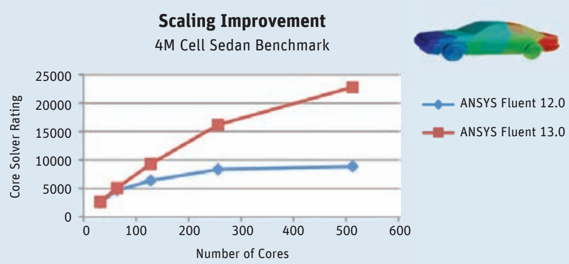

Hybrid parallelism in ANSYS Fluent yields impressive performance improvements — with scaling out to 512 cores for this 4-million-cell simulation of airflow around a car.

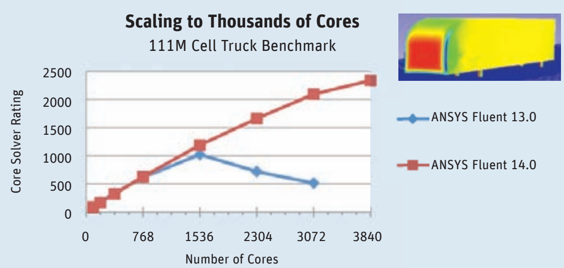

Ongoing software improvements enable scaling out to thousands of cores — and enable smaller problems to execute effectively on ever-larger core counts.

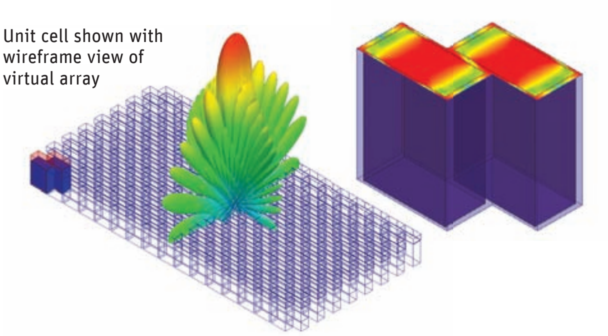

Distributed memory parallel processing enabled efficient electromagnetic simulation of this 256-element skewed waveguide array using ANSYS HFSS.

yielded up to 40 percent improvement in the speed of the sparse direct solver. These updates also mean that ANSYS 14.0 can take advantage of new AVX compiler instructions, which is expected to yield as much as 50 percent speedup when using the sparse direct solver on the latest processors from Intel and AMD, which support this new instruction set.

GPU Computing

GPUs are an exciting new technology for HPC, as they provide hundreds of processing cores that are capable of computational throughput far in excess of that achieved on today's multicore CPUs. For software developers, GPU computing represents a significant challenge, as the algorithms adapted to traditional CPUs must be reconsidered to effectively use GPU compute capacity while working within the constraint of the relatively limited memory available to the GPU. ANSYS 13.0 introduced GPU use to accelerate mechanical simulations using the shared memory solver [1]. Further innovations extend the use of GPU computing in ANSYS 14.0.

Specifically for mechanical simulations, ANSYS 14.0 introduces the ability to apply GPU acceleration to the distributed memory solvers — and to use multiple GPUs that reside in multiple machines within a cluster of computers. Depending on the workload, speed improvements of two times are achievable using a single GPU and continue to be significant when multiple GPUs are used. For fluids simulations, core solver acceleration is an ongoing focus of collaboration between ANSYS and the leading providers of GPU technology. Research on GPUs (which resulted in a beta feature in release 14.0) has led to substantial acceleration of specific fluid simulation submodels.

Adding Strategic Value to Engineering Simulation

ANSYS software development focuses on HPC performance numbers, providing the best possible return on a company's overall investment in HPC infrastructure. However, in the end, the value of operating in an HPC environment is rooted in what this speed enables: Running simulations at a higher level of fidelity and considering more design ideas translates into a long-term competitive edge. Using HPC, organizations can launch new products and design features more rapidly, at a lower financial investment and with higher confidence in ultimate product performance. ANSYS delivers best-in-class HPC performance to ensure the ability to maximize this strategic value. ■

[1] Beisheim, J. Speed Up Simulations with a GPU. *ANSYS Advantage*, **2010**, V4, I2, web exclusive.

Simulations and statistics for ANSYS 14.0 reported in this article were performed with a pre-release version of the software.

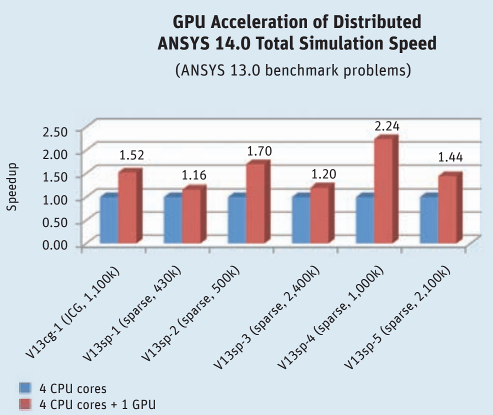

GPU computing has been extended to support distributed memory structural solutions in ANSYS 14.0 software, yielding single-machine speed increases of up to two times over a range of benchmark problems.

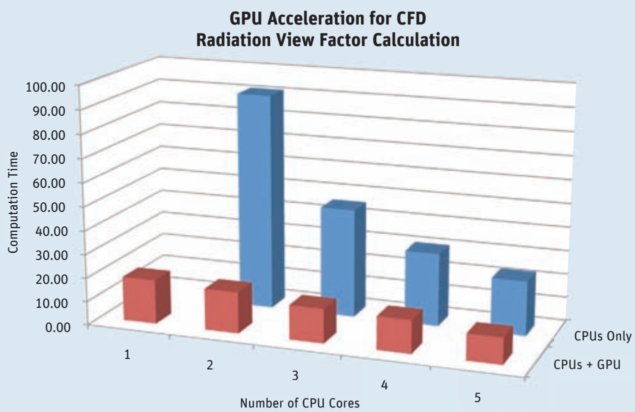

GPU computing applied to radiation heat transfer calculations in ANSYS Fluent 14.0 (using beta capabilities) demonstrates great potential to accelerate physical submodels.