

The Need for Speed

From desktop to supercomputer, high-performance computing with ANSYS 12.0 continues to race ahead.

Tuning software from ANSYS on the latest high-performance computing technologies for optimal performance has been — and will continue to be — a major focus area within the software development organization at ANSYS. This effort has yielded significant performance gains and new functionality in ANSYS 12.0, with important implications for more productive use of simulation by customers.

High-performance computing, or HPC, refers to the use of high-speed processors (CPUs) and related technologies to solve computationally intensive problems. In recent years, HPC has become much more widely available and affordable, primarily due to the use of multiple low-cost processors that work in parallel on the computational task. Today, clusters of affordable compute servers make large-scale parallel processing a very viable strategy for ANSYS customers. In fact, the new multi-core processors have turned even desktop workstations into high-performance platforms for single-job execution.

This wider availability of HPC systems is enabling important trends in engineering simulation. Simulation models are getting *larger* — using more computer memory and requiring more computational time — as engineers include greater geometric detail and more-realistic treatment of physical phenomena (Figure 1). These higher-fidelity models are critical for simulation to reduce the need for expensive physical testing. HPC systems make higher-fidelity simulations practical by yielding results within the engineering project's required time frame. A second important trend is toward *more* simulations — enabling engineers to consider multiple design ideas, conduct parametric studies and even perform automated design optimization. HPC systems provide the throughput required for completing multiple simulations simultaneously, thus allowing design decisions to be made early in the project.

Software from ANSYS takes advantage of multi-processor and/or multi-core systems by employing domain decomposition, which divides the simulation model into multiple pieces or sub-domains. Each sub-domain is then computed on a separate processor (or core), and the multiple processors work in parallel to speed up the computation. In the ideal case, speedup is linear, meaning that the simulation turnaround time can be reduced in proportion to the number of processors used. Parallel processing also allows larger problems to be tackled, since the processing power and memory requirements can be distributed across the cluster of processors. Whether performed on a multi-core desktop workstation, desk-side cluster or scaled-out HPC system, parallel

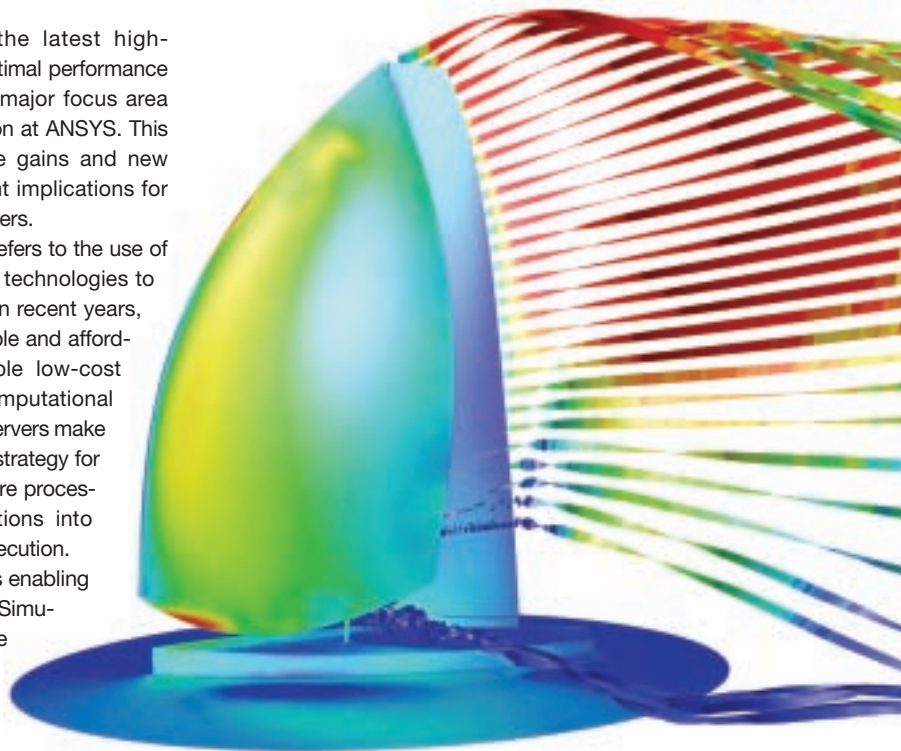


Figure 1. Simulations as large as 1 billion cells are now supported at release 12.0. This 1 billion-scale racing yacht simulation was conducted on a cluster of 208 HP ProLiant™ server blades. (For more information, visit www.ansys.com/one-billion.)
Image courtesy Ignazio Maria Viola.

HPC on Workstations?

While purists might argue whether workstations can be considered high-performance computing platforms, the performance possibilities for ANSYS 12.0 running on workstations are noteworthy. With the latest quad-core processor technology, an eight-core workstation running Windows® can deliver a speedup of five to six times for users of mechanical products from ANSYS (Figure 2) and over seven times for users of its fluid dynamics products (Figure 4). This means that parallel processing now provides tremendous ROI for both large engineering groups and individual workstation users, enabling faster turnaround, higher-fidelity models and parametric modeling. With release 12.0 and 2009 computing platforms, parallel processing improves productivity for all simulation types, from workstation to cluster, for mechanical or fluids simulations.

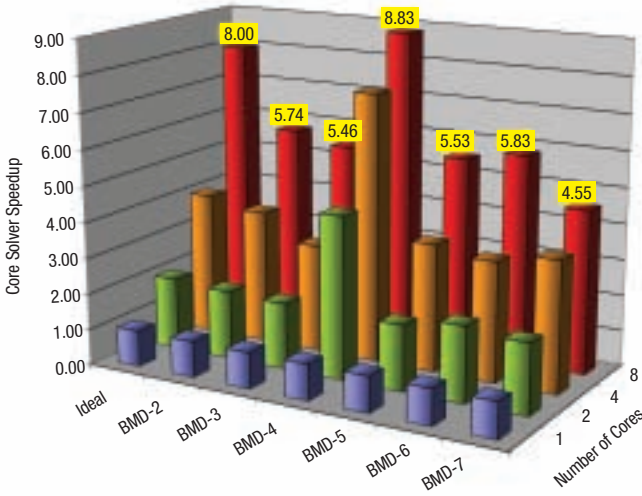


Figure 2. Speedup of Distributed ANSYS Mechanical 12.0 software using the 11.0 SP1 benchmark problems. Simulations running eight-way parallel show typical speedup of between five and six times. Data was collected on a Cray CX-1 Personal Supercomputer using two quad-core Intel Xeon Processor E5472 running Microsoft® Windows HPC Server 2008.

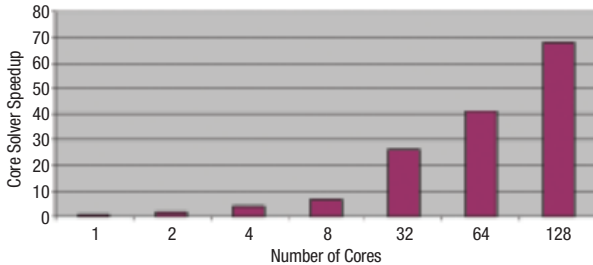


Figure 3. Scaling of a 10M DOF simulation using the ANSYS Mechanical 12.0 iterative PCG solver on a cluster of Intel Xeon 5500 Processor series. All cores on these quad-core processors are fully utilized for the benchmark.

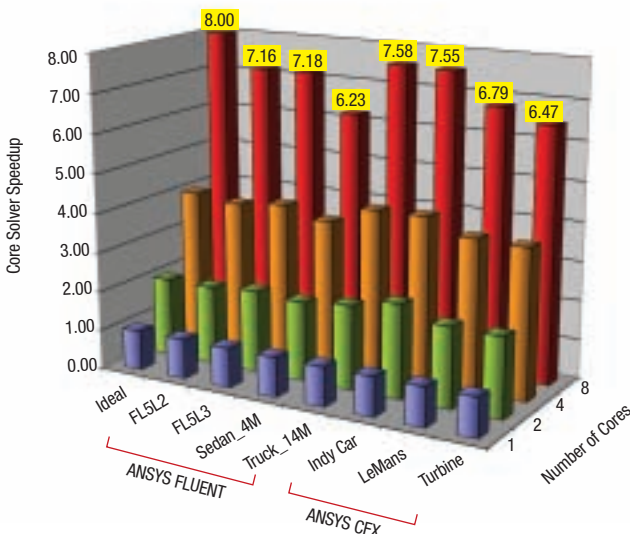


Figure 4. Scalability of ANSYS FLUENT and ANSYS CFX benchmark problems on the Intel Xeon 5500 Processor series quad-core platform. Simulations running eight-way parallel show typical speedup of over seven times.

processing provides excellent return on investment by improving the productivity of the engineers who perform simulation.

ANSYS 12.0 provides many important advances in areas related to parallel processing and HPC, delivering scalability from desktop systems to supercomputers. For users of the ANSYS Mechanical product line, release 12.0 introduces expanded functionality in the Distributed ANSYS (DANSYS) solvers, including support for all multi-field simulations, prestress effects and analyses involving cyclic symmetry. In addition, DANSYS now supports both symmetric and non-symmetric matrices as well as all electromagnetic analyses. Mechanical simulations benefit from significantly improved scaling on the latest multi-core processors. Simulations in the size range of 2 million to 3 million degrees of freedom (DOF) now show good scaling on eight cores (Figure 2). Based on benchmark problem performance, customers can expect to get answers back five to six times faster on eight cores. Even more impressive is the scale-out behavior shown in Figure 3, with a 10 million DOF simulation showing solver speedup of 68 times on 128 cores.

With turnaround times measured in tens of seconds, parametric studies and automated design optimization are now well within the grasp of ANSYS customers who perform mechanical simulations. These benchmarks are noteworthy, in part, as they show execution with all cores on the cluster fully utilized, indicating that the latest quad-core processors have sufficient memory bandwidth to support parallel processing for memory-hungry mechanical simulations. Software tuning has contributed to improved scaling as well, including improved domain decomposition, load balancing and distributed matrix generation. To help customers maximize their ANSYS solver performance, the online help system now includes a performance guide that provides a comprehensive summary of factors that impact the performance of mechanical simulations on current hardware systems.

Explicit simulations using ANSYS AUTODYN technology take great advantage of HPC systems at release 12.0. Full 64-bit support is now available, allowing much larger simulations to be considered from pre-processing to solution and post-processing.

For users of fluid dynamics software from ANSYS, release 12.0 builds on the strong foundation of excellent scaling in both the ANSYS FLUENT and ANSYS CFX solvers. These fluids simulation codes run massively parallel, with sustained scaling at hundreds or even thousands of cores. The release incorporates tuning for the latest multi-core processors, including enhanced cache re-utilization, optimal mapping and binding of processes to cores (for better memory locality and system utilization), and leveraging the latest compiler optimizations. The resulting ANSYS FLUENT and ANSYS CFX performance on the newly released Intel® Xeon® 5500 Processor series is shown in Figure 4, with outstanding speedup of over seven times for many benchmark cases. In addition, the new release delivers significant performance improvements at large core counts, the result of general solver enhancements and optimized communications over the latest high-speed interconnects. Figure 5 demonstrates

scaling achieved by ANSYS CFX software on a cluster of quad-core AMD processors. Nearly ideal linear scaling to 1,024 cores — and very good efficiency up to 2,048 cores — has been demonstrated with ANSYS FLUENT (Figure 6). Both fluids codes provide improvements to mesh partitioning that enhance scalability. ANSYS FLUENT software now provides dynamic load balancing based on mesh- and solution-derived criteria. This enables optimal scalability for simulations involving multiphysics, such as particle-laden flows. The ANSYS CFX code delivers improved partitioning for moving and/or rotating meshes, yielding important reductions in memory use and improved performance for turbomachinery and related applications. Finally, ANSYS FLUENT users will benefit from several usability improvements, including built-in tools for checking system network bandwidth, latency and resource utilization — all helping to identify potential scaling bottlenecks on the cluster.

Beyond solver speedup, the ANSYS 12.0 focus on HPC addresses issues related to file input and output (I/O). Both ANSYS FLUENT and ANSYS CFX software have updated I/O algorithms to speed up writing of results files on clusters, enhancing the practicality of periodic solution snapshots when checkpointing or running time-dependent simulations. ANSYS FLUENT includes improvements in the standard file I/O as well as new support for fully parallel I/O based on parallel file systems. Order of magnitude improvements in I/O throughput have been demonstrated on large test cases (Figure 7), virtually eliminating I/O as a potential bottleneck for large-scale simulations. ANSYS CFX improves I/O performance via data compression during the process of gathering from the cluster nodes, therefore reducing file write times. Proper I/O configuration is also an important aspect of cluster performance for the ANSYS Mechanical product line.

Recognizing that cluster deployment and management are key concerns, ANSYS 12.0 includes a focus on compatibility with the overall HPC ecosystem. ANSYS products are registered and tested as part of the Intel Cluster Ready program, confirming that these products conform to standards of compatibility that contribute to successful deployment (www.ansys.com/intelclusterready). In addition to supporting enterprise Linux® distributions from Red Hat® and Novell, ANSYS 12.0 products are supported on clusters based on Microsoft Windows HPC Server 2008. ANSYS has also worked with hardware OEMs, including HP®, SGI®, IBM®, Dell®, Cray® and others, to define reference configurations that are optimally designed to run simulation software from ANSYS (www.ansys.com/reference-configs).

As computing technology continues to evolve, ANSYS is working with HPC leaders to ensure support for the breakthrough capability that will make simulation more productive. Looking forward, important emerging technologies include many-core processors, general purpose graphical processing units (GP-GPUs) and fault tolerance at large scale. ■

Contributions to this article were made by Barbara Hutchings, Ray Browell and Prasad Alavilli of ANSYS, Inc.

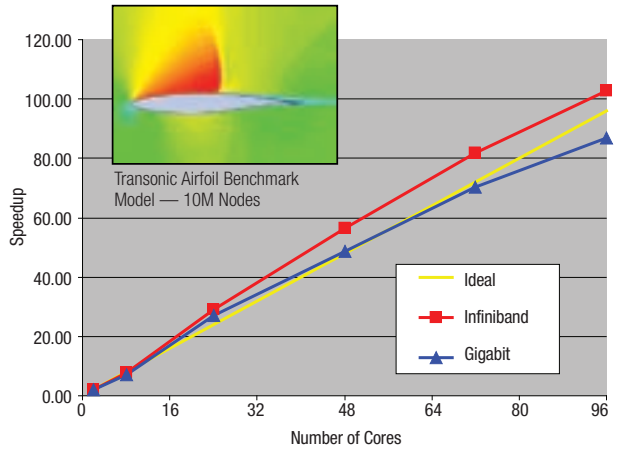


Figure 5. Scalability of ANSYS CFX 12.0 on a 10M node transonic airfoil benchmark example. Data was collected on a cluster of AMD Opteron™ 2218 processors, showing the benefit of a high-speed interconnect.

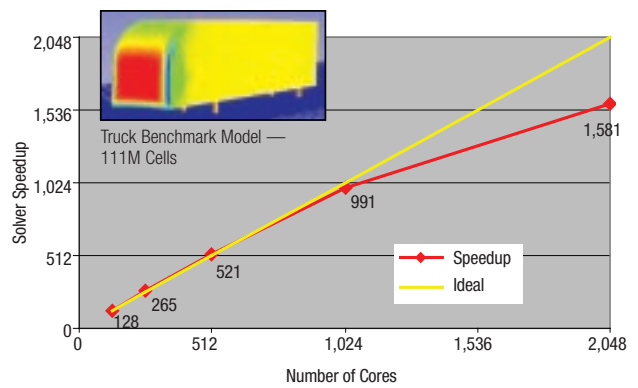


Figure 6. Scaling of ANSYS FLUENT 12.0 software is nearly ideal up to 1,024 processors and 78 percent of ideal at 2,048 processors. Data courtesy SGI, based on the SGI Altix® ICE 8200EX using quad-core Intel Xeon Processor E5472 with Infiniband®.

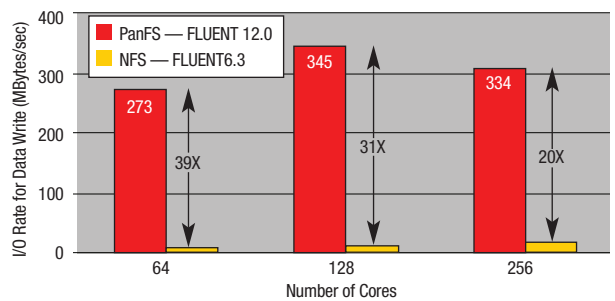


Figure 7. Parallel I/O in ANSYS FLUENT 12.0 using the Panasas® file system, compared to serial I/O in the previous release using NFS. Parallel treatment of I/O provides important speedup for time-varying simulations on large clusters.