Solutions Brief

Solving 111 Million Degrees of Freedom

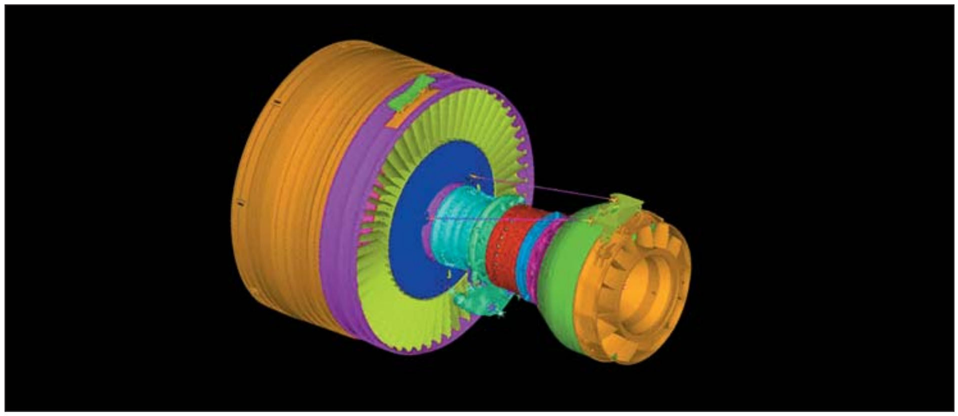

Full scale CAE simulation of complex detailed assembly; courtesy of ANSYS, Inc.

The ANSYS Big Memory Breakthrough on the Itanium® 2-Based SGI® Altix® High-Performance Computing Platform

For engineers who work with CAE models of complex, detailed assemblies, productivity has been limited by computing architectures that cannot hold and process large models in a timely way. The greatest limitation is memory size. Processing times have slowed to a crawl while data is passed to and from disk storage arrays because it has been impossible to bring complete large models into memory or to solve the large mathematical systems generated by larger models.

ANSYS, Intel, and SGI recently collaborated in a ground-breaking project that provides spectacular acceleration of ANSYS CAE model processing by combining major advancements in computing platform architecture, processor architecture, and ANSYS code. This combination effectively removes the model size and processing time limitations for current engineering design simulations and extends the capabilities of future more detailed and more sophisticated analyses.

At the International ANSYS Conference in May 2004, ANSYS used an SGI Altix server with six 64-bit Intel® Itanium® 2 processors and a single shared memory of 64 gigabytes to solve a structural analysis problem with 111 million degrees of freedom (MDOF) in just 8.6 hours of solver time. This demonstration made it clear that it is now possible to solve detailed models of aircraft engines, automobiles, construction equipment and other complete systems in hours instead of days, reducing design cycles from months to weeks. In the simplest terms, ANSYS customers are now enabled to build more detailed models and solve them in less time.

Breaking Through the Memory Barrier

The complex models being built today have outgrown commonly available compute technology. CAE users have been unable to exploit the great increases in processor speed because most computing systems have remained memory-limited. Systems based on 32-bit architectures have limited ANSYS users to 2GB of memory, and ANSYS code requires 1GB of memory per MDOF. This constraint has limited simulations to 2MDOF.

The NUMAlink™ architecture of the Itanium 2 processor-based SGI Altix platform has removed this barrier. The amount of memory that can be accessed for a given problem is no longer tied to the number of processors and is essentially unlimited. The NUMAlink fabric enables system memory capacity of

hundreds of gigabytes or even terabytes of memory per processor. SGI has already installed several systems with one, two, and even four terabytes of shared global memory and will have a system with a single 13-terabyte memory installed by year-end 2004.

The other major enabler for the ANSYS breakthrough is parallelism at all levels – in the SGI Altix platform architecture, in the architecture of the Intel Itanium 2 processor, and in the ANSYS code itself.

“ANSYS is very advanced compared to its competitors in terms of having a very parallel architecture,” says Principal Engineer Mike Greenfield of Intel. “They've done a very good job of using parallel processing across the entire application – not just the solver.” In 2002, ANSYS responded to the introduction of the Altix platform by initiating its big memory breakthrough project. ANSYS engineers first modified the ANSYS application code by removing 32-bit memory addressing limitations to exploit the 64-bit address space of the Itanium 2 processor. Even though ANSYS code has been supported on several 64-bit platforms for some time, the capability to address large data arrays exceeding 2 billion elements was needed to allow ANSYS code to practically scale to solve huge problems. These changes are included in the ANSYS 8.1 release that was announced at the May 2004 user conference.

“When you think of moving from 32-bit to 64-bit, it sounds like you're doubling memory access,” says Technical Fellow Gene Poole of ANSYS. "The truth is that you're increasing it by a factor of 2^{32}. The maximum memory addressable is increased by a factor of 2 billion – a larger number than anyone can afford to buy memory for. It solves memory access limitations in our program for the foreseeable future."

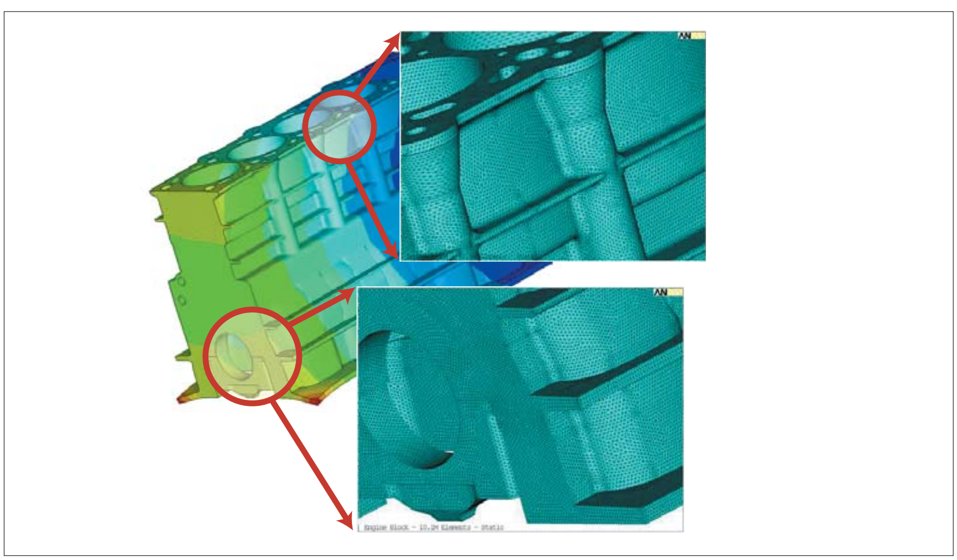

43.9 million DOFs with ANSYS SOLID92, Static Analysis; courtesy of ANSYS, Inc.

The other major advancement demonstrated by ANSYS at the conference was a new parallel solver designed to run using distributed memory parallel processing. “Only the solver ran in distributed mode in this first large memory breakthrough,” says Poole. “But solver time for a job of that magnitude is such a large proportion of processing time that we were able to reduce the total solution time for 111 MDOF to about eight hours.”

The new parallel processing improvements coupled with large memory gives ANSYS users the potential to process the largest models in core, eliminating the terabytes of I/O transfers that have stretched delivery time for large-model results to days and weeks. ANSYS is now developing a distributed version of its code, to be included in the 9.0 release, which will distribute the entire solution process. This highly parallel code architecture will improve scalability of the entire process and reduce time-to-solution.

The Intel® Itanium® 2 Microarchitecture

The Intel Itanium 2 processor gives the Altix platform superior floating point performance. Its Explicitly Parallel Instruction Computation (EPIC) architecture delivers up to four floating point operations per cycle and execute ANSYS instructions at an unprecedented level of parallelism and efficiency.

“It is one of the fastest processors available today at any price, and includes the Intel® Math Kernel Library (Intel® MKL), which does the matrix operations we need for our direct solver at near the theoretical peak performance of the architecture,” says Poole. “For those reasons it runs faster than anything else.”

“The ANSYS application is a very important one, and Intel continues to work with ANSYS engineers to ensure that it is highly optimized for the Itanium 2 architecture," says Greenfield. "We are constantly looking for opportunities to increase its performance on Itanium 2. This focus lets us make the application run much faster than it would unaided. The compiler, for example, gives us many permutations for building the code. We search constantly for the best combinations of choices to achieve peak ANSYS performance."

The explicitly parallel architecture of the Itanium 2 processor, including its massive execution resources, gives it the ability to execute multiple instructions simultaneously, enabling more users to access compute resources. Calculations and data analysis are accelerated, and large CAD and CAE models can be moved and stored more quickly. The parallel architecture also provides faster simulation and rendering times for automotive designs. The Itanium processor's internal resources include a massive on-die cache, 128 floating-point registers, and memory management features that enable more efficient management of large data sets.

The enterprise-hardened design of the Itanium 2 processor includes features that provide high reliability, availability, serviceability, and manageability. The Advanced Machine Check Architecture provides extensive error management in hardware, firmware, and operating systems to minimize data loss, data corruption, and downtime.

Intel's commitment to this flagship server processor is demonstrated by its strong roadmap. Multiple Intel design teams currently have six Itanium processors in simultaneous development with the goal of increasingly higher performance at reduced cost. By 2007, Intel plans to ship Itanium 2 processors with twice the performance levels of the Intel® Xeon™ processor line at the same price.

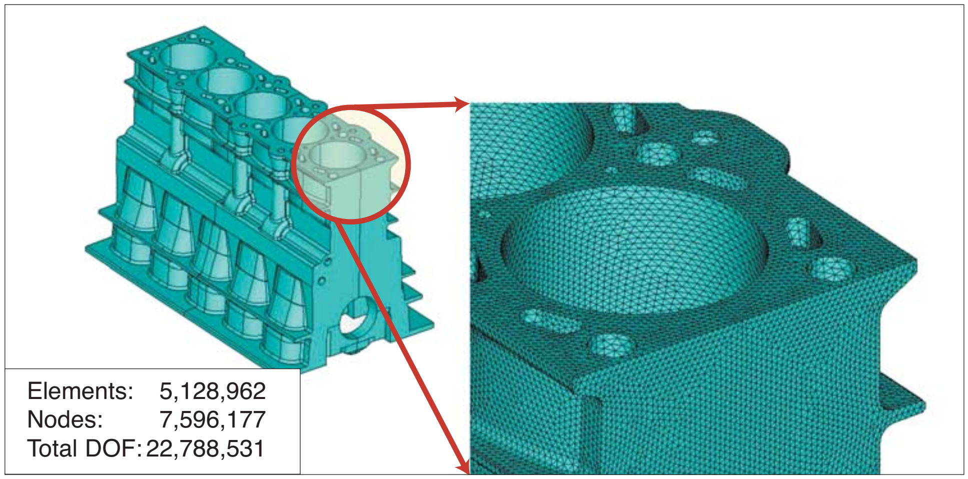

Parallel computing results: memory reduction; courtesy of ANSYS, Inc.

Big Memory Benefits for ANSYS Users

Now that ANSYS customers can use the virtually unlimited memory space available through the Altix architecture, they can focus on modeling complete assemblies using detailed models – the largest models ever attempted – to achieve breakthrough analysis results. But even with models of relatively modest size, big memory can make an enormous difference in throughput.

As an example, consider vibration-frequencies analyses of the turbine blades used in jet engines – analyses that require an ANSYS user to solve a model with several million DOF. These analyses require solution of mathematical systems of equations that require tens of gigabytes of storage. In the past these large systems would be stored on disk and accessed dozens of times during the analyses. Using the new large-memory capability of the ANSYS code on large-memory Altix systems to store the entire system of equations in memory can save terabytes of I/O and generate significant operational advantages:

Breakthrough performance levels. Running the entire model in memory for an I/O-dominated analysis can reduce solution time by a factor of 100. It has been demonstrated that ANSYS code running on the SGI Altix platform can solve a 100 MDOF model in a single work shift. "ANSYS code can also now run large modal (vibration) analysis with the entire mathematical system of equations in memory to dramatically increase performance," says Poole. "That is a unique capability."

No need for model size reduction. Companies spend many costly man-hours reducing the size of large models such as complete engines in order to run them in memory-limited environments. The ANSYS memory breakthrough enables them to run complete models with millions of degrees of freedom in very short run times, eliminating the delays, costs and potential errors of model reduction to deliver results of higher quality.

A solution to the I/O bottleneck. The challenge with running many large problems in memory-limited environments is constant I/O transfer to and from disk, which may amount to terabytes of movement for a typical analysis. Eliminating I/O delays dramatically reduces overall solution times. Big, complex problems can be run in hours.

More detailed, realistic solutions. Without memory limitations, analyses can be run at meshes exceeding 100 MDOF to yield highly detailed simulations that provide more accurate, realistic results.

ANSYS Code on the SGI® Altix® 350 Server

The Linux-based mid-range SGI Altix 350 system provides an ideal platform for ANSYS users who want to move up to big memory productivity. While 64-bit workstations and mainframe computers have been available for many years, main memory sizes of 100 GB or more were neither available nor affordable. Now the Altix expand-on-demand architecture enables users to independently scale processors, memory, or I/O to meet changing demands.

The Altix 350 server scales from 2 to 16 processors in a single system image, which covers the 8-16 processor range favored by ANSYS users. It includes the industry-leading 6.4GB/second NUMAlink interconnect. The Altix 350 system delivers its breakthrough performance at a price point of around $100,000, a modest investment compared to high-end workstations or memory-limited HPC systems that cannot provide the dramatic throughput of the Altix platform. It makes high-end computing a very affordable option for many companies.

ANSYS users typically require a gigabyte of memory per MDOF. On the Altix system they can scale to 192GB of memory to achieve the degree of detail and realism they need, or to increase the size of their models to include complete assemblies or products. ANSYS code has been performance-optimized for both the Altix and Itanium 2 architectures.

To Sum Up: A Breakthrough for ANSYS Users

ANSYS has become the first simulation company to break through the memory barrier and deliver analyses of over 100 million DOF on the SGI Altix platform. It can accommodate the largest models ever attempted, and provide results with much greater detail and realism in a dramatically shorter solution time.

SGI and Intel have collaborated to create the Altix platform that makes this breakthrough possible at an attractive price point. “Once ANSYS users with large models experience the performance and ease-of-use benefits of a large shared memory, it will be like moving up from dial-up internet service to broadband,” says Altix Product Manager Derek Robb of SGI. “They won't ever want to go back.”

“The enthusiastic response of our users to this breakthrough demonstration exceeded our expectations and confirmed our belief that this collaboration to break through the memory limitations of 32-bit computing would make a difference in the CAE community,” says Poole. “It was a highly successful demonstration of what our customers can do with a modest investment in hardware.”

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.925.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3691 [09.28.2004]
J14681