SGI and ESI Group Team to Deliver Scalable Crash Simulation Solutions

PAM-CRASH® scales to 512 processors on SGI® Altix® for unparalleled performance

Table of Contents

Driving R&D Productivity through Technology

Recent technological and industrial advancements have resulted in an integrated global environment where research and development (R&D), manufacturing, distribution, finance, and sales rely on technology to maximize efficiency and productivity. This tightly integrated global value chain offers significant advantages to manufacturers such as the ability to reach an enormous customer base, and product development infrastructures which allow management teams to quickly adjust to changing market requirements and regional capabilities. Despite recent technology-driven gains, however, R&D teams remain under great pressure to maximize efficiency and reduce time to market even further.

To this end, R&D groups are increasingly turning to virtual simulations to supplement or replace the lengthy process of creating and testing real prototypes. Automotive companies and other manufacturers in the transportation industry now rely on crashworthiness simulation software as an essential element of crash and safety testing. To meet this need, high performance computing (HPC) capacity in automotive companies is growing by a factor of 10 every 4 years. At the end of 2004, automotive companies had over 90 teraflops of compute server capacity, more than 55 percent of which is dedicated solely to crash simulation.

Many organizations have deployed large clusters—built from off-the-shelf computers connected by high-performance networks—for this work. While this strategy has proven successful for some workloads, communication between cluster nodes and limited memory capacity hampers the effectiveness of clusters on many problems.

As a result, leading manufacturers have now begun to realize that they need a comprehensive computing solution in which appropriate microprocessors are combined with a high performance system infrastructure and tuned software to further optimize the

product development process. By teaming with ESI Group, a leading provider of digital simulation software for prototyping and manufacturing processes, SGI is able to deliver a complete solution for the crash simulation marketplace capable of achieving unprecedented levels of performance to reduce the time needed for critical simulations.

SGI Altix, Intel® Itanium® 2 processors, and PAM-CRASH: A Winning Combination for Performance, Scalability and Reliability

The SGI Altix architecture is designed to scale in ways that commodity clusters simply cannot, without sacrificing the benefits that customers expect from a standard Linux environment The Altix 3700 Bx2 system, which utilizes the SGI NUMAflex® architecture and Intel® Itanium® 2 processors, can accommodate up to 64 processors in a standard Altix rack and up to 512 processors and 6TB of memory in a single system. NUMAflex provides global shared memory with extremely high bandwidth and very low latency between processors. All resources in an Altix node are managed through a single operating system instance, the industry's largest single system image. Up to four Altix nodes can be joined with NUMAlink to offer 2048 processors and 24TB of globally addressable memory for even greater capability. The SGI NUMAlink 4 interconnect creates a high bandwidth(12GB/s), ultra-low latency (<1microsecond) fabric to speed interprocessor communication and shared memory address, even between nodes.

As a result the Altix 3700 delivers:

- The fastest time-to-solution for capability applications and best job throughput for mixed capacity workloads.
- Flexibility to accommodate a wide range of HPC workloads.
- Superior ease of use in terms of application development, performance optimization and system administration.

SGI® Altix® dramatically reduces the time and resources required to run technical applications by accommodating extremely large data sets in a single, system-wide, shared-memory space. By holding more complex job geometries and entire workflows in memory, SGI Altix can significantly decrease the time needed to complete complex crash simulations or related work.

Global shared memory allows direct and efficient access to all data in the system's memory, without having to move data

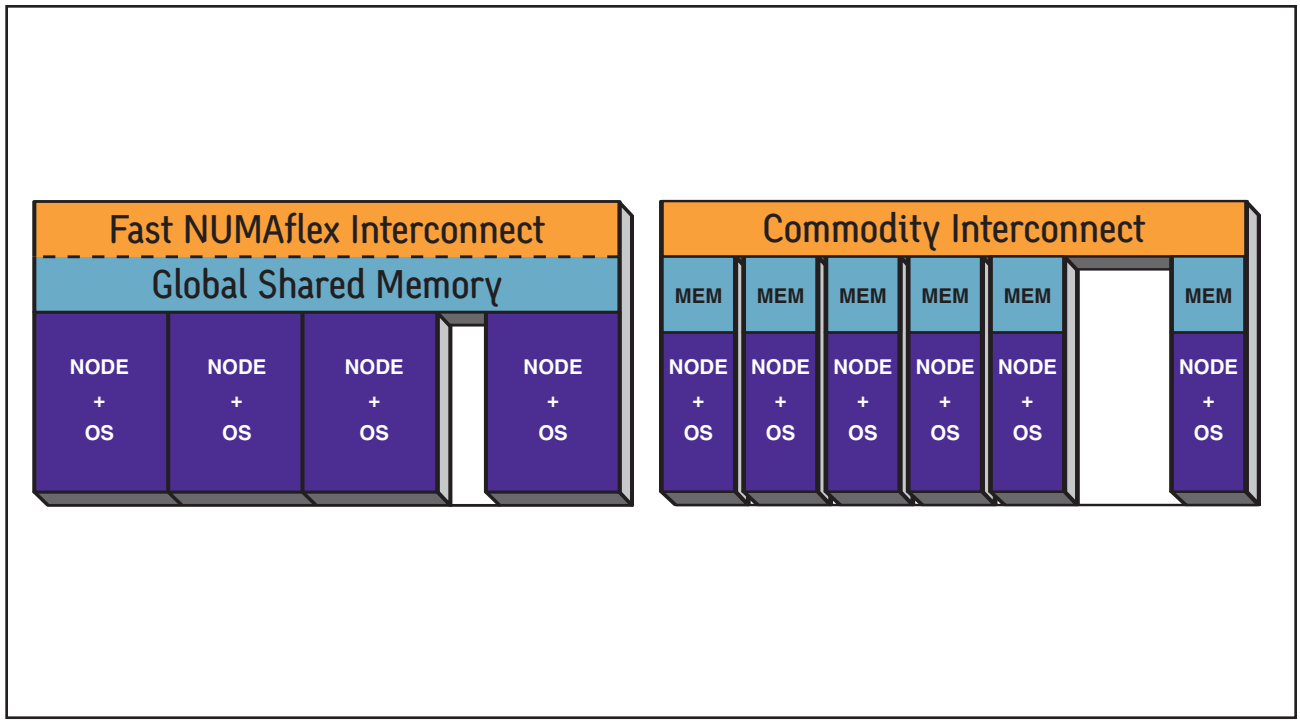

through I/O or networking bottlenecks. Clusters that lack global shared memory must instead pass copies of data back and forth, typically in the form of messages. This can greatly complicate programming and impede performance by increasing the time processors spend waiting for data.

Automotive companies need more than an economical computing platform. They need a platform that increases the productivity of engineers and the critical applications they rely on to do their work. The global shared memory and high-speed communication architecture of the Altix platform was designed with that goal in mind.

Intel® Itanium® 2 Processors

A key element of the Altix architecture is the Intel Itanium 2 processor, which gives the Altix platform superior floating point performance. By choosing Itanium 2 processors for the Altix architecture, SGI is well positioned to not only take advantage of the outstanding performance today, but to protect customer investments through the ability to upgrade over time according to the Itanium roadmap including upcoming multicore designs.

The Itanium 2 processor is based on the Explicitly Parallel Instruction Computing or EPIC architecture. Since the instruction set is explicitly parallel, the architecture attains high levels of parallelism inside the processor (so called "instruction level parallelism"). The Intel Itanium 2 compiler takes directives from the programmer to create highly parallelized machine code that takes the fullest advantage of the processor's multiple execution units and massive on-dye resources. The Itanium 2 processor is able to process 6 instructions per clock cycle versus 3 instructions per clock cycle for the Intel Xeon chip common in

"The Itanium 2 processor's EPIC architecture is perfect for large crash simulations. The Altix interconnect lets applications scale up to large numbers of Itanium 2 processors. At the same time, programming is dramatically simplified due to the large shared address space provided by SGI's NUMAflex architecture."

Tim Mattson Ph.D., Parallel Computing Evangelist, Intel.

commodity clusters. Itanium 2 processor actually significantly outperforms Xeon chips running at twice the clock rate.

"The Itanium 2 processor's EPIC architecture is perfect for large crash simulations" says Tim Mattson Ph.D., Parallel Computing Evangelist at Intel. "The Altix interconnect lets applications scale up to large numbers of Itanium 2 processors. At the same time, programming is dramatically simplified due to the large shared address space provided by SGI's NUMAflex architecture. This significantly lowers the barrier for institutional programmers and ISVs that want to port to the platform."

PAM-CRASH Software

PAM-CRASH is ESI Group's leading crash simulation software package. It enables engineers to focus their efforts on creating, evaluating and managing virtual prototypes under realistic conditions. The PAM-CRASH 2G environment facilitates model preparation and analysis while its solver architecture comes in a highly-scalable, parallel version. This parallel version enables overnight computation and intensive calculation even with the increased model size that can result from more detailed car body representations including trims and precise physical properties.

The parallel solver uses a distributed-memory parallel processing (DMP) model with increased processing efficiency as the primary performance objective. Good efficiency has been demonstrated up to 512 processors. Some options, such as contact interfaces, have been significantly reworked to further improve scalability for complex models.

"The new generation of our parallel solver, first delivered in 2002, offers unprecedented scalability for full car model computations. Crash simulation models that now require several millions of elements to take into consideration structural details and rupture issues, can be run overnight with PAM-CRASH."

Vincent Chaillou, President & COO Product Operations, ESI Group.

"The new generation of our parallel solver, first delivered in 2002, offers unprecedented scalability for full car model computations. Crash simulation models that now require several millions of elements to take into consideration structural details and rupture issues, can be run overnight with PAM-CRASH," says Vincent Chaillou President & COO Product Operations ESI Group.

These scalability improvements allow PAM-CRASH users to:

- Evaluate several variants in a single day
- Use automated optimization algorithms or perform design robustness checks in the short time required by industrial projects.
- Increase model size and/or refine the physics and still complete runs overnight.

Record-setting Performance on PAM-CRASH

SGI and ESI Group have a long history of pushing the performance envelope in crash simulation, and have collaborated to optimize the performance and scalability of PAM-CRASH solutions, starting with the initial NUMAflex implementation of the SGI Origin line. The partnership has now evolved to encompass the Linux-based SGI Altix to achieve a new level of performance from a solution that adheres to industry standards.

"The Altix server family offers great deployment flexibility. Altix systems can be tailored specifically to meet application needs in terms of compute, memory capacity, or I/O so you never have to pay for resources you don't need," says Christian Tanasescu, SGI Director of Engineering, Simulation Applications. "With global shared memory, no data transfers are required between nodes. Altix is the first system architecture in the industry to break the 1 microsecond barrier for MPI communication. This makes it possible to scale crash simulation applications such as PAM-CRASH to a very high number of processors. Because PAM-CRASH exchanges a large number of small messages, communication latency is crucial to performance."

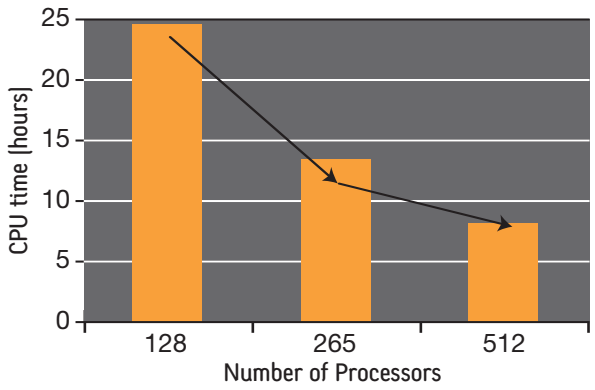

CPU times for the Neon car-to-car crash model

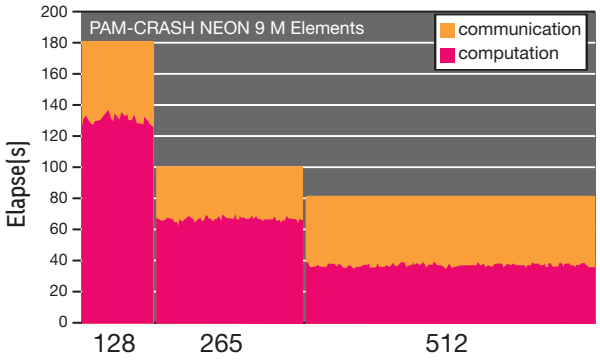

MPI profile of a sample of representative cycles

Case study: 9.1 million Element Neon Offset Crash

The performance of PAM-CRASH running on the Altix platform was tested using a 9.1 million element offset crash simulation derived from a National Crash Analysis Center (NCAC) public domain model of the Dodge Neon. The test simulates a car-to-car 50% offset crash with a single full-model contact interface; each car has an initial speed of 56km/h. The computation runs for 363,000 cycles, simulating a 120 ms crash event. An Altix system with 512 Itanium 2 CPUs (1.6GHz, 6 MB cache) and 512GB of RAM was used to execute all tests.

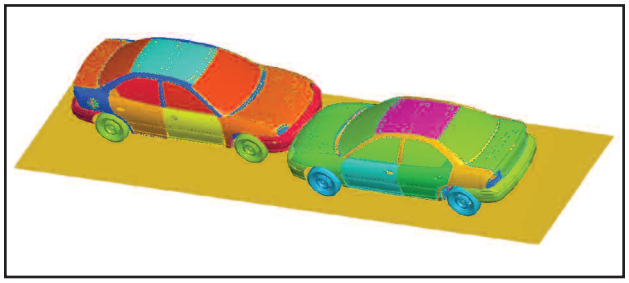

Neon Car-to-Car 50% offset crash using PAM-CRASH 2004 DMP

To determine the scalability of PAM-CRASH on Altix, the same run was performed using 128, 256 and 512 processors. A complete run on 512 processors took only 29,450s (8h11') on 512 processors, within the critical time range for overnight results. This run corresponded to a performance level of 338.20 Gflops, the highest ever achieved for PAM-CRASH.

The observed speed-up values were 1.83 from 128 to 256 processors and 1.64 from 256 to 512 processors. (Since the number of processors doubles in each case, a value of 2.0 would indicate perfect scaling.) These speed-ups are particularly impressive given that they include initialization time which is inherently proportional to the number of CPUs. A special optimization effort was undertaken by ESI Group and SGI to reduce initialization time, which might otherwise have jeopardized observed scalability gains. This optimization will be included in PAM-CRASH 2005, to ensure that commercial users can achieve the same level of scalability.

The chart above demonstrates the importance of communication latency in high CPU count PAM-CRASH computations. Load balance improves as the number of processors is increased, and communication scales with the computation. The 1 microsecond access time of Altix is the reason why PAM-CRASH scales so well up to 512 processors. This capability is unmatched in the market.

Conclusion

Important productivity gains in crash design are now possible with high-performance, scalable computing solutions; even 10-million-element models can be run overnight with ever-increasing precision. The SGI Altix 3000 supercomputer, based on Itanium2 processors, in combination with PAM-CRASH software from ESI Group, is a cost-effective crash simulation solution, offering the shortest restitution times for the largest models.

For more information, please visit:

http://www.sgi.com

http://esi-group.com

http://intel.com

3790 [August 2005] J14870