Meeting the High-performance Demands of Industrial Computing with Itanium® 2-based systems for SGI® Altix® Servers

Table of Contents

Itanium® 2-based systems for SGI® Altix® 3000 servers and MSC.Nastran™ software create a scalable, high-performance solution for solving problems with very large degrees of freedom in Finite Element Analysis (FEA).

Executive Summary

The demands on technical computing in Finite Element Analysis (FEA) are growing. Problems with hundreds of thousands to millions of degrees of freedom are becoming commonplace to successfully analyze designs, requiring a high-performance, highly scalable, yet cost-effective computing solution. While the cost of previous computing solutions, based on specialized processors and proprietary operating systems, have escalated, the SGI® Altix® 3000 server platform provides a cost-effective, high-performance solution to FEA and other industrial computing applications.

The Itanium® 2-based systems for the SGI Altix system, running the Linux* 64-bit distribution operating system (OS) and MSC.Nastran FEA software, create a unique, highly scalable solution for solving analysis problems with increasing degrees of freedom. Intel® Itanium® 2 processor performance, SGI® NUMAflex™ technology, and SGI's enhanced tools for the Linux OS, along with MSC.Nastran compiled for the Itanium 2 microarchitecture, enable reduced time-to-solution for large analysis problems.

This paper surveys the SGI Altix system and MSC.Nastran performance for FEA problems, and then describes the technologies behind the solution's performance that enables high productivity in industrial computing.

September 2004

Pushing the analysis envelope in MCAE

As business executives, customers, and government regulators increasingly press product manufacturers to produce better, more feature-rich, and safer designs for less cost, greater demands are placed on the product design phases. This is especially true in the case of analysis of complex mechanical designs, such as those found in automotive, aerospace, and other industries. For years, designers have relied on the computer for mechanical computer-aided engineering (MCAE) in Finite Element Modeling (FEM), Finite Element Analysis (FEA), and crash synthesis. And each year the demand on industrial technical computing requirements increases as MCAE pushes the envelope of mechanical design analysis.

The challenge of solving problems with ever-larger degrees of freedom results in much more intense analysis that the computing solution must complete. Hundreds of thousands to millions of degrees of freedom are common in automotive designs today. And aerospace designs can require hundreds of millions of degrees of freedom, while a billion degrees of freedom problems will be undertaken in the near future. Such scenarios create massive data sets and require intensive floating-point operations, which go beyond the capabilities of 32-bit computing with its limited addressable memory capabilities.

Problems of this magnitude demand very powerful computing platforms and software with advanced algorithms in order to produce accurate solutions in a reasonable amount of time. Accuracy and time to solution are key considerations when developing a computing platform solution for MCAE. Most of today's computing solutions in MCAE analysis are parallel processing systems, and platform hardware characteristics are specified to meet the rigorous demands of the very large problems encountered. These characteristics include

- Fast processor speed, high throughput, strong floating-point engine, and large cache size
- Easily scalable system
- Large memory capacity with 64-bit addressability
- High system bandwidth with low latency
- Very fast interconnect among nodes

Analysis software designed for parallel computing must maximize the ratio of computing to communications, and the operating system (OS) must take advantage of hardware features to ensure the fastest time to solution.

Traditionally, computing solutions for MCAE have been built on specialized 64-bit processor technologies and proprietary operating systems, resulting in very expensive systems. Today, with the Intel Itanium 2 processor and the Linux open-source operating system, powerful, parallel computing solutions for MCAE are being developed for mechanical design analysis requirements at a much lower cost. With hardware and software optimized for 64-bit Itanium 2 processors to quickly solve complex problems, these systems provide fast solutions for MCAE. One such industry-leading computing solution combines the Itanium 2 processor-based SGI Altix 3000 server platform with MSC.Nastran analysis software.

An industry-leading MCAE analysis solution

Both SGI and MSC have been providing solutions for virtual product development (VPD), including FEA, for many years. The SGI Altix 3000 platform and MSC.Nastran for Linux are recent evolutions in a long history of successful MCAE product releases. The SGI/MSC solution offers competitive price/ performance in a highly scalable, parallel processing platform.

MSC.NASTRAN™ BENCHMARKS ON ALTIX® (SINGLE CPU)

MSC Software measured time-to-solution of several MSC.Nastran solutions for problems with varying degrees of freedom, as listed in Table 1. These tests were run on different hardware platforms using the standard MSC.Nastran distribution for the particular OS of the system-under-test. For tests on the SGI Altix system, the standard MSC.Nastran distribution for Linux was used. The entire set of systems tested can be found on the MSC Software Web site.[1]

Note the complexity of the problems indicated by the degrees of freedom and the amounts of application memory used, scratch disk space consumed, and total I/O throughput. While the solutions require application memory space well within the limits of 32-bit computing solutions, the scratch space in all but one solution exceeded the limits of memory directly addressable by 32-bit systems. This means direct disk reads and writes for scratch data, which slows performance.

Figure 1 illustrates the time-to-solution (in seconds) achieved by a single-processor SGI Altix system running the standard Linux distribution of MSC.Nastran code. The standard distribution for Linux does not include optimizations for MSC.Nastran jointly developed by SGI and MSC Software to enhance performance on Linux. These tests provide baseline values for MSC.Nastran on the SGI Altix system.

Figure 1: Time to solution for MSC.Nastran™ standard Linux® distribution running on SGI® Altix® (single CPU) system (Source: MSC Software, 2004)

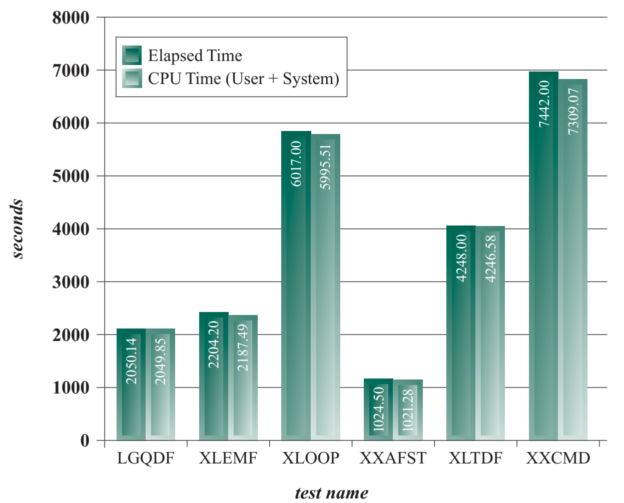

Table 1: MSC.Nastran™ v2004 serial test problems

Name	N Degrees of Freedom	Description	MSC.Nastran™ Solution	Memory Used	Scratch Disk Space	Total I/O (GB)
LGQDF	31,125	Cube with Interior	108	100 MB	600 MB	1,000
XLEMF	654,560	Car Body	111	400 MB	11 GB	400
XLOOP	486,573	Car Body	200	1 GB	28 GB	1,700
XXAFST	2,490,516	Propeller Housing	101	400 MB	11 GB	77
XLTDF	529,027	Car Body	108	450 MB	5 GB	209
XXCMD	1,584,622	Car Body	103	800 MB	43 GB	2,400

SCALABILITY BENCHMARK

While a single-processor system can provide results in an acceptable time on small problems, larger problems result in longer running times. Thus, today's application code is scalable, supporting parallel processing on multiple processor nodes to take advantage of computing platforms with greater available computing resources. Scalable systems—both in terms of software and hardware—can significantly improve time-to-solution, as illustrated in Figure 2.

The same MSC.Nastran solutions were run on a multi-processor SGI Altix server using MSC.Nastran code optimized for up to eight Itanium 2 processors. As the system is scaled up, results are achieved in about half the time.[2]

The SGI Altix 3000 platform running optimized MSC.Nastran code offers a highly scalable solution. The combination of SGI technologies, MSC Software code optimized for the Itanium 2 processor, and SGI enhancements for the Linux OS enable high performance for mechanical analysis.

Figure 2: Scalability testing results of MSC.Nastran™ SOL108 running MSC.Nastran optimized for SGI® Altix® 3000 series server and Intel® Itanium® 2 processor (Source: SGI internal testing, 2004)

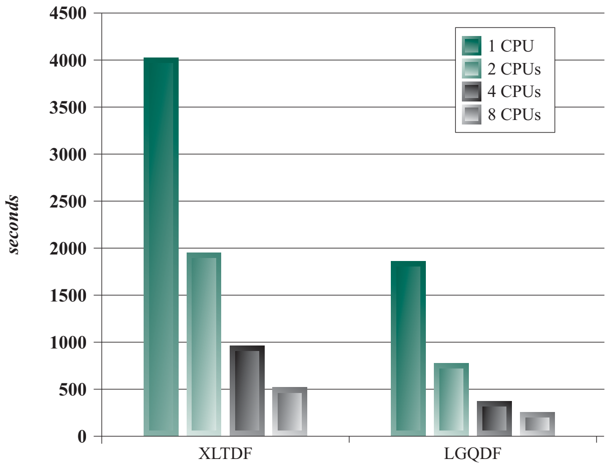

SGI Altix 3000 server platform

SGI has a strong history in MCAE. Its earlier solution for MCAE was the SGI® Origin® Server Series based on the company's 64-bit MIPS architecture processor and running the SGI® IRIX® operating system. The SGI Altix 3000 server platform is the next-generation parallel processing system based on the 64-bit Itanium 2 processor and running the 64-bit Linux operating system.

KEY SGI ALTIX SYSTEM TECHNOLOGIES

Three key technologies for the SGI Altix system create a high-performance, scalable computing platform for FEA: SGI NUMAflex technology, the Itanium 2 processor with 6 MB L3 cache, and the Linux 64-bit operating system.

With the Origin system, SGI introduced its SGI NUMAflex technology that enables non-blocking, shared access of a single large block of contiguous memory (up to the physically addressable memory capacity of the processor architecture and operating system) by up to 512 processors. NUMAflex technology creates a single server node sharing memory and running a single instance of the operating system.

Figure 3: SGI® Altix® 3000 architecture with NUMAflex™ technology

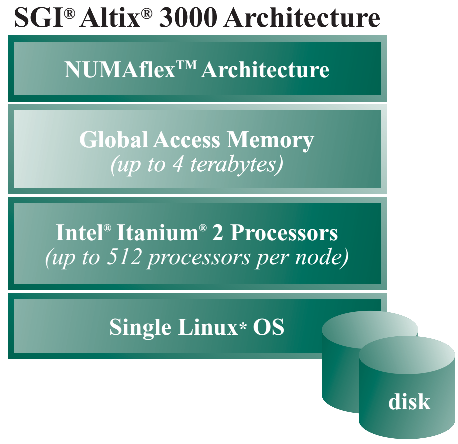

Table 2: Latency/bandwidth comparisons[3]

	InfiniBand Architecture	Myrinet*	Gigabit Ethernet	Altix®
End-to-End Latency	7.6 ms	8 ms	60 ms	1.8 ms
Bandwidth	822 MB/s	250 MB/s	125 MB/s	1165 MB/s

For very large Origin system configurations, with several multi-processor nodes, SGI developed SGI® NUMAlink™ technology, a very high-speed interconnect technology for large super-cluster systems with performance beyond that of currently available inter-connect technologies (Table 2). With NUMAflex and NUMAlink technologies, SGI can create high-performance, highly scalable systems using shared memory across hundreds of processors. SGI extended NUMAflex technology to the SGI Altix 3000 platform architecture, bringing the same shared memory architecture to a new family of computing platforms.

64-bit computing is a fundamental requirement for solving problems with very large degrees of freedom. These problems create massive data sets with very large memory requirements, often over the 4 GB limitation of 32-bit computing. The 64-bit addressability of the Itanium 2 processor not only allows very large memory capacity for application memory, it also allows the OS to create scratch disk space within the memory instead of on the storage system (note the scratch space required by the MSC.Nastran benchmarks in Table 1, on the previous page). Having scratch space in memory improves time-to-solution, since memory reads and writes are much faster than disk operations. The Itanium 2 processor's 6 MB L3 cache further boosts performance with fast access to data held in the large L3 cache. In addition, analysis problems are both floating-point and I/O intensive (review the Total I/O throughput of Table 1, on the previous page). The Itanium 2 processor's large computing resources include two floating-point engines that each execute two operations per clock cycle, multiple 128-bit registers, and 6.4 GB/s throughput to support the intensive operations of the MSC.Nastran software.

The adoption of Linux throughout industry and science, the continuously evolving Linux OS, and the advent of a 64-bit Linux distribution has made Linux a strong part of the industrial computing ecosystem and now a key component of 64-bit computing solutions. At Supercomputing 2003, Christian Tanasescu, Director of Simulation Applications at Silicon Graphics, Inc., highlighted the rapidly progressing adoption of Linux in the automotive industry through a survey of the top 20 high-performance computing installations in automotive. The combination of the standards-based, high-performance, 64-bit Itanium 2 processor and a highly reliable 64-bit, open-source OS has benefited automotive and other MCAE markets with improved price/performance of parallel computing systems, such as the SGI Altix 3000 platform.

SGI LINUX* TOOLS ENHANCE PERFORMANCE

SGI NUMAflex technology enables any processor in the system to access any memory, which could lead to increased memory access latencies if the processors running code and the memory where data is stored are not localized. Optimized system performance is maintained when processor affinity is maintained, that is, when CPU and memory resources are from the same node.

The standard 64-bit Linux distribution is not aware of the SGI NUMAflex technology, yet a task of the OS in multi-processor systems is to assist application code (MSC.Nastran code) to get processor affinity. SGI provides enhancements to the Linux OS with tools to help end-users establish affinity between user programs and CPU and memory resources. These tools include the following:

- **CPUSET**—soft partitioning of CPU and memory resources
- **RUNON**—runs a job on specific CPU(s)
- **DPLACE**—binds a job to specific CPU(s)

CPUSET enables partitioning of CPUs and memory in a large system. CPUSET limits the impact processes can have on other jobs running on the system, and they are not affected by processes running elsewhere on the system.

RUNON and DPLACE submit a job to a set of CPUs prescribed by the user. Except for I/O buffer cache, the default memory allocation policy for these commands is to use local memory first, then the nearest physical memory if the local memory is consumed. These features help maintain processor affinity. The DPLACE command also binds processes to prescribed CPUs. For I/O-intensive processes, this becomes important to keep the OS from migrating the process to another CPU after an I/O interrupt, which often occurs because the OS tends to forget the best CPU to run the process on.

These enhancements, designed to improve MSC.Nastran performance, are part of the MSC.Nastran software distribution for Linux when distributed with the SGI Altix 3000 system.

MSC.Nastran

MSC.Nastran is well established as an industry-leading FEA application, now ported to the Itanium 2 processor for the SGI Altix 3000 platform. MSC.Nastran is a scalable application that takes full advantage of the features of the SGI 3000 platform and NUMAflex technology.

Intel supported MSC Software to port and tune the MSC.Nastran code. MSC.Nastran performance was optimized by modifying the code to take advantage of Itanium 2 processor architectural features, such as the large L3 cache, making extensive use of Intel's Math Kernel Library optimized for the Itanium 2 processor, and by taking advantage of Intel compiler optimizations, including loop unrolling, software pipelining, and data prefetching.

ENABLING BINARY COMPATIBILITY

The industrial design environment often includes a heterogeneous network, where pre- and post-processing of FEM models, FEA solutions, computing, and database and file serving are all done on different computing platforms. Such an environment introduces binary incompatibilities in file formats, since UNIX* uses a big-endian format, while Linux and the SGI Altix 3000 platform use the little-endian format. MSC.Nastran offers several solutions to enable binary compatibility, making it simple for the user to move files across a heterogeneous network.

MSC.Nastran provides several tools, including TRANS and RCOUT2, to convert certain MSC.Nastran files from one format to another. For example, RCOUT2 can convert a MSC.Nastran OUTPUT2 file from the neutral format (ON2) to the native binary format (OP2) for a UNIX workstation. But, conversion operations can slow down the workflow; it is more desirable to deal with binary files directly. The Intel Itanium 2 compiler supports an F_UFMTENDIAN environmental variable that can be used in conjunction with a user program to convert commonly used MSC.Nastran Fortran unformatted binary files.

IMPROVING I/O THROUGHPUT WITH FFIO

With NUMAflex technology, the software and operating system can use memory as a cache and I/O buffer and perform parallel I/O asynchronously, boosting I/O performance. MSC.Nastran incorporates an SGI high-performance I/O interface called FFIO to improve I/O throughput. FFIO significantly improves I/O performance per job, as shown in Figure 4.

Figure 4: MSC.Nastran™ I/O job throughput (Source: SGI, 2004)

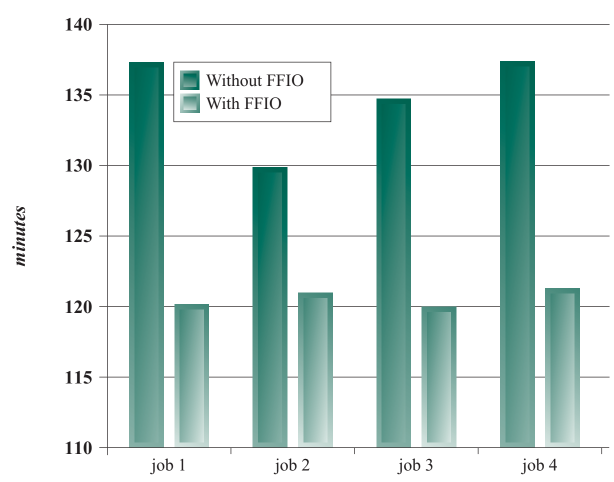

Conclusion

To enable high productivity in FEA applications, a high-performance, scalable computing solution that can quickly solve today's complex FEA problems with ever-increasing degrees of freedom is needed. The Itanium 2 processor-based, parallel-processing SGI Altix 3000 platform and MSC.Nastran FEA software form a cost-effective, high-performance computing solution for analysis and other MCAE applications, with fast time-to-solution on highly complex analysis problems.

LEARN MORE ABOUT THIS INNOVATIVE SOLUTION

For more information, see:

http://www.intel.com/itanium

http://www.sgi.com/industries/manufacturing/

http://www.mscsoftware.com

[1] www.mscsoftware/support/prod_support/nastran/performance/v04_sngl.cfm

[2] This level of scalability was achieved with Nastran* SOL108. Other sequences result in lower scalability.

[3] InfiniBand Architecture, Myrinet*, and Gigabit Ethernet data source: Voltaire*, 2003; Altix* data source: SGI, 2004.

1004/GJA/OCG/XX/PDF-3708

Order Number: 303790-001US

intel®

sgi®