[Volume 1]—Fall 2001

Eye on Innovation

Manufacturing Edition: SGI News

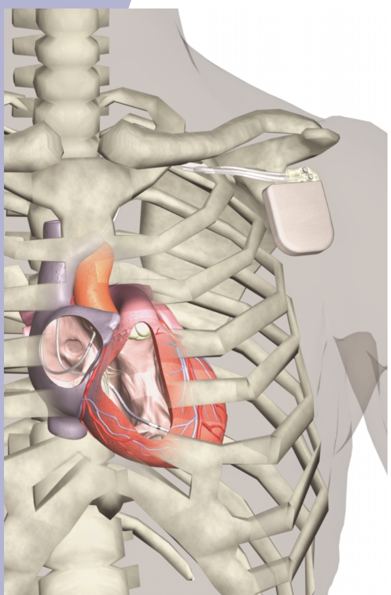

"Medtronic would not be achieving the high level of quality products that it is if it weren't for the visualization and simulation activities that are allowed by SGI technology."

—Tim Abraham, Surgical Simulator Development Coordinator, Cardiac Management Therapy Delivery Division, Medtronic

Medtronic Uses SGI™ Technology for Innovations in Life-Saving Medical Devices

Medtronic is the world leader in medical technology, providing lifelong solutions for people with chronic disease. Specifically, the company is the world's leading designer and supplier of pacing systems for the treatment of heart conditions, heart failure, and various cardiovascular diseases.

A Medtronic pacing system consists of three essential components: the pacemaker, the pacing lead, and the programmer. The pacemaker works in conjunction with the pacing lead, a special insulated wire that carries tiny electrical pulses from the pacemaker to the heart and relays information about the heart's electrical activity back to the pacemaker. These two components are surgically implanted inside the body. The third component, the programmer, is a specialized computer that is kept at the hospital or clinic and used by a physician or nurse to monitor and adjust the settings of the pacemaker.

The first Medtronic nonimplantable pacemaker was developed in 1949. Medtronic's first implantable therapy appeared on the market in the late 1950s and was designed to normalize the too-slow heartbeat of bradycardia patients. Since 1985, significant improvements in pacemaker patients' quality of life have been achieved by adding the ability for the pacing system to adjust automatically to changing activity levels, such as when the patient is exercising. In the early 1990s Medtronic introduced a therapy for people who go into tachyarrhythmia or tachyarrest, a condition in which the heart is beating too rapidly. It was about this time that Medtronic began using RISC systems that were based on SGI workstations with MIPS® R4000® processors. Medtronic needed to develop a higher level of knowledge about its products, how they behave, and how they could be developed or modified more quickly. The SGI 3D visualization environment proved to be the perfect solution for facilitating new product-development practices.

Medtronic's newest therapy, cardiac resynchronization, addresses heart-failure conditions. At present, only 4% of would-be heart transplant recipients ever receive a new heart. The majority of patients on the waiting list continue to suffer from the condition of their deteriorating hearts and/or from directly related complications—poor circulation, diabetes, kidney failure, and brain malfunction. A newly released system called InSync® is emerging as a promising treatment for some of the more than 20 million people worldwide now in the throes of progressive heart failure. In the United States, hospitalization for chronic heart failure totals more than 6.8 million days a year. The total cost of treating the disease is more than $38 billion annually, and the cost is expected to mount as the population ages. With InSync, Medtronic seeks to reduce the number of hospital visits and improve the quality of life for many of these severely compromised patients.

Computer-Based Modeling and Simulation

SGI systems for computer modeling and simulation, such as Silicon Graphics® O2® and Silicon Graphics® Octane® workstations and SGI™ Origin® family servers, are helping Medtronic to speed the design and development of new therapies.

Tim Abraham, surgical simulator development coordinator at Medtronic's Cardiac Management Therapy Delivery Division, explains, "The computer technology used is dramatically reducing design time and changes the way we think of the design process. We now are able to do more, faster, and have extreme confidence in the design decisions that we make. The simulation technology we have developed gives us critical traction required in this competitive marketplace."

The company uses a host of software products to explore new possibilities, such as I-DEAS®, Pro/ENGINEER®,

[continued on page 3]

IN THIS EDITION

SGI Architectural Direction to Address Customer Scalability Needs

Eng Lim Goh, Ph.D.
Senior Vice President and Chief Technology Officer, SGI

Manufacturers worldwide continue to invest significantly in high-performance computing (HPC) systems. The combined computational power from industries such as automotive, aerospace, high-tech, consumer products, and others has grown about tenfold in just four years. This power is deployed to conduct a wide range of simulations for designs that require structural integrity, efficient cooling and fluid flow characteristics, reduced noise, and minimum weight, among others. SGI research and development focus has been on developing low-latency, high-bandwidth interconnects and software. This, coupled with a strategy that targets only the technical and creative customers, has made us the leading high-performance computing systems provider for computer-aided engineering within the manufacturing industry.

The SGI™ NUMAflex™ modular design concept enables the SGI™ Origin® 3000 series of servers to scale to an industry-leading 1TB of single-address-space memory accessible by each of the 512 processors in the system. This allows designers and engineers to solve complex problems involving very large-scale, high-fidelity modeling and simulation in a drastically reduced time frame. Therefore, together with its modular scalable architecture, the SGI Origin 3000 series is well suited to the needs of major technical data centers.

A recent addition to the SGI Origin family, SGI Origin 300 was designed to extend this NUMAflex modular approach to lower-cost and space-saving configurations that can be tailored to the specific high-performance computing needs of engineering departments or perhaps even individual engineers. As a stand-alone computational system, SGI Origin 300 packs the power of up to four MIPS® processors and 4GB of memory into an ultradense, 2U (3.5-inch) form factor. Up to eight of these 2U modules can be connected to create a 32-processor single-address-space memory system, fitting into only half a rack.

This new product can have an impact on the product-development process by, for example, allowing car companies to give their best engineers the personal power and freedom to explore more alternatives earlier in the design cycle, particularly while their models are still in their initial stages. At the significantly lower entry price, even customers with smaller budgets can build a robust, high-performance, high-throughput computing environment (in shared-memory and clustered modes) for simulation.

Across the industry, even more computational power is expected, as engineers are starting to look at the kinds of problems that are well beyond the limits of current technology. Recent focus has been on high-fidelity modeling and analysis of only a few design alternatives. However, to be able to come up with truly revolutionary new products, engineers now want to understand the behavioral response of hundreds of design parameters involving multiple disciplines and for entire families of designs.

In response to these demands, we go back to our basic technology vision, which centers on the three fundamental systems-performance factors: computational power, bandwidth, and latency. Computational power, the first factor, will continue to improve as it has in the past, consistently confounding those who predicted that Moore's Law cannot continue. Bandwidth will also continue to improve, particularly when it has multiplexing techniques working for it. However, latency has physics working against it, as light speed is known to be limited to one nanosecond per foot. Therefore, eventually, the most powerful computer may have to be the densest.

Consequently, we have directed our research and development efforts toward the performance of the MIPS processor, system interconnect, capacity, and density in order to deliver a highly scalable yet balanced system. It would not be inconceivable that within five years we will be able to offer 1 TFLOPS of computational power and terabytes of single-address-space memory (not memory that is just distributed or clustered) all in just one or two racks. SGI will continue to advance the field of high-performance computing for scalable technical applications to enable true simulation-based design.

Eye on Innovation

CUSTOMER

[continued from page 1]

ABAQUS®, Maya®, and customized tools written by Medtronic for use in conjunction with these applications.

Modeling the Environment

In developing complex medical devices intended for implantation, modeling the device is only half the equation. Medtronic also models the environment in which the product is being placed, in this case the heart. A primary focus is to provide better, more reliable products for its patients. Medtronic pacing-system research and development focuses on uncovering new innovations within its products. At present there is much study surrounding pacemaker lead materials, and Medtronic plans to meet the challenges of continuous innovation through use of simulation and modeling technologies.

According to Abraham, "New surgical simulators will need to be able to model complex anatomies as well as physiologic behavior and cardio physical properties. In addition to procedural simulator training, it is also hoped that the models will be able to predict design performance prior to production."

In the future the advancements in heart modeling will result in a rapid increase in model complexity. Conditional slider bars can then be used to quickly modify the condition of the model by enlarging the heart, enlarging the valve, increasing or restricting the blood flow, and then mapping the conditions to best represent the type of environment in which the lead should be placed. The models will be used for research and development of new products and for training purposes, as well as in clinical studies for product-comparison analyses.

ABAQUS, from HKS, is Medtronic's primary mechanical modeler. The company also uses Maya, from Alias|Wavefront (a division of SGI), the same application that is used by the entertainment industry for special effects in movies such as *Hollow Man*. Maya provides a highly realistic 3D representation of what is actually occurring in the heart.

Trends in Performance and Model Complexity

Through the years, SGI hardware and software solutions have helped Medtronic minimize the time required to test models while maximizing product development speeds. Using SGI™ 2000 series systems, the complex jobs ran three weeks. With the introduction of a new SGI Origin 3000 series system with a 600 MHz processor, that time dropped to just one week. With further advancements in processor speed and computing systems architecture working in combination with optimized code and newly developed algorithms, courtesy of ABAQUS, Medtronic hopes to achieve overnight turnaround on highly complex jobs within the next few years.

Among the many benefits of these faster computers being used by the scientists are being able to solve more problems and the ability to add features to their models or increase the number of elements without worrying about bringing the computer to its knees. According to Abraham, "Electrophysiologic models increase the complexity by a factor of nearly 10. If end users drive the model to its end point, where I believe they want to, the complexity of the models we run today may well be just a fraction of those we are running within three years."

Compute and Storage Capacity

To ensure adequate compute capacity, SGI is working with Medtronic to divide the workload between their SGI 2000 series and SGI Origin 3000 series systems. SGI is also helping to implement a scheduling program and making sure CPU sets are in place to optimize utilization and performance. With the release of ABAQUS 6.2, scalability is now up from four CPU jobs to six to eight CPU jobs depending on the type of elements used in the job.

The SGI 2000 series system, with an InfiniteReality2™ graphics pipe, also serves as Medtronic's primary rendering machine for any type of data sets, such as volume data or Maya 3D models. The SGI 2000 series system is complemented by the real-time volume rendering capacity of the Octane workstations. In addition to its rendering capabilities, the SGI 2000 series system is a great tool for animators because of the speed at which it rasterizes images.

Complex 3D models take up a lot more space than their contemporaries. And with regulatory mandates requiring companies to have an audit trail for the data generated as new products are created, data storage capacity is fast becoming a crucial component of pacing system development. To meet the requirement, Medtronic has added nearly 3TB of storage to its SGI 2000 series and SGI Origin 3000 series systems and bumped up disk storage capacity significantly on its SGI workstations.

Next Steps

SGI computer modeling and simulation are helping move Medtronic in other new and equally inspiring directions. The Medtronic surgical simulator project now under way strives to improve the quality of patient therapies by using 3D modeling to provide physicians with advanced levels of technical and training support. Medtronic envisions its virtual surgeries being distributed to physicians in hospitals and clinics via a

[continued on page 8]

Benefits of SGI Origin 300
Technology for MCAE Applications

The application of mechanical computer-aided engineering (MCAE) provides essential tools for the efficient design of products and processes. Industries such as automotive and aerospace and a variety of general manufacturers benefit from MCAE applications that enable reduced design-cycle time and costs and improved design quality.

MCAE applications extend the geometry-definition capabilities of mechanical computer-aided design software and offer methods for evaluation of a design's functional behavior under "in-operation" conditions. This is accomplished through application of numerically intensive MCAE software with high-performance computing (HPC) system technology.

This article examines the benefits of a new HPC server, SGI Origin 300, which was developed to advance current capabilities of technical HPC, including MCAE applications.

Introduction

The combination of MCAE application software and HPC system technology provides engineers with an increasingly competitive advantage in today's global manufacturing industry. MCAE applications assist engineers with predictions of design function and performance by simulation of the mechanical loads applied during a design's operation. Examples of these are thermal and mechanical stress, vibration, impact loads, and fluid-induced pressure, among others. This variety of load conditions gives rise to design complexities and trade-offs when target objectives for product safety, fuel economy, environmental impact, and consumer appeal are considered with attempts to reduce time-to-market and increase competitiveness.

Developments in HPC and visualization technology continue to rapidly advance MCAE simulation capabilities of engineers globally. SGI is an established leading supplier of this technology with the company's SGI Origin and SGI™ Onyx® product families. Evidence of SGI success was reported in a recent study issued by D. H. Brown & Associates entitled "Symmetric Multiprocessing Dominates High-Performance Computing in CAE," which concluded that SGI was the platform of choice of respondents. This conclusion is consistent with similar industry views that include most MCAE software providers, who acknowledge SGI as their most strategic hardware partner.

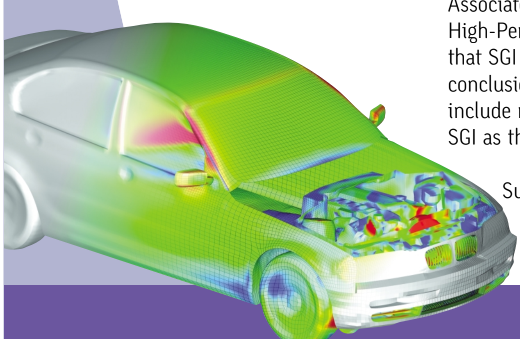

Success with MCAE applications is the result of SGI's ability to offer a scalable SGI™ Origin® 3800 single-system image (SSI) with up to 512 processors and 1TB of single address space memory, pushing MCAE simulations to achievements of a historic level. At the core of this remarkable capability is a scalable IRIX® operating system and the cache-coherent SGI™ NUMA (nonuniform memory access) system architecture. These HPC technologies combine to offer industry and manufacturing research organizations a high-availability, nondegrading, and efficient application environment that ensures turnaround and throughput are delivered in support of hundreds of simultaneous users with a demanding mix of MCAE applications.

SGI continues to invest in the MCAE community's shared vision and success with the introduction of SGI Origin 300, an HPC server that enables further advancements for MCAE applications. SGI Origin 300 delivers the existing application advantages of IRIX and NUMA with complete SGI Origin 3000 series compatibility, yet in a package that is significantly smaller than a similar SGI Origin 3000 series configuration. While the SGI Origin 3000 series offers large-scale data center HPC, SGI Origin 300 scales from 2 to 32 processors and offers a midrange HPC solution as either a departmental SSI server or a large HPC cluster of these servers.

HPC Characteristics of MCAE Software

From a hardware and MCAE software algorithm perspective, there are roughly three types of MCAE application disciplines to consider: implicit and explicit finite element analyses (FEA) for structural analysis and computational fluid dynamics (CFD) for fluid flow simulation. There is further segmentation of implicit FEA since structures (and associated algorithms) at rest (static) and those in motion (dynamic) behave differently. This gives rise to four distinct types of MCAE software application behavior.

Generally speaking, MCAE software applications exhibit a range of various HPC resource demands for each of the four segments described. An examination of these demands aids in a proper characterization of what features are required from an effective HPC system architecture. For example, the desired features for MCAE from a RISC system architecture include fast processors with a Level 2 cache, large addressable memory, high memory-to-processor bandwidth rates, high disk-to-memory I/O rates, and a low-latency interconnect that provides efficient parallel scalability to hundreds of processors.

A closer examination of the different demands from each of the four segments highlights the importance of a balanced HPC system architecture. For example, implicit FEA for static load conditions requires a fast processor for effective turnaround, in contrast to dynamic response, which requires high rates of memory and I/O bandwidth with processor speed as a secondary concern. In addition, FEA modeling parameters such as the size, the type of elements, and the load condition of interest all affect the resulting execution behavior of implicit and explicit FEA applications.

Explicit FEA benefits from a balance of fast processors for the required element force calculations and high rates of memory bandwidth necessary for the efficient contact resolution that is required for nearly every structural impact simulation. CFD also requires a balance of memory bandwidth and fast processors, but benefits most from parallel scalability. Each segment has inherent complexities with regard to efficient parallel scaling, depending upon the particular parallel scheme and system architecture. While CFD scales efficiently to hundreds of processors, explicit FEA scales to 50 and implicit FEA to less than 10.

A system architecture that can achieve high parallel efficiency becomes increasingly important as algorithms for MCAE software applications have developed such capability. Most commercial MCAE software employs a distributed memory parallel (DMP) technique for compatibility across the range of RISC architectures available. Other techniques include shared memory parallel (SMP) and hybrid parallel schemes that take advantage of both DMP and SMP within a single computation. The scalability of SMP algorithms is limited to the number of processors offered in an SSI, meaning scaling beyond the maximum 32 processors of an SGI Origin 300 server would require implementation of a DMP technique.

Most MCAE software is carefully designed to avoid major sources of parallel inefficiencies, whereby communication overhead is minimized and proper load balance is achieved. For MCAE software that utilizes a DMP technique with a message-passing library such as MPI, development of an SGI NUMA-aware MPI is included for MCAE software and is transparent to the user. This MPI further reduces communication overhead when scaling to a large number of processors, which is achieved by a reduction in latency that is more than threefold improved over public-domain MPICH.

SGI Origin 300 System Technology for MCAE

The SGI Origin family has a NUMA multiprocessor architecture that is a breakthrough implementation of conventional SMP architectures. The SGI NUMA architecture distributes memory to individual processors in order to reduce latencies that inhibit high bandwidth and scalability. At the same time, all memory is globally addressable to enable high-resolution MCAE modeling and simplify MCAE algorithm development. A single image SGI Origin 3800 system offers the largest SMP system currently available.

The SGI NUMA architecture was introduced in the SGI Origin 2000 server in 1995 and later advanced with the SGI NUMAflex modular design concept of the SGI Origin 3000 series. Now the same high-bandwidth and low-latency NUMA architecture is available in SGI Origin 300, yet with a significant cost-performance advantage for MCAE applications. This is achieved, among others, by reducing the size of Level 2 cache of the MIPS processor from 8MB to 2MB, while keeping the performance penalty limited to 15%, even with the most cache-intensive MCAE applications.

The modular building block of the SGI Origin 300 system is a node that contains two or four MIPS processors, corresponding memory up to 4GB, and a connection to a portion of an I/O subsystem. The hub interface to the node is the distributed memory controller, and nodes are connected together via the NUMAlink™ cable or module in a maximum SSI configuration of 32 processors that requires only half a rack. Alternatively, these same node modules can be clustered with a choice of scalable interconnect networks to much larger processor counts and system configurations.

Conclusions

A discussion was provided on the HPC technology requirements of MCAE applications, including characterizations of the performance behavior typical of four types of conventional MCAE simulations. The HPC technology offered with the new midrange SGI Origin 300 server is well suited to the demands of each MCAE application segment. In particular, the choice offered with SGI Origin 300 as a moderately configured SSI for departmental deployment, or as a large cluster for scalable DMP applications, provides substantial performance benefits to MCAE applications for a range of modeling requirements.

SGI Origin 300 design breakthroughs in cost-effectiveness and compact packaging will further expand the use of MCAE and HPC to include a variety of applications in the mainstream of product and process development. With SGI Origin 300, the entire suite of MCAE applications can now achieve capability levels that are nearly equivalent to the industry-leading SGI Origin 3000 series up to 32 processors, but at roughly half the cost. This new SGI technology development for the MCAE community demonstrates the company's continued commitment to delivering valuable leadership to the manufacturing industry.

Dassault Systemes and SGI Deliver
SGI™ Immersive Solutions Review Based on V5 Architecture

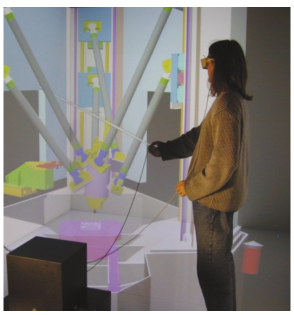

Dassault Systemes, the premier global software developer of product life-cycle management solutions, and SGI have embarked on a CAA V5 development partnership to create high-end immersive solutions in hardware, planning, devices, software, and customization that are fully integrated within the ENOVIA DMU navigator environment.

“This great product is the result of a long fruitful collaboration between Dassault Systemes and SGI on the V5 visualization infrastructure. By leveraging this infrastructure in a virtual reality environment, our customers will benefit from best-in-class SGI experience and recognition in the field of immersive solutions. Delivering immersive solutions is a key element of Dassault Systemes’ product life-cycle management strategy to make the virtual a reality,” says Severin Lanfranchi, ENOVIA portal director, Dassault Systemes.

The newly forged development partnership ensures long-term product integration within CATIA V5, punctuated by timely and compatible future product releases. Both Dassault Systemes and SGI development-partnership industries stand to benefit from the latest technologies.

Breaking New Ground
The first product to emerge from the partnership is the SGI Immersive Solutions Review Based on V5 Architecture, a value-added software module that works with CATIA V5, which supports digital product definition and simulation. Optimized when used on Silicon Graphics® Onyx2® or SGI™ Onyx® 3000 series servers with single or multiple graphics pipes, the SGI Immersive Solutions Review Based on V5 Architecture allows real-time interaction with design data, similar to what end users would experience in the physical world or in a way that is impossible in the real world. Additional real-time visualization techniques exclusive to the SGI Immersive Solutions Review Based on V5 Architecture include SGI™ ClearCoat™ and SGI ClearCoat™ 360. These programs bring virtual designs yet another step closer to realism by adding reflective or natural lighting to an object, making it appear just as it would in the real environment.

Seamless Data Integration
For the first time, multidisciplinary workgroups can work with native CAD data within an immersive environment without leaving their CATIA interface. Entry to the immersive environment—SGI™ Reality Center™ rooms, powerwalls, workbenches, and desks—is accomplished with point-and-click ease. Users can then navigate their way through the scene using 3D input devices.

The SGI Immersive Solutions Review Based on V5 Architecture delivers real-time data that can be easily evaluated. Advanced interactions such as cutting through a model or measuring distances are infinitely easier to perform, and in the future kinematics relationships can be verified or disproved with a single entry. Other key enhancements include the user’s ability to select parts, annotate problems with specific 3D markers, and document the entire session.

Because there are no data exchange problems, seamless integration of virtual reality technology into existing process flows is now possible. Previously, data from 3D styling tools or 3D-CAD systems had to be made compatible with VR software systems, causing a loss of semantics and making it difficult to establish a feedback path for updating data in the underlying CAD model from within the VR software.

The Real Beneficiaries
The most significant beneficiaries of the SGI Immersive Solutions Review Based on V5 Architecture are manufacturing industries such as automotive, aerospace, and shipbuilding that require extensive product testing from the early stages of development through actual deployment. Compared to the often-daunting task of building multiple life-size physical prototypes, digital 1:1 scale mock-up reviews (or any scale, for that matter) can begin much earlier in the development process, providing a tremendous cost savings over the product-development life cycle while accelerating the product’s time to market significantly.

Immersed in the Future
In the years ahead, immersive projection technology will continue to play an increasingly significant role as the visualization technology of complex and design-oriented products. SGI will continue to collaborate with software application leaders such as Dassault Systemes to develop new and complete solutions sets that enable multidisciplinary workgroups across a range of industries to design, simulate, and visualize their ideas. The collaborative efforts will further enhance the engineering and manufacturing techniques of companies worldwide, enabling them to bring more advanced products to market faster and at a lower cost.

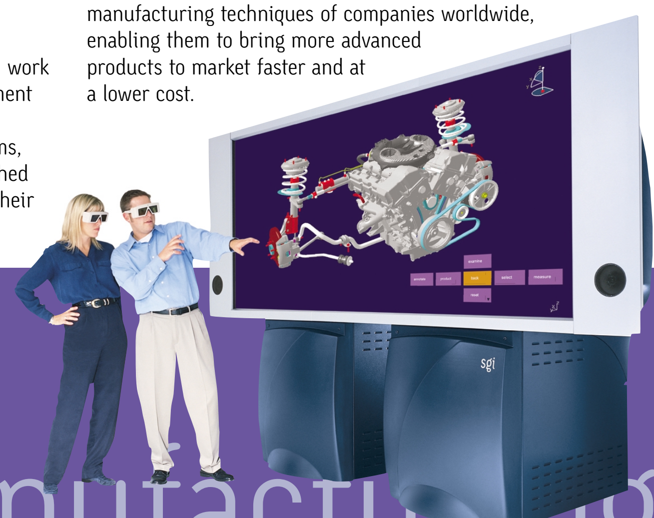

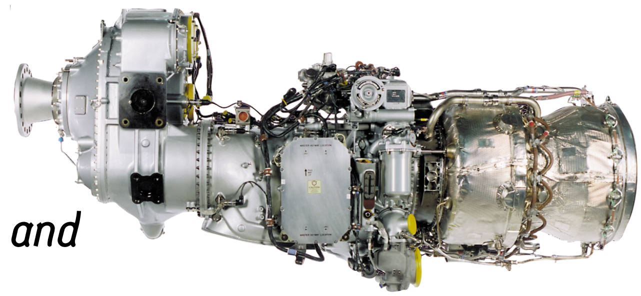

SGI behind Lockheed Martin and Pratt & Whitney JSF Award

On October 26, 2001, the Pentagon awarded the Joint Strike Fighter (JSF) contract to Lockheed Martin. The win signaled a go-ahead for the company to produce an initial 22 aircraft as part of a U.S. Department of Defense System Development and Demonstration (SDD) phase.

To power the JSF aircraft, the Pentagon selected the Pratt & Whitney® F135 engine, a derivative of the company's F119 engine now used in the military's F-22 fighter aircraft. It was an equally big win for Pratt & Whitney, the world's leading manufacturer of gas turbine aircraft engines.

Instrumental in helping both Lockheed Martin and Pratt & Whitney land the JSF contract were SGI advanced visualization and high-performance computing (HPC) technologies. Throughout the design selection process, Silicon Graphics Onyx family systems, SGI Origin family servers, and optimized modeling software used by these two leading aerospace companies enabled them to reduce costs, minimize risk, and maximize the accuracy and performance of their respective designs.

"Our JSF approach [enabled by SGI visualization technology] will radically reduce the cost of sustaining U.S. airpower by ensuring affordability during SDD, production, operations, and support and by achieving operational excellence throughout the program," said Tom Burbage, executive vice president and general manager, Lockheed Martin JSF Program.

The Road to Success

In 1994, the U.S. Air Force, Navy, and Marine Corps, together with British allies, announced a mission to develop and field an affordable next-generation strike fighter aircraft—the Joint Strike Fighter. Lockheed Martin was one of two finalists selected to vie for the $200 billion JSF program, which would eventually create a fleet of some 3,000 aircraft.

JSF program requirements stated that the aircraft were to be built around three distinct design specifications—conventional takeoff and landing, carrier variant, and short takeoff/vertical landing. Additionally, all three designs were to share key high-cost components, including propulsion systems, and have a "cost commonality" of 70% to 90%. The goal: realize significant cost savings in the development, manufacture, and maintenance of the aircraft.

With affordability and performance issues at the core of the JSF program, Lockheed Martin and Pratt & Whitney chose to rely on SGI advanced visualization and HPC technologies, respectively. Powered by SGI Onyx family systems and SGI Origin family servers, the technologies helped both companies generate higher quality, lower cost, and more competitive designs that could be simulated, prototyped, and visualized in 3D in their entirety before any assembly began.

Virtual Development, Real Collaboration

Lockheed Martin addressed JSF design requirement issues through a Virtual Product Development Initiative (VPDI), a program in which everyone involved in the JSF project had direct access to all relevant information and the most current design iteration.

Combined, the VPDI and JSF programs utilized over 450 SGI workstations to run CATIA®, Deneb, and EAI applications as well as several Silicon Graphics Onyx2 systems and SGI™ Origin® 2000* servers spread across multiple disciplines. SGI visualization tools enabled project teams to review, rework, and reintroduce optimized component designs while the project was still in a virtual state, thereby avoiding costly and untimely design changes in the latter stages of aircraft development.

According to Mary Ann Horter, program manager for VPDI, Lockheed Martin, "We have focused on eight key high-payoff initiatives, which include a combination of process change, hardware, and software. In some areas we are seeing reductions in cycle time and cost savings in excess of 70% to 80%. Virtual simulations have been key components in helping us meet these targets."

Testing Flightworthiness

Because the JSF will use a higher percentage of composite materials than previous military aircraft, Lockheed Martin used SGI visualization technology in conjunction with its own Laser Ultrasonic Technology system to verify, part for part, the material flightworthiness of the JSF. According to Horter, "In addition to reducing test cycle times by more than 90%, this cutting-edge technology virtually eliminates setup, simplifies operator interface, and provides design feedback to enhance affordability."

"SGI is proud of the role it has played in helping Lockheed Martin to create the most combat-capable, efficient, and

JSF Award

[continued from page 7]

lethal family of multirole strike fighters in the history of military aviation. The pilots who will fly this low-cost, next-generation fighter aircraft will benefit from its superior performance," said Bob Bishop, chairman and CEO, SGI.

Propulsion Optimization by Design

In its efforts to secure the JSF win, Pratt & Whitney utilized HPC technology for finite element structural analysis and computational fluid dynamics (CFD) modeling software to ensure accuracy and performance in the engine design.

SGI has a long history as a supplier of HPC solutions to various Pratt & Whitney divisions. SGI HPC and solver technology has enabled Pratt & Whitney Canada to develop effective simulation models of mechanical, thermal, and aerodynamic loads on engine components and systems. The company has also used the technology to move beyond traditional modeling disciplines to those such as acoustics or even multidiscipline interaction and optimization.

According to Bernard Proulx, director of Component Engineering at Pratt & Whitney Canada, "We see SGI high-performance computing as key to the successful development of full engine simulation and modeling capabilities as we reach our goal of a virtual engine design environment."

Proulx reports that "[SGI HPC and solver technology has improved] simulation turnaround performance as much as 20-fold, enabling our analysts to increase their modeling resolution for improved accuracy." In short, modeling a full-size engine on SGI systems and using the models for CFD evaluation will reduce engine-development costs, allow for the direct transfer of engine design information to manufacturing, and significantly reduce performance certification time.

At present, Pratt & Whitney Canada employs SGI Origin 2000 and SGI Origin 3000 series systems, with hundreds of MIPS processors and more on the way.

Visualizing the Future

Advanced visualization and high-performance computing technology solutions from SGI enable collaborative decision making, expedite insights to complex problems, and streamline project analyses. From advanced 3D modeling and simulation, digital prototyping, and digital manufacturing to group visualization, SGI transforms product-development processes, thereby enabling the aerospace industry to arrive at design optimizations in less time, experience enormous cost savings, and reduce the time-to-market for new aircraft.

*SGI Origin 2000 is now marketed and sold as the SGI™ 2000 series.

Medtronic

[continued from page 3]

central SGI server using the unique SGI visual serving technology.

A similar effort was undertaken in recent years when Medtronic created a surgical simulator on an OpenGL® API-based system to train physicians on the delivery of a heart failure lead. Current simulator drawbacks, however, include the ability to use only a limited number of polygons to represent the model as well as issues involving reliability. High-performance visualization servers from SGI will be able to deliver the millions of polygons needed to accurately simulate placement of a lead anywhere in the heart.

Intrinsic Value

Using SGI hardware and a range of advanced modeling and simulation technology, Medtronic has increased the insight and the understanding of very scientific data. Now, people who may not understand the complexities of a product design as described in a white paper, or decision makers who may not realize its value in the marketplace, can go to a computer and gain valuable product insight through visualization tools that communicate the ideas in three-dimensional clarity. According to Abraham, "The most successful people within the company are regularly utilizing these technologies."

In summing up what SGI brings to the research and development table, Abraham says, "The most significant contribution that SGI provides for Medtronic is value...the kind that stands behind the computer, even more so than what's in the computer. SGI delivers not only the physical technology of the computer, but backs it with an intellectual and scientific support system as well as a hardware support system that is second to none. It's about much more than just cycles in a computer, it's about the partnership."

For more information on Medtronic's products, therapies, and services, visit its Web site at *www.medtronic.com*.

sgi™

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America (52) 5267-1387
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

www.sgi.com/manufacturing/eoi

3156 (12/2001) J13106

manufacturing