



SGI news



Quarterly Corporate Newsletter



Welcome to the second edition of the SGI Quarterly Corporate Newsletter. This newsletter is designed to showcase what SGI does best. In recent quarters SGI has sharpened its focus on the core competencies that have defined us—high-performance computing, complex data management, and visualization. These technology focal points represent our heritage and remain at the core of everything we do. As we celebrate our 20th year, we know that no one else in the world has been innovating the way we have over the past two decades, and we believe that now, more than ever, we provide a level of performance and innovation unequaled in the industry.

As you read this newsletter, you will see that SGI's customers are indeed diverse, including such markedly different organizations as manufacturing companies, government agencies, medical institutes, weather forecast centers, film studios, and geoscience companies. But despite their obvious differences, they all have one thing

in common: the need to compute and manage large amounts of complex data and then visualize the results in a realistic time frame and scale. In each of our targeted market segments we can point to substantial expertise and experience in providing these kinds of solutions to an installed base of world-class customers.

Moreover, there's widespread acknowledgment on the part of our customers that SGI stands alone at the forefront of innovation when it comes to providing this range of systems. In an increasingly digital world—with more complex data to compute, more concepts to visualize, and more information to manage—the ability to interpret and gain insight from large amounts of complex data translates into a world of opportunities for SGI and our customers.

Greg Estes
Vice President, Corporate Marketing

SGI and NASA: Extending the Boundaries of Supercomputing

The NASA Advanced Supercomputing [NAS] division, which is part of Ames Research Center, provides supercomputing capabilities for many of NASA's most important projects. From modeling the aerodynamics of potential space shuttle replacements to predicting Earth's climate far into the future, these critical projects encompass some of the most challenging computational problems ever undertaken.


The goal of NAS is to increase the use of supercomputers for important research by making them readily available, reliable, and easy to program and use, enabling scientists to focus on research, not computing. NAS is also bringing the benefits of supercomputing to areas of research where it has previously been underutilized.

According to Bill Feiereisen, NAS division chief, "NASA researchers depend on NAS to provide advanced supercomputing capabilities. Because reducing the time

required to execute critical software is essential to continued progress, getting the most from each budget dollar spent on hardware and programming is critically important."

In recent years NAS has turned increasingly to SGI to provide production supercomputers to meet its computational needs. For NASA workloads, shared-memory, single system image [SSI] SGI® Origin® family systems have proved superior to large, clustered supercomputers, providing far greater performance on NASA projects with less programming effort and at significantly less total expense. SGI and NASA collaborated to extend the limits of SGI Origin family systems, culminating in the delivery of the first 1,024-processor SGI® Origin® 3800 system in July 2001.

"The 1,024-processor Origin 3800 system allows us to accomplish in hours jobs that literally took months or [continued on page 2]



IN THIS EDITION

1 **Government**
*SGI and NASA:
Extending the
Boundaries of
Supercomputing*

4 **Sciences**
*Neurosurgeons
Perform Highly
Complex
Separation of
Conjoined Twins*

5 **Energy**
*Austin
GeoModeling
Helps Maximize
the Potential of
Existing Oil Fields*

7 **Manufacturing**
*Visualization
Benefits for the
Automotive Industry*

9 **Media**
*SGI Sweeps
Academy Award®
Visual Effects
Nominations for
Eighth Year*

[continued from page 1]

even years on previous systems. Because of the demonstrated linear scaling of NASA software on the SGI Origin family platform, we expect to achieve even greater things in the future. The computing capability offered by the Origin family platform will fundamentally change the way NASA carries out many of its most critical projects," said Jim Taft, co-director, Terascale applications group.


Replacing the "Gold Standard"

For years, NASA's gold standard for supercomputing was the Cray C90™, a vector supercomputer rated at 16 gigaflops. [A gigaflop is one billion floating-point operations per second.] The Cray C90 demonstrated about 25% efficiency running NASA workloads, delivering about 4 GFLOPS when executing OVERFLOW, one of NASA's most important software programs. When it became clear that follow-ons to the C-90 were not available, NAS began exploring other supercomputing solutions.

Over a period of years, a variety of clustered supercomputing solutions were tried from a number of different vendors, but—for NASA workloads—none of these came close to the performance of a Cray C-90, let alone exceeded it. Even though these systems looked great on paper, with theoretical performances of hundreds of gigaflops, they achieved efficiencies of only a few percent on important NASA software and failed to compete with the C-90s already in place. Because of low efficiency, clustered systems large enough to achieve acceptable performance on NASA workloads would be prohibitively expensive if they could be built at all.

NAS purchased its first SGI Origin family system—a 128-processor SGI® Origin® 2000 server—in 1996 and was immediately impressed with the results. After just a few months spent porting code, the Origin 2000 system exceeded the performance of a Cray C-90 on OVERFLOW. More important, NASA software scaled linearly on the Origin 2000 architecture. Performance increased predictably as processors were added.

NASA immediately entered into a memorandum of understanding with SGI to collaborate on the development of larger systems. A 256-processor Origin 2000 server soon followed. That system was fully operational within a few weeks of installation and could run OVERFLOW at 20 GFLOPS—five times the speed of a C-90. This system was later followed by a 512-processor Origin 2000 system and, most recently, the 1,024-processor SGI Origin 3800, which runs OVERFLOW 40 times faster than the C-90. Despite this impressive performance, the 1,024-processor Origin 3800 system cost NASA less than half as much as the C-90 did in the early 1990s.



SGI facilitated the purchase of the Origin family systems at NAS with special financing arrangements, enabling NAS to obtain equipment and begin important projects while spreading repayment across multiple funding appropriations.

SGI also established a special funding contract for NASA, simplifying the procurement process for all parties.

After a few months of operation, the Origin 3800 system is already achieving processing efficiencies approaching those seen with the C-90, proving it is possible to achieve high efficiencies on nonvector machines, something that has generally eluded application designers in the supercomputing world.

Building a Virtual Wind Tunnel

Historically, 20% of C-90 CPU cycles at NAS have been dedicated to running OVERFLOW and other important computational fluid dynamics [CFD] software. CFD is used to model the aerodynamics of advanced aircraft and space vehicles passing through Earth's atmosphere and is critical in the design and simulation of rocket motors.

The goal at NAS is to provide sufficient computing power to enable scientists and engineers to test their designs in a "virtual wind tunnel." Traditional wind tunnel testing is time consuming and expensive—requiring the construction or alteration of physical models for each test—and even the most powerful wind tunnel cannot simulate the conditions of the launch and reentry of space vehicles.

Until the arrival of the 1,024-processor SGI Origin 3800 system, the computing power needed to make the virtual wind tunnel a reality was unavailable. Complete modeling of one aircraft configuration during landing required up to a year on a C-90, imposing serious limits on the use of simulation. The Origin 3800 system can run the same configuration in a matter of hours, so, for the first time, scientists can routinely use simulation to validate their designs under varying conditions. "The Origin architecture has created a revolution in computational fluid dynamics at NASA and will fundamentally change the way aircraft are designed in the future," said Taft.

The ability to do advanced simulation has already proved its value in NASA's mission to design a new reusable launch vehicle to replace the space shuttle. During simu-

[continued on page 3]



[continued from page 2]

lation of the X-37 drone, designed to be dropped from the space shuttle to test reentry, a serious flaw was discovered that would have led to catastrophic failure. Millions of dollars and months of time were saved because of the advanced capabilities enabled by the Origin 3800 architecture.

Predicting Climate Change

In the study of climate change, NASA is reaping similar performance benefits from the 1,024-processor SGI Origin 3800 server. After a three-week porting effort on the climate model of NASA's next-generation climate modeling system, the model ran five times faster than on any other system ever tested. The Origin 3800 architecture again demonstrates superior scalability as more processors are used to execute the climate model. By comparison, clustered systems often reach a peak and then slow down as more processors are added.

With the success in enhancing the atmospheric model performance, NASA has moved on to ingest more satellite observations into the data assimilation system. The use of the 1,024-processor SGI Origin 3800 system with state-of-the-art assimilation will help improve our understanding of the cause and consequence of the change in Earth's climate system.

NUMAflex™: The Architectural Advantage

The 1,024-processor SGI Origin 3800 server is the largest system ever built with an SSI and shared memory. All processors access all system memory directly. This is in sharp contrast to clustered solutions in which a separate instance of the operating system is needed for every few processors and each processor has direct access to only a subset of total memory.

The patented SGI® NUMAflex architecture makes it possible to scale the number of processors well beyond the level that has been possible in other shared-memory designs. Each processing node has up to four processors and a local pool of up to 8GB of memory. Instead of the traditional backplane design, NUMAflex uses crossbar switches and high-speed cabling, allowing each node direct access to the memory in other nodes with a relatively slight increase in latency versus accesses to local memory [hence the designation NUMA—nonuniform memory access].

NUMAflex uses standard, modular building blocks called bricks that allow systems to scale independently in different dimensions over time, providing unprecedented levels of flexibility, resiliency, and investment protection. Various types of bricks can be added as

needed to tailor a system to the exact capabilities required by the application. As an added advantage, SGI® Origin® 3000 series systems provide an extremely small physical footprint relative to other comparable systems because of the efficient modularity of NUMAflex.

MLP: A Simpler Programming Model

SGI supports a variety of programming interfaces for use on Origin family servers, allowing scientists and engineers to choose the best method for tackling a particular problem. SGI designs each programming interface to take maximum advantage of the unique features of the shared-memory Origin architecture, making SGI the leader in shared-memory programming.

NASA programs the 1,024-processor SGI Origin 3800 system using Multi-Level Parallelism [MLP], a programming method developed at NAS specifically for use with shared-memory NUMA architectures. MLP uses both coarse-grained parallelism in the form of UNIX® forked processes and fine-grained parallelism in the form of lightweight OpenMP™ API threads. All communication between processes occurs via shared memory rather than message passing.

The message passing interface [MPI] is the standard programming model for clusters in which all communication occurs via messages passed between processors. [MPI is supported on SGI Origin family servers and has been implemented to take advantage of underlying shared memory to provide low latency and high bandwidth.] MPI is an appropriate programming model for various highly parallel problems, but for many of the problems tackled by NASA, MLP offers significant advantages over MPI.

For example, scientists have recently become interested in small protein molecules being studied for their pharmaceutical applications. Modeling the dynamics of proteins usually involves fewer than 50,000 atoms. On a 1,000-processor machine this means that each processor is responsible for no more than 50 individual atoms.

The performance of MPI on clustered systems when simulating such a system quickly degrades because of the large number of messages that must be passed to calculate interactions between each atom and other atoms not resident on the local processor. [Each message has latency typically in excess of five microseconds.] Because of the submicrosecond latencies of shared-memory access on

[continued on page 12]





Neurosurgeons Perform Highly Complex Separation of Conjoined Twins Using SGI® Technology

Separating conjoined twins is one of the most demanding and complex medical procedures. A team of doctors in Singapore separated twins joined at the head and brain thanks to rehearsals using the Dextroscope™, a neurosurgical planning system developed by Singapore-based Volume Interactions and powered by high-performance SGI® Onyx® family systems. The Dextroscope system transforms two-dimensional images of patients' bodies into three-dimensional graphics, allowing neurosurgeons to plan the best way to perform complex surgeries. The combination of the Dextroscope system and SGI technology has been used successfully in more than 40 operations at the Singapore National Neuroscience Institute.

Operation Lasts 97 Hours

The 97-hour operation to separate the Nepalese twins began on Friday afternoon, April 6, 2001, and was completed the following Tuesday morning. Jamuna and Ganga Shrestha were joined at the top of their heads and shared the same brain cavity. The separation surgery was complicated by the fact that their brains, partially fused, shared many overlapping blood vessels. Two teams of surgeons worked around the clock to successfully complete the operation.

Planning for the operation was conducted collaboratively on two different continents with two key surgeons simultaneously observing the same data—in this case, virtual versions of the babies' heads. With the help of the Dextroscope system, Dr. Keith Goh, co-leader of the surgery team in Singapore, was able to discuss complex questions with Dr. Benjamin Carson, a surgeon based at Johns Hopkins University in Baltimore, Maryland.

The neurosurgeons utilized the Dextroscope system to explore a virtual world of surgical scenarios well in advance of the actual operations. They manipulated the 3D visualization system with stereoscopic goggles, a toggle bar, and a stylus. The twins' biological features were viewed from different angles and "sliced" into with a virtual scalpel that behaves exactly like a real one.

Harnessing Technology

"SGI is proud to be powering Dextroscope," said Lawrence Lee, general manager, SGI Singapore. "The application helps neurosurgeons to make better decisions on surgical procedures, resulting in greater successes. This is in line with our corporate objective to harness 3D graphical technology for the benefit of humankind." Dr. Luis Serra, president and chief technical officer of Volume Interactions, was equally impressed with the systems' performance. "Since the application involved huge volumes of data, we needed sophisticated hardware to render the interactive 3D graphics. The SGI Onyx family of systems was the best platform available. With SGI hardware advances providing more texture memory, we should be able to deliver better, more accurate, and complete data in our solutions."

Bracco S.p.A., a chemical-pharmaceutical and biomedical group, and a world leader in the diagnostic imaging sector, recognized Volume Interactions' strength in visualization and 3D interaction and recently acquired a majority stake in the company.

Going Home

On November 16, 2001, the twins were released from Singapore General Hospital and flown back to Kathmandu, Nepal. Dozens of well-wishers, doctors, and nurses said goodbye to the twins at Singapore's Changi Airport. Having been captivated by the twins' plight, the people of Singapore had donated hundreds of thousands of dollars to pay their medical expenses. Doctors performed the surgery for free, and Singapore Airlines paid for travel expenses.

The girls are staying in Kathmandu in order to be close to a major hospital for ongoing treatments. Their parents and grandparents will also stay in the capital city rather than return to their home village of Khalanga, which is a 24-hour bus ride from Kathmandu.

[continued on page 12]



SGI Technology and Austin GeoModeling Help Maximize the Potential of Existing Oil Fields

Fossil fuels are a finite resource. As the cost and difficulty of developing new oil and gas fields in ever more remote regions of the globe continues to grow, getting the most from existing fields becomes increasingly important.

Austin GeoModeling, Inc. [AGM], has been developing software and providing consulting services to help geoscientists maximize resource recovery from existing oil and gas fields for the past five years, bringing the advantages of 3D visualization and interpretation capabilities to geologists' desktops. This technology is helping AGM customers stay competitive by enabling them to get the most from their investments in existing fields.

SGI has been an important partner for AGM in these endeavors, providing technical assistance and support to help AGM expand the market for its products and services. According to Robin Domisse, chief executive officer of AGM, "As a small software vendor in a large and complicated worldwide market, the support and assistance of SGI has been


critical to our success, allowing us to increase the quality and scope of our product."

Real-Time Interpretation to Maximize Oil Field Potential


The likely locations of fossil fuel reservoirs are determined based on seismic analysis, providing an initial interpretation of geological features below the earth's surface. After a well has been drilled, logging tools measure various rock and fluid properties in the layers that the drill has penetrated. Integrating well log data with seismic data dramatically increases the understanding of the reservoir, providing a much more detailed picture of the underlying geology.

AGM developed its flagship application, Recon, to maximize the utility of well log data. Using Recon, geoscientists can correlate the data from hundreds of wells to solve the most complex problems of geological interpretation, producing a high-resolution, 3D image of the field under study. Recon provides three coupled views: a base map view showing a surface map of well locations, a cross-sectional view showing the logs for various wells in a cross section of the field, and a 3D view showing underlying strata. Changes made in one view are automatically reflected in the other two.

[continued on page 6]



Recon integrates well log data, seismic horizons, and production data in an integrated interpretation environment.



[continued from page 5]

Recon was developed using OpenGL®, the premier environment for developing portable, interactive 2D and 3D graphics applications. AGM's strong commitment to OpenGL contributes greatly to the overall performance and feature set of Recon. Originally developed by SGI, OpenGL has now become an industry standard.

With Recon, geoscientists can cycle through an interpretation workflow faster and more frequently, increasing the accuracy of the interpretation. These timesavings reduce the risk of drilling decisions and allow real-time updates to the geological model as new wells are drilled.

Because of its advanced, coupled visualization and interpretation capabilities, Recon makes it easier to maximize the potential of existing fields. Recon has increased the value of every field where it has been applied.


Unparalleled Services for Reservoir Characterization

In addition to direct sales of Recon, AGM provides advanced consulting services. The expertise of AGM professionals has been used in more than 20 reservoir characterization studies worldwide.

In one important study Recon was used to create a 3D framework for Saudi Aramco's largest reservoir in the Ghawar field of Saudi Arabia. Combining four study regions into a single Recon project, a new sequence-stratigraphic framework was created, consistent across the entire 120-mile-long field containing 3,000 wells.

On the second interpretation cycle, sequence-stratigraphic clinoform structures were identified, validating a completely new interpretation of the subsurface geology. Within days field-wide correlations involving hundreds of wells were revised to incorporate this discovery—the first new geological interpretation of the world's biggest oil field in 40 years.

In another study AGM used Recon to incorporate 3D production history data with sequence-stratigraphic correlations in a large, mature West Texas field, leading to a more accurate geological model. The results were validated by a successful recompletion [a process to increase the production of existing wells] and infill drilling program that increased the value of the field by 670%.



“SGI technical personnel have incomparable expertise in the development of visual applications. Over the past five years AGM has received tremendous support from SGI to improve Recon and market it successfully. We simply couldn't ask for a better partner than SGI.”


—Robin Dommissé, Chief Executive Officer of AGM

SGI: A Partner for Success

From the beginning AGM has depended on SGI hardware for its unparalleled visualization capabilities and data throughput. The analysis of a large field begins with several megabytes of well log data for each well. This data can be complemented with production data for each well, as well as 3D seismic data sets for the region. The resulting data set can be extremely large and difficult to manage. Because of SGI's incomparable ability to manage large visual data sets, the combination of Recon and SGI systems is a natural one for AGM.

As a small software developer in a big market, AGM faces significant challenges. With SGI's assistance, AGM is now able to offer Recon on the full line of SGI visual workstations, as well as Silicon Graphics® Onyx2®, SGI® Onyx® 300, and SGI® Onyx® 3000 series visualization systems for the ultimate in visual realism. AGM also has recently been able to showcase Recon's unique capabilities in SGI® Reality Center™ facilities with a demonstration showing real-time interpretation using thousands of well logs.

“SGI technical personnel have incomparable expertise in the development of visual applications. Over the past five years AGM has received tremendous support from SGI to improve Recon and market it successfully. We simply couldn't ask for a better partner than SGI,” said Dommissé.



Visualization Benefits for the Automotive Industry



Leaders in the automotive industry have advanced beyond conventional CAD modeling and are using 3D visualization throughout the complete vehicle development process. Modifications to components, assemblies, and even fully integrated vehicle systems can be made without committing to a final product specification until the very last minute.

Several major automotive companies, including Renault, DaimlerChrysler, and Toyota, continue to report success with Silicon Graphics Onyx2 systems for visualization applications that make virtual product design a valuable alternative to costly physical prototypes.

This article features examples from these automotive companies in diverse areas such as styling, occupant safety, and manufacturing, where the SGI visualization environment enabled an interdisciplinary and collaborative framework that significantly reduced time, improved design quality, and provided cost savings.



Competitive Styling

At the Mercedes-Benz car development facility in Sindelfingen, Germany, an SGI Reality Center facility allows designers at DaimlerChrysler to introduce new concept vehicles to their colleagues in a 3D simulated space. New design concepts can be examined as they operate under everyday traffic conditions or compared side by side with competitive vehicle designs.

“Simulations and visualization allow us to examine a greater number of variants in a shorter period of time and at lower costs,” said Hans Joachim Schopf, chief engineer for Mercedes-Benz passenger cars. “As a result, we can rule out nonviable options more quickly.”

Imperfections and inconsistencies in the shape of vehicles can be detected early in the process, enabling many key design decisions to be made without the development of expensive clay modeling.

Virtual Crash Testing

Virtual reality technology not only gives DaimlerChrysler a significant return on investment in cost, time savings, and product quality, but also reduces material waste and environmental impact while at the same time allows manufacturers to improve occupant safety and crashworthiness.

“We can work through a lot more variations than is possible in real tests without having to destroy a Mercedes every time,” said Bharat Balasubramanian, a director at Mercedes-Benz car development. With a Silicon Graphics Onyx2 InfiniteReality2™ system configured with 60 CPUs, 15GB of main memory, and 14 independent graphical high-performance pipes, data-rich, full-scale crash tests can be simulated and visualized.

Virtual crash tests allow a more detailed examination from inside the vehicle than is possible in a physical crash test. For instance, overlapping components or assemblies can be visually “removed” to reveal how concealed components were affected in the crash. “From inside the vehicle we can see considerably more on a virtual level than we could show in a real crash using the most advanced high-speed camera,” Balasubramanian said.

[continued on page 8]

[continued from page 7]

Design Collaboration

Visualization also enables increased collaboration between different disciplines. For example, if design conflicts arise between styling and the manufacturability of a particular body feature, with access to a lifelike model in a team discussion these can be resolved quickly. Visualization allows an investigation of the aerodynamic effects for certain body shapes and the ability to manufacture the body panels for those shapes. Collaboration means time savings. Using visualization to explore different design and engineering options together, engineers at Renault said the time required to review, resolve, and incorporate a change to a powertrain component, such as an engine cylinder head, was reduced by as much as 55%.

Improved Manufacturability

At Toyota, visualization is improving communications on the assembly floor. Using a video conferencing system combined with ENVISION® software from DELMIA, a division of Dassault Systemes, designers and assembly line managers can look at simulations of individual components as they are being virtually installed into a vehicle during assembly. Engineers and factory workers in Japan and overseas can see and talk to each other while simultaneously viewing the same information.

Before introducing this Visual and Virtual Communication [V-Comm] system into the manufacturing process, different factories relied on faxes to announce assembly procedure changes. As a result, changes sometimes took two months to be uniformly implemented. With V-Comm, the information is reviewed and discussed in real time. Suggestions and feedback from floor workers are received and implemented long before the vehicle moves into the mass production.

“Simulations and visualization allow us to examine a greater number of variants in a shorter period of time and at lower costs. As a result, we can rule out nonviable options more quickly.”

—Hans Joachim Schopf, Chief Engineer for Mercedes-Benz passenger cars

These are just a few examples of how advanced visualization technology from SGI has helped car companies bring new, better-quality models to market faster, at lower cost, and with more options for their customers. The way forward for the automotive industry will require a technology environment that efficiently manages increasing complexity in the vehicle development process. SGI continues to invest in visualization and other technologies to provide the automotive industry with leadership required for meeting the challenges ahead.



SGI Sweeps Academy Award Visual Effects Nominations for Eighth Year in a Row



Artificial Intelligence: A.I.

For the eighth consecutive year all films nominated for an Academy Award for Distinguished Achievement in Visual Effects relied on the power of SGI high-performance computer graphics systems to create the year's most unforgettable images.

The amazing worlds and cyborgs of *Artificial Intelligence: A.I.* and the harrowing air and sea battles of *Pearl Harbor* were made possible by the creative artistry of Industrial Light + Magic (ILM), a division of Lucas Digital Ltd. The Academy of Motion Picture Arts and Sciences (AMPAS) also nominated Weta Digital, a New Zealand company that built a state-of-the-art production facility based on SGI systems, to produce the film trilogy

of *The Lord of the Rings*. The first film of the series, *The Lord of the Rings: The Fellowship of the Ring*, was not only nominated for visual effects but also garnered 12 other Academy Award nominations, including Best Picture.

In addition, two films nominated in the brand new AMPAS category of Best Animated Feature Film were also created using SGI visual workstations and servers: Disney/Pixar Animation Studios' *Monsters, Inc.*, and PDI/Dreamworks' *Shrek*, both critically acclaimed for advancing physical and emotional realism in computer animation.

Artificial Intelligence: A.I. and Pearl Harbor

SGI IRIX® OS-based compute power provided the multi-Academy Award-winning artists of ILM with the means to create, design, animate, photo-realistically light, and render the amazing array of visual effects that enhance two of the nominated films. Using proprietary software written on Silicon Graphics® O2® visual workstations as well as off-the-shelf 3D modeling and animation software, ILM
[continued on page 10]



The Lord of the Rings: The Fellowship of the Ring

[continued from page 9]

completed 200 shots on *A.I.* and 160 shots on *Pearl Harbor*. ILM, which has used SGI systems since 1988, uses more than 500 Silicon Graphics O2 visual workstations networked to an 800-processor SGI® 2000 series system with 400GB of storage.

ILM artists working on *A.I.* used a variety of software from Alias|Wavefront, an SGI company, relying heavily on PowerAnimator™ software for modeling, Maya® software for procedural animation, and Avid® Softimage® software primarily for animation, supplemented with a large amount of ILM's proprietary software—all running on Silicon Graphics O2 systems.

"*A.I.* had one of the broadest ranges of visual effects that I've ever been exposed to," said Doug Smythe, ILM associate visual effects supervisor for *A.I.* "We did literally every kind of shot you can think of, including simple wire removal, blink removal, and digital creature work of all kinds; digital enhancement of live action; all the CG eyeballs and prosthetics in the *Flesh Fair* shantytown; complete digital environments such as the excavation; combinations of miniatures with live action and computer graphics, such as *Rouge City*; and, of course, the furry teddy bear.

ILM's visual effects work for *Pearl Harbor* included three main sequences: the aerial tour de force of the Battle of Britain; the devastating Pearl Harbor attack (including populating the sky with CGI planes, the 3D all-digital ships in Battleship Row, the water they float on, and a library of digital sailors); and many shots in the Tokyo raid sequence at the movie's climax.

Using Silicon Graphics O2 systems, artists adapted ILM's proprietary fluid dynamics, originally written on O2 systems for last year's Academy Award nominee for visual effects *The Perfect Storm*. For the volumes of smoke in *Pearl Harbor*, ILM adapted its fluid dynamic volumetric software to generate very complex smoke plumes, which



Pearl Harbor

filled the battle sequences. All the procedural animation tools were written in Alias|Wavefront™ Maya, which also handled the crowd placement for the digital sailors. ILM proprietary software, written on Silicon Graphics O2 systems, was also used for crash dynamics, lighting, and rendering.

"All of our code was written in the software we developed for the SGI O2 systems. They're great," said Michael Bauer, ILM's CG supervisor on *Pearl Harbor*. "They are integral to our work."

The Lord of the Rings: The Fellowship of the Ring

Weta Digital has been working on the *Lord of the Rings* film trilogy for more than four years, producing more than 1,200 visual effects shots. The films each require dozens of digital human and humanoid characters plus lead creatures who are entirely CGI-created, vast CGI landscapes, battle scenes with hundreds of thousands of animated characters, and more special effects than you can shake a wizard's staff at.

In excess of 150 artists, keyframe animators, modelers, digital paint artists, motion editors, composers, and numerous software engineers were provided with 150 Silicon Graphics® Octane® visual workstations, which run Alias|Wavefront Maya as the facility's core 3D application; an eight-processor Silicon Graphics Onyx2 system running Discreet inferno for compositing; and 80 dual-processor Silicon Graphics® 330 and Silicon Graphics® 230 workstation based on Linux®, which are used for a combination of paint, rotoscoping, and compositing duties. Two SGI file servers using SGI® TP9400 storage arrays, StorageTek Tape Robots, and SGI® 2000 series server technology provide a combination of 4TB of online storage and more than 20TB of nearline storage as a global storage repository to support workstation information sharing.

"What we're about is the ability to move large amounts of information around the facility all day, every day, and we rely on SGI to help us do that. Ninety percent of our equipment is SGI," said Jon Labrie, chief technical officer of Weta Digital. For nearline and offline storage Weta Digital uses DMF, the SGI hierarchical storage management system, which is already managing 50TB of information. "SGI DMF has greatly simplified our management of the thousands of tapes needed to store the bulk of the data," Labrie added.

From the beginning of preproduction, Weta Digital also used the Octane visual workstations to write extensions to Maya, creating proprietary technology including Massive, a [continued on page 11]

[continued from page 10]

custom-built crowd animation or “artificial ecology” system developed on IRIX and now ported to Linux, which draws from a huge database of motion-capture data.

“We’re using Massive for battle animation scenes with hundreds of thousands of fighting, screaming, and dying orcs, elves, and all the other magical and fantastical creatures that appear in *The Lord of the Rings*,” Labrie continued. “For these sorts of graphical challenges we prefer to work in the world of IRIX and UNIX. The graphics engines available to us on the SGI platforms make our jobs easier. We are so thrilled with the latest increase in performance and quality with the new Octane2 system that we’ve enlisted it for our next production.”

Weta’s primary rendering resource is based on SGI® I100 and SGI® I200 Linux OS-based servers. Having started with 32 dedicated processors, the facility currently runs 192 dual-processor SGI servers rendering frames 24 hours a day, seven days a week.

Monsters, Inc.

SGI’s longtime customer Pixar Animation Studios created a monster box office hit, Disney/Pixar’s *Monsters, Inc.*, using Silicon Graphics® Octane2™ IRIX OS-based visual workstations. In preparation for the production of *Monsters, Inc.*, Pixar Animation Studios purchased 250 Silicon Graphics Octane2 visual workstations in July 2000, barely a month after the powerful new line of high-performance visual compute workstations was introduced. Pixar used SGI workstations for interactive graphics applications throughout the studio, including 3D modeling, painting, compositing, and animation. Silicon Graphics Octane and Octane2 workstations were also used to create and run new lighting tools designed for the film.

Building the final character models for *Monsters, Inc.* was a difficult and complex task, performed using the Silicon Graphics Octane2 workstations. “Sulley [the monster voiced by John Goodman] is covered with shaggy blue fur consisting of millions of individual hairs,” said Darwyn Peachey, vice president of technology, Pixar Animation Studios. “All that hair was positioned and styled interactively on the Octane2 workstations, and the motion of the hair was computed using our proprietary algorithms for dynamics simulation.”

In addition to the 250 Silicon Graphics Octane2 visual workstations, Pixar Animation Studios uses a Silicon Graphics® Onyx® system with InfiniteReality® graphics for high-resolution display and playback.



Shrek

Shrek

DreamWorks’ latest all-CG animated feature, *Shrek*—entirely created with SGI technology—has been critically hailed as a high-water mark in human character animation, both in movement and the translucent quality of the skin. To bring some of the film’s most memorable “human” characters—Princess Fiona and Lord Farquaad—as well as Shrek and Donkey to life the talented artists at PDI/DreamWorks used mostly Silicon Graphics O2 desktop workstations, the same machines they used for *Antz*.

The best thing about relying on SGI technology “is the solidity of the overall platform, which has really made SGI workstations attractive to us over the years,” said Dan Wexler, a member of PDI/DreamWorks’ R&D team. “The most powerful features of the SGI systems are their robustness, their stability, and the total solution that they provide in terms of being able to integrate three-dimensional graphics, video, and audio.” Although *Shrek* relied on a great deal of traditional SGI IRIX hardware, it was also the first DreamWorks production to embrace Linux technology. Linux OS-compatible platforms from SGI helped make PDI/DreamWorks’ transition as painless as possible. “We’ve traditionally been on IRIX operating systems and have used a wide variety of SGI machines for a long time,” Wexler reiterated. “Our SGI machines gave us an advantage in terms of being able to port software to Linux before the work on *Shrek* started.”

Rendering *Shrek*’s complex imagery was also a joint hardware effort. “We have a whole bunch of SGI® Origin® 200 servers, but the majority of the render farm on *Shrek* was in fact based on the Linux OS,” DreamWorks’ head of technology, Ed Leonard, states. “We purchased 168 SGI I200 2U dual-processor Linux OS-based machines. Ultimately our decision to use Linux [or any other technology] is based on its ability to bring more horsepower to the making of great films, and SGI certainly played a key role there—and in the making of *Shrek*.”


Neurosurgeons

[continued from page 4]

Singapore has been growing as a center for medical and biotechnology research, and the success of the operation has boosted the city-state's reputation as a center for progressive medical treatment.

SGI Contributions

SGI continues to make groundbreaking contributions to the health care and medical fields. Our solutions include the hardware for computer-aided, computer-guided surgeries and simulations such as that performed on the twins, as well as high-performance graphics systems for diagnostic imaging devices, medical image management, and picture archiving/communications solutions.



The Dextroscope technology is one example of the many ways that SGI technologies facilitate virtual and enhanced visual procedures for improved preoperative planning and simulations. As a result, the need for more invasive diagnostic and surgical procedures is reduced. In addition to neurosurgery, the Dextroscope can be used in other industries that would benefit from the visualization of three-dimensional multimodal data sets, enhanced interaction, and the capacity to support planning. These industries include geophysical exploration, drug design, research, microchip defect inspection, and cell biology visualization.

SGI and NASA

[continued from page 3]

Origin family servers, the Origin 3800 platform with MLP runs small atom count problems faster than any other system on the planet.


For NASA, MLP has proved significantly easier to program than MPI, in which the programmer is typically responsible for analyzing and decomposing the problem into fine-grained parallel threads. With MLP, the compiler is capable of generating most of the parallel code with limited user directives. MLP also permits simple and highly efficient dynamic load balancing. Work can be shifted quickly and automatically between CPUs in response to load, providing better efficiency and therefore better performance.

NASA has also found that its legacy programs are much easier to port to MLP than to MPI. Relatively few code changes are

necessary to port vector applications to MLP, and applications are far easier to debug than MPI programs. This has translated into a cost reduction of millions of dollars for NASA, adding to the price/performance benefit of the SGI Origin 3000 series architecture. For NASA, MLP and Origin family systems have proved a big win in application performance, programming simplicity, and overall expense.

A Future So Bright

Although NASA and SGI are excited by the new computing capabilities incorporated in the 1,024-processor SGI Origin 3800 system, they are far from ready to rest on their laurels. Plans are already in the works to scale the SGI NUMAflex architecture to even higher processor counts without sacrificing the benefits of SSI and shared memory.



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

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



©2002 Silicon Graphics, Inc. All rights reserved. Specifications subject to change without notice. Silicon Graphics, SGI, Origin, IRIX, Onyx, Onyx2, Octane, O2, OpenGL, InfiniteReality, and the SGI logo are registered trademarks and NUMAflex, OpenMP, Reality Center, Octane2, and InfiniteReality2 are trademarks of Silicon Graphics, Inc., in the U.S. and/or other countries worldwide. Linux is a registered trademark of Linus Torvalds. UNIX is a registered trademark of The Open Group in the U.S. and other countries. Maya is a registered trademark of Silicon Graphics, Inc., exclusively used by Alias|Wavefront, a division of Silicon Graphics, Limited, and Alias|Wavefront and PowerAnimator are trademarks of Alias|Wavefront, a division of Silicon Graphics, Limited. All other trademarks mentioned herein are the property of their respective owners. Artificial Intelligence: A.I.—©2001 Warner Bros. & DreamWorks LLC. All Rights Reserved. Photo courtesy of Industrial Light + Magic. Pearl Harbor—©2001 Touchstone Pictures and Jerry Bruckheimer, Inc. All Rights Reserved. Photo courtesy of Industrial Light + Magic. Shrek—©2001 PDI/DreamWorks. Photo courtesy PDI/DreamWorks. Monsters, Inc.—©2001 Disney Enterprises, Inc./Pixar Animation Studios. All Rights Reserved. The Lord of the Rings—©2001 New Line Productions, Inc. All rights reserved. NASA images courtesy of NASA. Medical images courtesy of Bracco S.p.A and Volume Interactions. Oil drilling images courtesy of Austin GeoModeling. Virtual mechanic and assembly images courtesy of DELMIA. Virtual car image courtesy of DaimlerChrysler.