sgi™

The Academy Awards:
Visualizing the Impossible

the magazine of imagination, innovation, and insight

synergi

Issue 3—Spring 2001
U.S./Canada Edition

What is Artificial Intelligence? Can a robot play soccer? Discover the technology behind the fully autonomous humanoid robot players that are competing for the RoboCup Federation on page 13.

Contents

Executive Editor
Andrew Spybey
andrews@sgi.com

Editor-in-Chief
Milva Naguib
milvan@sgi.com

U.S./Canada Editor
Mar Garcia
mgarcia@sgi.com

Contributors
Grant Ellis, Jennifer Hand, Yvonne Kennedy, Ron Magid, Angus Malcom

Design/Production
Splash of Paint (U.K.)

Repro/Printing
Acorn Press Ltd. (U.K.)

Corporate Headquarters
SGI
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
USA
Phone: (650) 960-1980

Welcome to the spring edition of SYNERGI, the magazine of imagination, innovation, and insight from SGI.

The cornerstone of SGI's strategy for Linux involves enabling its hardware, software, and services to run on Linux; developing strategic alliances; and contributing to the open-source community to help accelerate Linux in the marketplace. As a result, SGI is considered by many within the Linux community to be one of the strongest supporters of the open-source movement.

In the News section, we uncover the new SGI 1100 server, the latest addition to the SGI 1000 series. We also present the new products and solutions from SGI introduced at NAB 2001.

When power problems strike, computers can cause major process problems. Turn to page 8 to find out how to protect your data with a simple solution.

Interested in virtual time travel? Hold on to your seat and get ready to experience the world's largest virtual reality theater, powered by SGI. To learn about this theater, turn to page 17.

SGI and its partners have once again swept the Academy Award nominations in the area of visual effects. For seven years straight, every Academy Award finalist has relied on SGI systems to create award-winning films. Find out which movies were nominated and how the special effects were created on page 20.

In the Wired Up section we delve into the world of Media Commerce, where digital content is created, managed, delivered, and transacted over a vast range of networks. It's the world of the future that is transforming marketplaces today and making new products and new ways of doing business possible.

Also in this issue: the next big thing in artificial intelligence; multidisciplinary design optimization; and the birth of the online photo album.

I hope you enjoy reading this latest edition of SYNERGI. If you'd like to see particular topics covered in future issues, send your comments to *synergi_us@sgi.com*.

Mar Garcia
U.S./Canada Editor
SYNERGI

Investing In *Linux*®

SGI's ongoing commitment to the Linux platform and the open-source community was reinforced at the recent LinuxWorld conference in New York.

The cornerstone of SGI's strategy for Linux involves enabling its hardware, software, and services to run on Linux, developing strategic alliances, and contributing to the open-source community to help accelerate Linux in the marketplace. As a result, SGI is considered by many within the Linux community to be one of the strongest supporters of the open-source movement. The company has consistently demonstrated its commitment to moving Linux forward through the development of its IRIX® operating system and by applying its IRIX innovations to enhance its Linux solutions.

At the recent LinuxWorld Conference in New York, SGI launched a suite of Linux OS-based solutions for clustering, Internet, storage, and database applications.

These applications include SGI Internet Server™ for e-commerce, SGI™ Internet Server Environment (ISE), SGI™ Advanced Cluster Environment (ACE) for Linux, and the SGI™ 1100 server. This new suite of SGI solutions is targeted toward technical, creative, and business professionals in the Internet, sciences, manufacturing, government, and education markets.

SGI 1100 Server
The 1100 server is a new addition to the Intel® processor–based, rack–mount–optimized server family from SGI. Certified to run Windows NT®, Windows® 2000, or Linux, the SGI 1100 server is a powerful yet affordable solution designed for easy manageability, maintainability, and expandability and can be deployed as an Internet, e-commerce, rendering, or cluster computing nodes solution. See page 6 for more details.

SGI Internet Server for E-commerce
A powerful, thin 1U serving solution complete with industry-leading SoftCart® e-commerce software from Mercantec that enables up to 150 storefronts per server, SGI Internet Server for e-commerce is designed for service providers. It provides a completely integrated, fully configured Linux OS–based solution that allows the providers' merchant customers to build attractive, user-friendly virtual storefronts for Internet retailing.

The server offers full integration capabilities with secure payment systems and back-end commercial applications. Moreover, merchants can also take advantage of better inventory tracking with electronic data interface message formatting. And, with integrated support for multiple virtual stores, SGI Internet Server provides exceptional ROI within minimal rack space.

SGI Internet Server Environment (ISE)
Designed to enable service providers to integrate new solutions into existing infrastructures quickly and affordably, ISE includes advanced management, monitoring, and security tools with powerful integrated installation scripts and a Web-based user interface—all on a single CD. ISE provides basic Internet infrastructure and a foundation to support the messaging, Web-serving, and e-commerce applications typically required by service providers. A powerful, Internet-serving software solution, the Linux OS–based ISE is qualified to run on the SGI 1100 server.

Kasenna™ MediaBase for Linux Servers from SGI
Kasenna MediaBase is the only carrier-class broadband software platform for managing, distributing, and streaming rich-media content over the Internet, intranets, and virtual private networks at bit rates ranging from 28.8Kb per second to 15Mb per second to multiple client platforms. In combination with SGI Internet Server, Kasenna MediaBase supports multiple streaming formats simultaneously, including RealVideo®, RealAudio®, QuickTime™, MPEG-1, MPEG-2, and MP3.

But SGI offers more than solutions: with more than 200 trained Linux service engineers, it is uniquely positioned to provide the services that are integral to a complete Linux solution. Linux education and training courses are offered in the format that meets customers' needs, and the company offers the broadest range of support in the industry, from Mission-Critical Support to single-incident WebPacks that can be purchased and used over the Internet.

For more information on the range of Linux products from SGI, visit *www.sgi.com/linux/*

Developing *Strategic Alliances*

SGI has formed a strategic alliance with Turbolinux, Inc., a leader in high-performance Linux OS–based solutions.

The alliance agreement, which includes an SGI equity investment in San Francisco-based Turbolinux, means that SGI and Turbolinux will work closely in a number of areas, including engineering, distribution, business development, and marketing.

The two companies are confident that collaboration in these strategic areas will significantly advance Linux for high-end computing, and they are also discussing joint research and development projects, including high-performance computing and management software for high-density computing.

The alliance builds on existing collaboration between the two companies. During the past year, SGI and Turbolinux have worked cooperatively on the Itanium™ for Linux project, and at the LinuxWorld Conference & Expo in San Jose in August 2000, the two companies announced an agreement to incorporate breakthrough SGI Pro64™ compiler technology into the upcoming commercial release of Turbolinux®, for Itanium.

Both SGI and Turbolinux have aggressive Itanium roadmaps with proven expertise in this area, and the incorporation of SGI compiler technology into the new Turbolinux release has substantially increased Itanium performance on Linux.

For more information:
www.turbolinux.com/
www.sgi.com/linux/

Minimum Footprint,

Maximum Performance

The new SGI 1100 server is the latest addition to the SGI 1000 series, a comprehensive and cost-effective family of servers that merges SGI's unrivaled expertise and innovation in scalability.

The SGI 1100 server sets the standard in compute performance per square foot. Packing maximum performance in minimal space, this thin and powerful server excels in application environments requiring compute power in a space-constrained area.

The optimized design of the SGI 1100 server delivers high performance for both system and critical application needs. The server is highly scalable, supporting up to two high-performance Intel Pentium® III 800 MHz or 1 GHz processors and up to 2GB of SDRAM. Customers can match or expand performance capabilities as needed to track changing application requirements. The new system also features one or two IDE hard-disk drives and a 64-bit, 66 MHz, full-length PCI slot in an ultracompact 1U chassis, as well as two integrated Ethernet LAN controllers for fast setup in any network. All these features make the SGI 1100 server ideal as an Internet, digital publishing, or rendering server or as a compute node in a cluster environment.

The new SGI 1100 server is shipped with Linux and certified to support Windows NT and Windows 2000, which allows users to run the operating system that best suits their solution. High-performance I/O options include Myrinet, external Fibre Channel RAID, Gigabit Ethernet, and Ultra2 SCSI. These choices are very valuable, since they allow matching system I/O performance to application demands.

The SGI 1100 server features remote management capability (including remote monitoring, alert, console, and viewing), an emergency management port (including remote power on/off/reset, remote diagnostic and BIOS, and event paging capabilities), plus easy-access internals, making it extremely simple to service and manage.

Technical computing professionals will appreciate the affordability and power-packed design of the SGI 1100 server as a space-saving solution for a broad range of technical and creative applications such as electronic design automation, bioinformatics, rendering, digital publishing, Internet infrastructure, Web serving, and e-commerce. Customers that deploy the SGI 1100 server can easily rack and stack these affordable, ultradensity servers to better utilize valuable real estate and built-in networking while providing critical data protection.

"With this announcement, SGI leverages its expertise in high-end technical computing, system optimization, and flexible industry solutions that scale as our customers' application requirements and businesses grow," said Greg Goelz, vice president of marketing for volume products, SGI. "SGI extends this world-class ability to the SGI family of industry-standard Intel architecture–based servers, bringing high-performance features and functionality to mainstream users."

For more information:
www.sgi.com/servers/1100/

New

Media Commerce™ Solutions

At the NAB 2001 conference held in Las Vegas, Nevada, April 23–26, SGI demonstrated its Media Commerce solutions that enable customers to create, manage, deliver, and transact digital content over a vast range of networks.

SGI introduced new DMediaPro™, digital asset management, and media serving products that range from the Silicon Graphics Zx10™ and Silicon Graphics® Octane2™ visual workstation product lines all the way up to the SGI™ Origin™ 3000 series servers and SGI™ Onyx® 3000 series visualization systems. SGI also highlighted media streaming, interactive television, and high-definition solutions. In addition, Kasenna MediaBase Enterprise and Network Editions for SGI IRIX® OS–based servers made their debut.

The following new products were presented.

SGI Media Server™ for Broadcast, New MPEG-2 Model
SGI Media Server for broadcast is a completely integrated solution for serving video, audio, and data that is based on open video, data networking, file transfer format, and storage technology. The new MPEG-2 model of SGI Media Server for broadcast takes advantage of the revolutionary distribute data, view video approach—managing video as data and distributing it over data networks. The high throughput of the server facilitates simultaneous ingest, data network–based file transfers to the server, and playout to air of multiple video channels. SGI Media Server is a powerful and versatile solution for acquisition, commercial insertion, spot playback, and distribution between and within networked facilities. The server can also interface with browsing and asset management systems—streamlining the process of repurposing content for Web and broadband distribution.

New Silicon Graphics Zx10 Visual Workstation with DMediaPro™ DM1
SGI introduced the new video-enabled configuration of the Silicon Graphics Zx10 VE visual workstation. This system supports uncompressed 8- or 10-bit YUV (with alpha) standard-definition video input and output. This digital media solution provides the only integrated Windows OS–based solution that offers groundbreaking price/performance capabilities for virtual sets, virtual actors, and real-time graphics applications. This will be sold as a bundled solution, with two different configurations: rack-mount chassis and tower chassis.

New Silicon Graphics Octane2 Visual Workstation with DMediaPro™ DM2
The Silicon Graphics Octane2 visual workstation with DMediaPro DM2 delivers unrivaled video quality, versatility, and performance. The graphics quality and system bandwidth of Octane2 plus the high-definition (HD) capabilities of DM2 create today's most powerful desktop HD solution for visual effects, editing, and compositing. The first desktop workstation to feature uncompressed high-definition 8-bit RGB 4:4:4, this dual-definition (HD/SD) option will be sold as a production-ready bundle with Octane2, VPro™ V12 graphics, and Video Breakout Box.

New SGI Onyx 3000 Series and SGI Origin 3000 Series Systems with DMediaPro™ DM3
The new DM3 option for the SGI Onyx 3000 series and SGI Origin 3000 series supports HD and SD uncompressed video in and out with real-time colorspace conversion. The SGI Onyx 3000 series and SGI Origin 3000 series systems with DMediaPro DM3 take advantage of the scalable computing and high-bandwidth capabilities of the SGI™ NUMA architecture to deliver state-of-the-art film mastering, virtual sets, and multistream compositing applications. The new High-Definition Graphics-to-Video Output (HD GVO) for the SGI Onyx 3000 series also supports high-definition and standard-definition video out from InfiniteReality3™ graphics and is ideal for real-time graphics and virtual set applications.

SGI StudioCentral™ Library 3.0
StudioCentral Library 3.0 is the first digital asset management system for enterprise-wide deployment. It immediately benefits any organization wanting to better manage the creation and delivery of content. Unlike other out-of-the-box solutions, StudioCentral Library is engineered from the ground up to seamlessly scale to global, multiplatform usage, harness the content-sharing power of fiber-optic wide area networks, and rapidly integrate other off-the-shelf asset management software. The robustness of StudioCentral Library 3.0's new Asset Management Protocol will reduce the time to market for software development companies, application service providers, and IT departments.

When power problems strike, computers—the star of modern business practices—can cause major process problems. It's distressing, and it can be very costly.

Will Your Data Survive

The Next Power Interruption?

With the current rolling blackouts throughout California, more than ever the need for uninterruptible power systems (UPS) is crystal clear. The state's electricity crunch is creating havoc in many different places, especially for companies that need to keep their servers running even when the electricity goes out.

Companies in Silicon Valley and throughout the world whose business survival depends on the reliable processing of data cannot afford to have their files corrupted or their data processing interrupted due to power irregularities. The impact is felt worldwide, and chances are that your data is at risk as well.

The Cost of Downtime

E-commerce companies that rely on smoothly functioning Web sites lose anywhere from $1 million per hour to $1 million per minute during a power outage due to lost orders and a decline in workforce productivity.

Although it was not caused by a blackout, a November 2000 power outage at Seattle-based e-tailer Amazon.com showed how costly even a brief outage can get. During the Thanksgiving holiday weekend, Amazon suffered a series of outages. Investment firm Thomas Weisel Partners estimated that one 20-minute outage deleted roughly 20,000 product orders and $500,000 in revenue.

In June 1999, shares of eBay Inc. plunged 30% in the wake of a site crash that cost the online auction house millions of dollars. The company revealed that a nearly 24-hour crash of its auction site was expected to reduce sales by 10% in the second quarter of that year. Early this year, eBay's site suffered once more from an outage that had a financial impact on the company.

More Facts and Figures

Each year, the typical computer location experiences 36 spikes, 392 voltage fluctuations, 128 voltage surges, and 15 blackouts[1]. More than half of the service calls and downtime experienced on computers and networks can eventually be traced back to the utility power or mains. Recent estimates show that annually, damages add up to $26 billion in lost data, time, and damaged hardware due to power failures[2]. According to estimates by the Electric Power Research Institute in Palo Alto, California, power outages cost U.S. businesses $29 billion annually. Ninety percent of those outages are the result of distribution system failures. Also, the average cost of an hour of downtime for leading online retailers is estimated to be $78,191[3].

According to IDC's Technology Integration Panel Study of 1998, 31% of all server failures are caused by power problems (e.g., power failure). In addition, power blackout occurrences increased 150% from 1972 to 1994. It's shocking to imagine the exponential increase in power blackouts and brownouts for the years 2000 and 2001. This costly downtime can be avoided with a simple solution: an uninterruptible power system.

The Uninterruptible Power System

A UPS is a battery backup system for electronic equipment. The UPS positions itself between the raw mains and the electrical equipment being protected. It conditions and filters the power, eliminating any electrical spikes or surges that may cause damage to the server and its data.

Many people assume that the electricity coming from the utility is pure. However, there are a multitude of power problems such as sags, spikes, brownouts, and surges that can affect power quality. Sensitive electronic equipment such as computers, hubs, and routers are susceptible to power variations. Not only does a UPS provide backup power in the event of a power failure, it also regulates the incoming electricity before it reaches the connected equipment.

The units, which range in cost from under $100 to tens of thousands of dollars, keep computers and other high-tech equipment running even when the electricity goes out. Interest in such products has exploded in California with the recent power shortages and rolling blackouts. Commonly available uninterruptible power systems can be as small as lunchboxes; larger enterprise systems for data centers can fill up a whole room.

The level of UPS protection you need varies with the size of your organization and the criticality of the systems. Some companies have one large UPS in the basement of the building that protects the power supply for the entire building.

The larger UPS is often backed up by a generator, which can provide longer term uptime. Other companies may just protect the network server or individual workstations where the data is vital to their business.

The Power Protection You Need

SGI offers a full line of UPSafe™ solutions, including power management software, to protect your electronic equipment against power problems—even in the most critical circumstances. Protection during extended blackouts includes a user-defined, planned shutdown for complex server–client environments. By providing conditioned power at all times, you maximize the productivity of your company's homogeneous or heterogeneous IT infrastructure. Included in UPSafe solutions is a three-year support contract designed to complement the support contract purchased for your SGI system. SGI offers UPSafe solutions with power capabilities matching the current range of SGI products, including highly scalable, modular, and redundant solutions for products as flexible as the SGI™ 3000 family. Custom solutions are also available.

For more information:
www.sgi.com/services/productivity/

[1]Source: National Power Laboratory
[2]Source: Power Electronics Application Center
[3]Source: *ComputerWorld*

Optimizing Digital *Product Design*

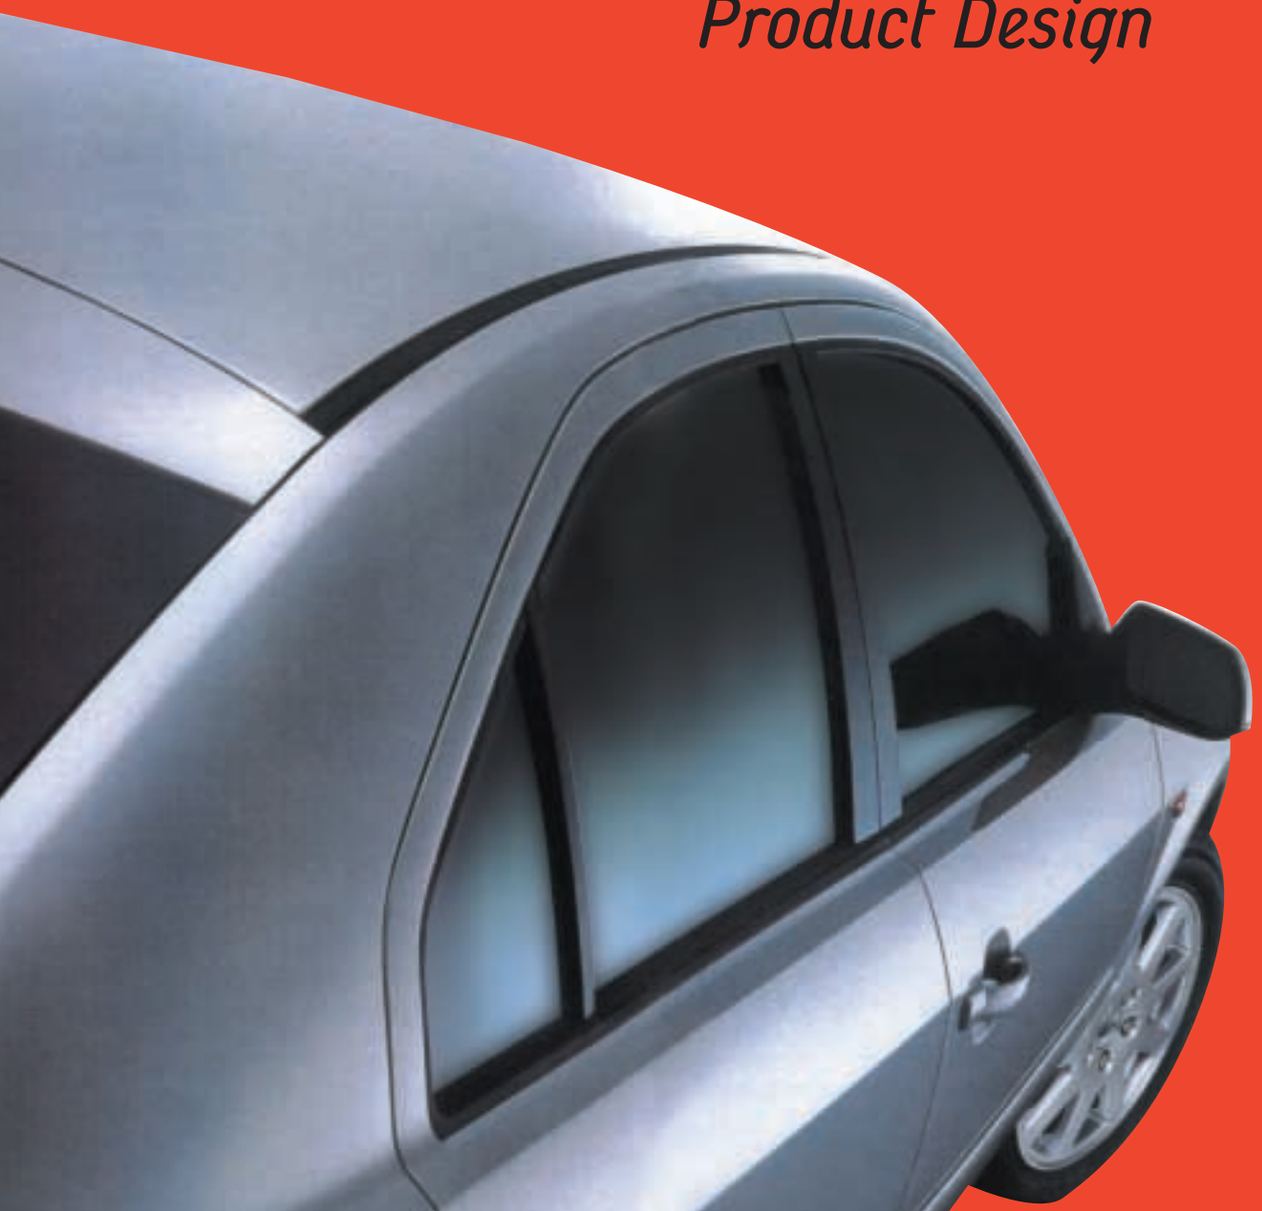

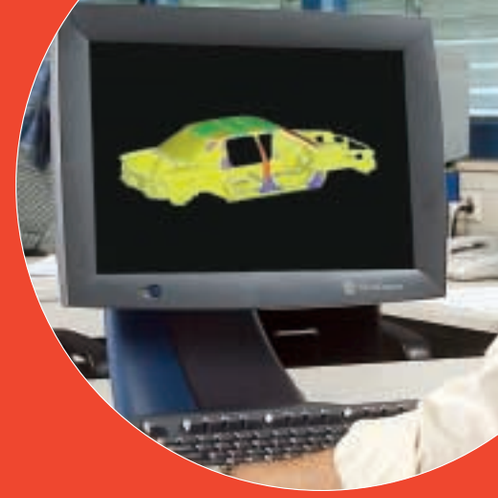

In manufacturing industries, a new product design involves intensive collaboration between teams with specialized disciplines.

In the design process, specialized teams often have conflicting considerations, such as aerodynamics, structures, propulsion, and controls, as well as costs and returns on investment. In an engine-blade design, a thin leading edge might be desirable from an aerodynamics point of view, but structurally, a high degree of stress would result.

A successful design requires manufacturers to integrate the parameters and come up with an overall optimum design across disciplines. Often designs are passed between the product teams several times until the differences are minimized and a mutually acceptable solution is found.

In the aerospace industry, multidisciplinary design optimization (MDO) methods have been effectively deployed to overcome the drawbacks of this sequential design process. Algorithms have been developed to compute data from different disciplines and provide an optimal solution.

"MDO has been implemented in the aerospace industry for 30 years or so," said Dr. Srinivas Kodiyalam, of SGI, who led a joint research project in 1999 with NASA Langley, Ford Research Laboratory, and Engineous Software, Inc., to apply MDO methods to the automotive industry.

Kodiyalam said that until recently, MDO was not feasible for applications in the automotive and other manufacturing industries, partly because these industries typically did not invest in developing in-house software tools. "The automotive industry is heavily reliant on commercial off-the-shelf, vendor-provided tools," he said.

The good news is these key technologies are now emerging as commercial tools. "Traditional CAE and CAD vendors now support the capability to do optimization," Kodiyalam said. "Software vendors focused primarily on process integration, simulation, and optimization environments also are emerging."

Another challenge in applying MDO to automotive manufacturing is the sophisticated high-fidelity models that have evolved as the standard in the industry. Without superior computing power, elapsed computing time for such detailed models could take years. "New technology is great, but such technologies can't be useful unless computing time can be reduced to decrease the product design cycle," Kodiyalam said. "That's where high-performance computing machines come in."

In a recent joint project between Ford and SGI, Kodiyalam and the Ford team set out to reduce the elapsed computing time to optimize a vehicle system for minimum weight while meeting international crashworthiness standards and Ford's standards for noise, vibration, and harshness (NVH) levels. The design Ford provided violated both the NVH and safety targets. Changing the design solely for NVH attributes might adversely affect the safety targets.

The computer intense, nonlinear analysis required for optimizing this particular design was daunting. "Using MSC.Nastran and RADIOSS, Ford estimated almost three years of computing time using a single processor on SGI Origin 2000*," Kodiyalam stated. "On an SGI Origin 3000 series system with 256 processors, we reduced that computing time to less than two days."

"At Ford, we want to bring quality products to market faster," said Ren-Jye Yang, senior engineer at Ford. "With SGI and NASA, we were able to demonstrate that with high-performance computing, a conventionally impractical multidisciplinary optimization problem becomes feasible and thereby impacts the product design cycle."

Not only was elapsed computational time reduced, Kodiyalam said the team was able to turn an unfeasible design into a feasible one. The integrated technology is capable of presenting previously unimagined solutions and could change the overall design process. "Manufacturers can reverse engineer products," he noted. "They know what they'd like to see, and MDO shows them how to optimize the design."

Kodiyalam said the technology creates the possibility for remarkably unique product innovations, a departure from derivative or heritage designs. "MDO analysis will come up with a design solution that engineers haven't thought of," he said.

MDO technology has a wide range of applications. "I can see as much application in a general manufacturing world as there is for automotive design," he said. "The set of algorithms is applicable to the biomedical world, to general manufacturing, and to the financial world."

For more information:
www.sgi.com/manufacturing/mdo/
www.sgi.com/manufacturing/automotive/

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

Images courtesy of Ford Research Laboratory.

Artificial Intelligence:

The Dream of Smart Machines

The fifth RoboCup soccer championship for robots, to be held in Seattle August 2–10, 2001, is a unique event designed to foster international research in artificial intelligence (AI).

The event is organized by the RoboCup Federation, a Swiss-based organization that brings together more than 150 universities, national research institutions, and private corporations from 35 countries with a focused objective: by 2050 a team of fully autonomous humanoid robot soccer players, complying with official rules, shall win a game of soccer against the human world champions.

However, soccer supremacy is not the ultimate objective. The RoboCup Federation's aim is to apply this evolving technology to greater industrial and social needs such as disaster rescue, transportation, construction, space exploration, and telecommunications.

SGI has signed a three-year worldwide sponsorship agreement with the RoboCup Federation and is providing a number of visualization systems to support the organization's activities. Sony Corporation, which funds the Sony Legged Robot League, is also a major sponsor.

The founder and chairman of the RoboCup Federation is Hiroaki Kitano, senior researcher at Sony Computer Science Laboratories, Inc. and a director of the ERATO Kitano Symbiotic Systems Project, Science and Technology Agency of the Japanese government. His other research interests include computational molecular biology, evolutionary systems, and engineering use of the morphogenesis (development of form in organisms) process. SYNERGI asked him to share some thoughts on the evolving area of AI.

How would you define a robot?
I suppose the general perception of a robot is probably a mobile machine with movable components, operated by a certain level of intelligence and with a certain level of autonomy. However, there is no agreed-upon definition.

When did you first become interested in AI?
When I was an undergraduate studying physics, I had to do a lot of data processing. This led me to create an adventure game with an astronaut stranded on an unknown planet. Players had to program the agent to move around. At about the same time, I became interested in the need for a machine to interpret human language. Then I discovered that AI was a major field, very premature and very young, with lots of exciting opportunities.

How did you come up with the idea for RoboCup?
In 1992 I organized a small workshop in Japan for researchers to think about the most important way in which AI can have an impact on society and mankind. Other fields of research have clear challenges—for example, the human genome project. We were not sure what would be our major challenge. We needed an interesting project that would attract people. We felt that soccer would be a good target, and we ended up creating an international event.

How are you using SGI technology in RoboCup?
SGI provides hardware, software, networking, connections, and support for the duration of the event. We port everything onto Linux. There must be no disparity, no unfairness in the availability of computing power.

We want to encourage a diversified approach. We let participants compete, then the winners are asked to disclose the source code so others can benefit. The rate of evolution is therefore very rapid. Everyone is on the same ground after each event so everyone has to keep on working. It is very different from the style of Formula One where everyone tries to keep their secrets. In fact, it is very like the Linux community. It is not just a competition, it is a joint project. Our final goal is not to play soccer but to contribute to society.

How do you get robots to work together?
This is where engineering capability is very important, and each team has a different approach. Some teams have a program for players in certain field positions to take a proactive approach when a certain event or sequence of events happens. Others focus on reacting. There is no explicit planning for teamwork but it will emerge. Others program the robot for repetition in order to train the robot to behave in a certain way and learn over time.

Teamwork is a major challenge. We will probably find that there are several ways that will be applicable. We do not want to define one way. If you select only one way then evolution is stopped.

In the year 2050, what will robots need most in order to beat a team of humans?
Safety. We need to have a human team that will agree to play. Our goal is not to create a powerful robot that wins on sheer force. We want to have a game, an interaction. We want professional players to feel that it is worth their time and effort.

The second major issue is artificial muscle. Soccer robots need to walk, run, and jump. Gears and motor-based systems may be too fragile. In addition, metal and plastic, the materials currently used, are not suitable. We need a material that is both soft and strong.

We also need to teach robots how to make decisions. They may be near the goal and able to kick the ball into the goal but there might be a human being in front. They need to make a decision to pass the ball so that they do not hurt the human.

Likewise, in disaster recovery, the objective might be to pull a person out of some debris. However, doing this might cause the victim physical damage. A robot needs to be able to determine when it is appropriate do something else first, such as move the debris.

Do you think that it's important to replicate human thought and movement?
I think robots should have a different body structure so that they can do what humans cannot do. It would be redundant to have a robot do the same as humans. However, if we have a biped robot then it can be used in the same living environment as humans. If it can reach light switches, for example, it has some value.

If we focus on a specific aspect of human behavior then we can attain excellence with robots. However, humans have many integrated aspects; they operate reasonably well on various fronts. It is extremely difficult to emulate this and get a robot to do multiple tasks. No human can beat a computer, but no machine can do all the things humans do. Machines are very specialized, specific tools. They are robust but may not necessarily adapt to a changing environment, whereas humans have the ability to adapt.

There are parallels with the animal kingdom; robots will have their own ecological niches. Robots that are snake-like, for example, will have their own usefulness.

How will robots make a positive contribution to society?
Robots will augment human capability by doing things that humans cannot do. If a robot can withstand pressure or heat or go into a space that a human cannot go into, then it is very useful. There is also the possibility of robots helping with the care of handicapped or sick people. People may not wish to be helped directly by a robot but robots can help the caregivers—for example, they can carry and hold items. We are only limited by our imagination in developing and using robots for the wealth of mankind.

Presumably you can program robots to be self-learning. How do you prevent them from learning in the wrong way?
That's an interesting question. It's like how do you raise your kids not to do anything wrong. I'm not quite sure to what degree we can make robots self-learning—we cannot program them for every eventuality. Right now we can program them to learn something but it is within very limited parameters. There is a lot of scope to do things the wrong way. However, robots are not going to cause any danger to human beings at the moment.

Science fiction writers often play on the human fear that robots will take over the world. How near or far from reality is this concept?
I believe that mankind is not so stupid. As more sophisticated robots come on the scene there are two possibilities: a robot may malfunction and do damage or a human can program a robot to damage. Each of these cases would involve only one robot so the potential for harm is limited. I think that in the long term governments will bring in regulations for owning and operating robots.

Some people say robots will evolve by themselves and will need to be self-sustainable. That would not be easy but it raises an interesting point. If robots evolve to hurt humans, they will have less of a chance of staying here. Evolutionary robots who are friendly will survive.

Science and tools can be positive or negative. In the end it depends on human wisdom. The automobile, for example, is extremely useful, yet it kills many. In fact, our most dangerous technology is very simple—a sharp edge. We need it for our daily survival but we have evolved to handle it. A robot is no exception. We need to be careful; it is not technology that counts but wisdom.

How would you define what separates a human being from a robot?
That's a difficult question. Robots are nothing like humans. However, as robots evolve, the situation may be different. In some ways, we can say that a human is a molecular robot. What makes a life form? From a scientific aspect it is the level of complexity in a system. If we look at robots in this way they are still emerging. If we think about the spiritual and intelligence aspects, one thing humans have is self-awareness. We know who we are. How a robot starts to recognize that it exists—now that's a big challenge.

In the future robots may have that kind of perception. But we need to be careful. If we create that kind of self-awareness, a robot will no longer be a tool, it will be another species. That raises a whole new issue.

What progress do you think AI will make in your lifetime?
It is only 50 years since the first digital computer so it is difficult to predict the next 50 years. We will probably be using AI on a daily basis, taking it for granted and not worrying about it at all. There will be lots of progress, but this will still be very small compared with the level of human intelligence.

For more information:
www.robocup.org/

Images courtesy of Sony Corporation.

Immersive environments can transport the human consciousness backward or forward in time, to remote locations or to other galaxies. No place can be too far, no time too distant.

Virtual Time Travel: An Unforgettable *Experience*

Images provided by Imaging Media Research Centre, KIST.

Immersive environments can transport the human consciousness backward or forward in time, to remote locations or to other galaxies. During fall 2000, visitors to the Kyongju World Culture EXPO 2000 streamed into the 651-seat Cyber Images Hall, the world's largest immersive and interactive virtual reality theater.

There they waited before an immense 89- by 26-foot cylindrical screen to see the exposition's most dazzling presentation: a virtual walk through the ancient Buddhist Silla Kingdom, generated in real time by a Silicon Graphics®, Onyx2®, visualization system. The visitors were guided through seven historical venues depicting the ancient Korean city of Sorabol. They pressed buttons on their chairs to vote for the routes that they wanted to follow or to control, for example, the movement of butterflies in cyberspace. Detailed, textured 3D images, sounds, smells, and physical sensations transmitted through the seats with subwoofer speakers made the walk-through an unforgettable experience. The life of the remarkable Silla culture of 1,300 years ago became a reality. This experience was a message about harmony—between tradition and modernity, humanity and nature, technology and culture.

The Kyongju World Exposition: Encouraging Reconciliation and Peace

The Kyongju World Culture EXPO 2000, held in Korea's old city of Kyongju from September 1 to November 26, was a remarkable celebration of Korean history and culture. Federal and provincial governments sponsored the event. The main theme was Breath of the New Millennium, but its secondary theme, Encounter and Harmony, describes its spirit more precisely. By taking visitors back in time to the era of the Silla Kingdom, it encouraged Koreans to draw on the strength of their cultural heritage to promote reconciliation and peace. The era of the liberal, progressive Sillans is viewed as the Golden Age of Korean culture, and its center was Sorabol—the Kyongju City of today.

The exposition included a great variety of exhibits, events, and performances on its themes, ranging from fusion art performances to an outdoor opera. But the most dazzling attraction was the virtual re-creation of ancient Sorabol in Cyber Images Hall, an SGI™ Reality Center™ facility. A six-pipe, 14-CPU Silicon Graphics Onyx2 InfiniteReality2™ system drove six BarcoReality 9300 LCD projectors, double-stacked to deliver front-projected, edge-blended, passive stereoscopic images to the huge cylindrical screen. The Korea Institute of Science and Technology (KIST), Korea's leading think tank, designed and implemented the immersive and interactive virtual reality theater. KIST worked with local software companies to develop the application software, which ran on top of IRIS Performer™. Additionally, KIST developed hardware that enabled a guide to navigate through the virtual world and make the experience unique for every audience.

The Experience: Virtual Time Travel

The virtual experience was entitled Into the Breath of Sorabol. First, the Cyber Images Hall audiences experienced the events that created the Silla Kingdom. Then they began a highly realistic tour of reconstructed architecture of the Silla era, all created on the basis of historical and archaeological knowledge.

"We wanted to allow visitors to choose the routes they took through the ancient city of Sorabol," says Dr. Hyoung-Gon Kim, head of the Imaging Media Research Center, KIST. "To do that, we needed a high-performance and high-quality visual supercomputer system capable of generating high-resolution imagery in real time; that's why we chose the Silicon Graphics Onyx2 system to power the Cyber Images Hall. The audience feedback confirms that we made the right choice."

"Kyongsangbuk-do, a southeastern province of Korea, is a traditional leader of Korean national culture," says Jung-Bae Lee, president of the Kyongju World Culture EXPO 2000 Organizing Committee. "We are proud to present a new paradigm for cultural events in the 21st century."

The Silla forefathers adopted the elements of a variety of continental and maritime cultures and harmonized them with the unique flavor of their own culture. The power of this philosophy is evident in the seven experiences available to the theater's audiences, which included a visit to the Hwangryongsa (Royal Yellow Dragon) Temple, completed in 553 and the largest of the Silla temples. Audiences were deeply impressed by a virtual reconstruction of the dramatic nine-story wooden pagoda in the temple towering 80 meters, the tallest one in Asia. The pagoda was completed in 645; its nine stories represented the nine countries surrounding Silla and symbolized Buddha's generosity to the countries. The pagoda was razed in 1238 by Genghis Khan, and currently only the foundation stones remain. Other experiences included visits to a 207-foot-long bridge of Moon Spirit (Woljeong-Kyo), Moon Castle (Wolsung), a royal pond of Anapji, and the Sokkuram Grotto Shrine, an ancient mountaintop shrine to Buddha that includes 38 superb wall carvings and a remarkable domed ceiling made of cut stones.

Interactive and real-time capabilities, combined with high-resolution immersive imagery, create a much richer understanding of this ancient culture. "Just as the Silla people of 1,300 years ago embraced and assimilated other cultures, the organizers of the Cyber Images Hall embraced cutting-edge technology to educate the world about Korean cultural heritage," says Sheila Pickett, SGI museum market manager. "This is how future generations will be educated about their own cultural heritage and the heritage of other cultures around the world. We are proud to have SGI Reality Center technology play a role in bringing this ancient culture back to life."

For more information:
www.sgi.com/entertainment/museums/

SGI and its partners have once again swept the Academy Award nominations in the area of visual effects. For seven years straight, every Academy Award finalist has relied on SGI systems to create award-winning films.

The Academy Awards:

Visualizing the Impossible

From the stunning megalopolis of ancient Rome brought to vivid life and a deadly invisible man appearing and disappearing in anatomically correct layers to the terrors of the sea unleashed on a fishing trawler off the Eastern seaboard, all of these unforgettable images were captured using SGI systems.

The amazing images featured in *Gladiator*, *Hollow Man*, and *The Perfect Storm* would not be possible if it were not for the groundbreaking effects produced courtesy of England's Mill Film, Sony Pictures Imageworks (SPI) and Tippett Studio, and Industrial Light & Magic (ILM), respectively. Most exciting is the fact that SGI continues to be the platform of choice for these remarkably innovative contenders.

Gladiator
For visual effects supervisor John Nelson (*City of Angels*) and a team led by Mill Film's visual effects supervisor Tim Burke, the task presented by director Ridley Scott on his epic period piece, *Gladiator*, was formidable to say the least. They needed to re-create the grandeur that was Rome without bankrupting the studio, which literally happened when Hollywood revisited the Golden Age for 1965's *Cleopatra*. Fortunately, Nelson and the team at Mill Film used their SGI hardware to devise awesome visuals through the magic of digital technology that didn't break the bank and netted them an Academy Award nomination. "The concept of Rome as an empire is just overpowering, and the city itself, at the time when this story takes place, had a million people," says Nelson. "It took another 1,000 years for another city to have a million people, so our main goal was to broaden the scope of the story, beyond the sets, with huge crowds and vast scale."

SGI had a strong presence even when cameras went on location. Mill Film chose to use Silicon Graphics® O2® workstations for on-set previsualization. Then, using O2 workstations, Indigo2 Impact™ workstations, and Silicon Graphics® Octane® workstations for programming and 3D animation, they created the epic Rome sequences. The elaborate imagery was rendered on a dedicated renderfarm of nine SGI™ Origin™ 200 dual-processor machines with a gigabyte of RAM per processor and four Linux 500 MHz dual-processor machines.

The epic war sequence that opens *Gladiator* is a stunning re-creation of a battle between the Roman legion headed by General Maximus (Russell Crowe) and the German fighters. It was a difficult lesson in logistics.

"One of the big things we had to do in that sequence was show the technological superiority of the Roman war machine and also the scale of the battle," Nelson says, "except we had 1,500 extras and we wanted to create the appearance of 10,000 soldiers."

To bring this battle to life, Mill Film created a spectacular pan across the battlefield from three locked-off shots using six Octane workstations running Discreet flame and two Silicon Graphics Onyx2 systems running Discreet inferno. The three tiles were later pieced together by Mill Film's artists into a single seamless plate, over which a virtual pan was added. "They did several passes on each tile, then filmed the Romans in the background, the midground, and the front. In addition, layers of marching Romans were added. Then they layered in the Germans in the background, left and right, and shot the pyro explosions going off amidst the Germans separately, without any actors in the plate. Then Mill Film synced all the practical elements, created a virtual pan, and added lots of flaming arrows and fireballs. And underneath all of this firepower you see wave after wave of the advancing Roman army. It was a long, arduous process to broaden the scope and scale of the story," said Nelson.

Hollow Man
Paul Verhoeven's postmodern riff on H.G. Wells's *Hollow Man* boasts the closest thing to a digital human—albeit an invisible one—courtesy of SPI and Tippett Studio, the same team responsible for the award-nominated work in Verhoeven's *Starship Troopers*. The result is the most stunning transformation from visibility to invisibility ever. Rather than doing a traditional invisibility effect, Scott Anderson, SPI's visual effects supervisor, convinced Verhoeven to let Kevin Bacon's performance drive the creation of the elaborate invisibility gags. Consequently, Bacon participated in every shot, wearing blue, green, or black costumes and matching makeup so he could eventually be removed from the plates, leaving his invisible alter ego to wreak award-worthy havoc.

R&D on *Hollow Man* began in mid-1998. Anderson and his 150-plus Imageworks team quickly realized that in order to create an invisible man on-screen, they were going to have to build a completely digital human from the inside out. SPI's eye-popping transformation animation utilized Maya® software from Alias|Wavefront. Compositing was run on Octane visual workstations, O2 workstations, and SGI™ Onyx® family visualization systems. Tippett Studio, which animated the shots of an invisible Kevin Bacon, whose form was revealed by such elements as smoke, fire, water, and blood, relied primarily on Maya and Renderman, along with a host of in-house tools, to produce the images. Both Tippett's and SPI's images were then rendered on SGI™ 2000 series and Silicon Graphics® Challenge® 10000 servers. "We've taken the best of medical science, applied it to what we need to do, and used off-the-shelf software to control the process, but the process itself is entirely proprietary," Anderson says. "We've written huge bodies of software to help with everything from the removal and integration of Kevin into the film to the animation and structure of the human forms, which should look like an entirely organic process. The film explores the transformation almost like a scientific study where you get to witness and savor the process of the body transitioning from visibility to invisibility. That was tough because you can do very complex work that in the end still looks like a dissolve. Paul wanted it to feel like Kevin's body was dissolving in a way that was logical and beautiful, so we made each part of his anatomy disappear individually. No one's ever experienced a dynamic, sentient being from the inside out before. It's certainly the most complex film I've ever worked on."

The Perfect Storm
In Wolfgang Petersen's true-life sea saga of the doomed fishing ship Andrea Gail, *The Perfect Storm* demanded an unparalleled mastery of the raging deep. Weather and the elements are traditionally the hardest things to replicate via visual effects, but ILM marshaled a legion of Silicon Graphics O2 and Silicon Graphics Octane workstations to bring a storm of biblical proportions to the big screen.

ILM chose to use fluid dynamics, a computer simulation of real-life phenomena, to digitally replicate the massive walls of water resulting from the impact of three deadly weather systems. Technical directors did this on Silicon Graphics O2 workstations and then rendered the images on SGI 2000 series systems. The result was a wet and wild digital ride on the high seas, arguably the most spectacular storm sequence ever.

But ILM couldn't just run the simulation and call it a day. "That's just not how you stage a shot in a major motion picture," says ILM's reigning weather man, visual effects supervisor Stefen Fangmeier (award-nominated for *Twister*). "On top of the big rolling waves, we applied all sorts of tricks to create the little 2-foot and 3-foot waves that actually build on top of those 80-foot waves. Over that, we devised five or six different levels of particle simulations to simulate the white foam and a windswept look. We couldn't just take static textures and slap them on our waves. Everything had to move with the ocean, which made it very difficult."

Making the Impossible a Reality
In 2000, the Academy of Motion Pictures once again recognized special effects created via SGI systems. The Academy Award-nominated visual effects of *Gladiator*, *Hollow Man*, and *The Perfect Storm* are the product of the imagination and skill of a myriad of artists from Mill Film, Sony Pictures Imageworks with Tippett Studio, and Industrial Light & Magic. These artists all used Media Commerce solutions from SGI to make their impossible imagery a reality.

For more information:
www.sgi.com/features/2001/mar/oscar/

Powering *Astrophysical* Research

The University of Leicester, home to the new UK Astrophysical Fluids Facility (UKAFF), hosts the third largest academic supercomputer in the UK and the largest dedicated to a single area of science, based on SGI Origin 3000 server series systems.

UKAFF is a £5.9 million project funded jointly by the UK Higher Education Funding Council for England (HEFCE) and SGI, with further support from the Particle Physics and Astronomy Research Council and the Leverhulme Trust. The facility is managed and directed by Professor Andrew King of the University of Leicester in conjunction with Professor James Pringle of the University of Cambridge.

The SGI NUMAflex™ concept is at the heart of the SGI Origin 3000 series. NUMAflex is based on a unique modular approach consisting of a brick-style architecture for constructing small to very large systems from a common set of building blocks. This modular approach allows users to build the optimum configuration one component at a time and adopt new technologies that map to their specific business or research needs as they emerge. In contrast, traditional high-performance computers need to be replaced frequently to keep pace with competitive demands and technological changes—a costly and cumbersome process.

This burden is totally impractical for some organizations and particularly so in the academic world, where institutions and research projects are funded by public or public/private initiatives; university departments not only must justify their investments but have limited cash reserves and need to deploy resources carefully.

However, in academia as in commerce and industry, staying ahead of the competition is critical. Academic research in the 21st century is no longer knowledge for knowledge's sake. Today it is an essential practical discipline that is critical in securing industrial and commercial competitive advantage. It also serves a national purpose through providing credibility and prestige for excellence within a particular field. To exemplify, understanding how turbulence affects the computational fluid dynamics of a plane wing at various speeds will allow plane component manufacturers to design more effective wings in the future. Manufacturers may therefore choose to sponsor an academic group specializing in CFD to conduct this aspect of its R&D work. At the same time the group is able to further its own research. This type of collaboration is beneficial to both parties, while benefits to industry also have implications on a national level.

Additionally, to specialize in a particular discipline and show exceptional prowess in this area helps to raise the profile of a nation. Investment in academic research projects can therefore have a wide-ranging positive impact for the national good, and the government backs research via the Research Councils and through the HEFCE's Joint Research Equipment Initiative (JREI).

With this in mind, it is clear that investment in the latest technology and equipment is essential; otherwise a competitor with more funding will advance its research far faster. Access to supercomputing facilities is crucial today for all science-related disciplines from medicine and surgery to astronomy and geophysics. Use of such technology allows researchers to extrapolate their findings to a degree unprecedented prior to the introduction of the high-powered machines available today.

The new SGI Origin 3000 servers series offers a solution. Although the SGI Origin 3000 series is a superior product and therefore the initial cost outlay is still significant, the modular nature of the design means that the core system can be added to or altered where necessary to accommodate new technological developments.

The system also incorporates a number of advanced features, including high processing speeds and exceptional scalability. These features make the SGI Origin 3000 series ideal for computationally intensive work such as astrophysical research, a specialist area at the University of Leicester. The strength of the SGI Origin 3000 series and its perceived value for the UKAFF was sufficient for the University of Leicester to be allocated funding for the project. This came from the JREI, with matching funding provided by SGI.

Scientists at Leicester will be using the newly installed system for calculations, tackling some of the most complex mysteries of the universe. The facility will also be available to astronomers across the UK who make calculations for world-leading projects in theoretical astrophysics. The speed of this machine will allow them to calculate much more complex problems than ever before possible. Dramatic advances can now be made within research into matters as yet not fully understood, such as black holes and colliding stars.

SGI Origin 3000 series servers are the only servers currently scaling to the processing levels required by astronomers while retaining a shared-memory architecture, making computational methods far more efficient. The SGI Origin 3000 series system to be used by UKAFF contains 128 MIPS® R12000™ processors in a single shared-memory configuration based on SGI NUMA technology.

To keep the SGI Origin 3000 series system safe from unplanned power interruptions, an SGI UPSafe uninterruptible power supply was installed as part of the SGI Managed Services solution. The 60,000VA UPSafe solution provides some 30 minutes of run time during power surges, brownouts, and power failures. In addition, in the event of a power failure the system will automatically trigger a controlled shutdown, allowing it to halt any strings that are running so that they can be resumed when power is returned.

Andrew King, professor of astrophysics at the University of Leicester, commented, "The availability of this high-speed machine for general use by astronomers means that the UK will retain its world-leading position in theoretical astrophysics. This has implications elsewhere: many calculations made in astrophysical research are similar to those performed in industry, for example, the effect of a bird strike on an airplane engine. UKAFF will now be able to foster the growth of expertise in this area in the UK."

Andrew Grant, higher education business manager for SGI in the UK, commented, "SGI is honored to have been chosen by the UKAFF consortium for this exciting project. The unique scalability of the SGI Origin 3000 series, coupled with advanced SGI visualization technologies, allows for unprecedented amounts of data to be analyzed. This will ultimately result in new insights in many areas of theoretical astrophysics."

The official opening of the facility by Dr. John Taylor, director general of the UK Research Councils, was October 31, 2000. At the launch event, a number of exciting new projects were discussed by some of the professors who have been using the new systems. One such project involves astronomers who have been following the trailblazing birth of a cluster of stars like the sun, in unprecedented detail.

The computational work involved in this project was led by Dr. Matthew Bate of the Institute of Astronomy at Cambridge University. Stretching the UKAFF supercomputer to its limits during its initial tests, he was able to simulate the collapse of a cool gas cloud as it shrank inward under its own weight and broke into about 100 individual stars. Most of the gas was expelled back into space by heat from the stars, which themselves emerged as a cluster of newborn stars like the Pleiades in Orion. King is currently working on a number of projects that will benefit from his use of the SGI system. For example, he is investigating stellar collisions, when a black hole swallows a neutron star, which can result in a gamma ray burst. In this event, the amount of energy produced is equivalent to doubling the light output in the entire universe for a brief period.

Additional research projects at UKAFF are expected to yield further interesting findings, which will help Leicester University to keep ahead of competitive research organizations across the world in theoretical astrophysics.

The Birth Of The Online *Photo Album*

You live in London, and you've just taken pictures at your daughter's graduation party with your new digital camera. You'd like to share them with the rest of your family members, who live in New York, Amsterdam, and Zurich.

You download the pictures from your camera to your PC, then download them to your photo album on your photo dealer's Web site, which resides on an SGI™ Origin™ 3200 server at ShareAPhoto. You crop them, brighten up the pictures that are a little on the dark side, add a few captions, and frame the first picture with a cap, gown, and diploma. You enter e-mail addresses, click Send, lean back, and smile.

In a few minutes, your far-flung relatives will receive an invitation to visit your album, check out the photos, and add their own comments. If Uncle George and Aunt Lucy want prints of your shots, they can order them with a few clicks. There will be high-quality prints on photo paper in their mailbox in just a few days.

The Shoe Box of the Digital Generation
ShareAPhoto, a London, England-based business-to-business application service provider (ASP), is using SGI server technology to solve the well-known problems that surround family photography: disappearing negatives, shoe boxes in hard-to-find places, and seldom-seen albums. Digital photography has solved some of these problems, but ShareAPhoto's founders knew there was room for improvement.

"With the digital camera, you can print out your pictures or e-mail them around, but people want more," says Bradden Wondra, co-founder and CEO of ShareAPhoto. "They want to create albums or home pages without any technical skills, and they want to be able to get the same quality prints they get from a film camera. That's what we've enabled in our ShareAPhoto products and services, which are based on SGI server technology."

ShareAPhoto, which was launched in February 2000, gives ISPs and Web portals a sticky application that brings visitors back again and again. It also gives traditional photo-finishing shops a turnkey entrée into the digital revolution by hosting their customized, branded Web sites on ShareAPhoto's SGI servers. Customers create photo albums on these Web sites using digital photos, scanned-in photos, or high-resolution scans of film negatives, which their photo shop or film processor can post directly to their albums. They can then give their creativity free rein. ShareAPhoto has developed proprietary technologies to provide a wide range of image editing, retouching, text, and special effects tools.

Hardware Solution: SGI Internet Server Technology
All this movement of high-resolution images requires high-availability, high-throughput Web server systems and massive storage capacity. ShareAPhoto's founders looked at their vendor options and chose SGI. "We evaluated other technologies, but we found the scalable SGI NUMA technology to be superior," says co-founder and CTO Mark Berry. "It allows us to give our customers responsive Web sites, despite the size of the files we're moving. We also found that SGI systems give us lower total cost of ownership and maintenance, partly because SGI can fill our needs with fewer machines."

The Next Phase: Video Albums Online
On October 10, 2000 at the streamingmedia.com conference in London (*www.streamingmedia.com/europe/*), ShareAPhoto announced the next phase of its ASP program: ShareAVideo. "We're adding Real and Windows streaming media formats, which people can easily view across the Internet," says Wondra. "Using digital video (DV) streaming media technologies from SGI, ShareAVideo converts video straight from the DV camera into Internet-ready streaming media. The consumer never has to think about the complexity of formats and standards." Home users with DV cameras can download video to their computers and use any browser to quickly upload files to their multimedia albums. The files are converted automatically to the appropriate format and bandwidth for viewing by others, leveraging the power of industry-leading digital media technology from SGI.

As they deal with the rapid growth of their enterprise, Wondra and his colleagues are pleased with their selection of SGI as a technology vendor. "Apart from sheer product performance, it's giving us great marketing support, which we need at this point in our development. SGI is the industry leader in the digital media space."

For more information:
www.shareaphoto.com/

Dossier

Measuring the Earth with *Precision*

At the Texas headquarters of TerraPoint, a group of visitors from Mitsubishi Corporation in Japan, wearing stereo CrystalEyes® goggles, gathers before a Silicon Graphics Octane workstation.

They need precise digital measurements of the entire city of Tokyo, including the tiny lanes and buildings of the old city, from which they will derive the data-mapping products they offer their customers. They're interested in TerraPoint's precise digital models of terrain and urban landscapes, but they have doubts about its ability to resolve the details they want. The demonstration begins with the workstation screen displaying a high-level view of a model of New Orleans, Louisiana. The operator uses the mouse to fly the visitors down into the French Quarter, and the group suddenly finds itself looking into an alley so narrow that one could easily touch both walls simultaneously. There is an appreciative intake of breath from the visitors.

Mitsubishi Space Imaging: Tokyo in Detail
TerraPoint light detection and ranging (LIDAR) data is highly desirable for building 3D urban simulation models. It provides approximately 450,000 data points per square kilometer, with typical accuracy of 60 cm or better in the XY plane and 30 cm or better in the vertical dimension. Space Imaging recently pointed out to its foreign affiliates that where U.S. Geological Survey data was not available, LIDAR data could be a more efficient solution than stereo-pair analysis for building digital elevation models. Mitsubishi Corporation's Satellite/Information Systems Unit hired TerraPoint to map all of the urban areas of Japan for its 3D project. Mitsubishi will sell LIDAR-derived data products to a variety of clients, including Japan Space Imaging, which serves customers in automotive, government, and other markets.

TerraPoint: High-Fidelity Digital Earth Surface Data
TerraPoint LLC, a Transamerica Group company, was formed to convert LIDAR technology developed by NASA and the Houston Advanced Research Center to commercial applications that deliver precision digital-terrain data. TerraPoint Airborne Laser Topographic Mapping System (ALTMS), carried along a flightline at 140 knots by a dedicated twin-engine aircraft, directs 20,000 laser pulses a second to the ground from a height of 3,000 feet.

The pulses are reflected back to the aircraft by ground features. On the ground, LIDAR, GPS, and Inertial Measurement Unit (IMU) data from the flightline is downloaded to a pair of SGI Origin 200 servers. TerraPoint runs internally developed postprocessing software on the SGI Origin 200 servers to merge the three inputs into a single data set. The GPS data fixes the aircraft's geographic position in the XYZ planes; the IMU data describes the aircraft's pitch, yaw, or roll. Together they provide a precise measurement of ground features that has attracted customers worldwide whose businesses demand high-resolution modeling and mapping of terrain features.

Compute Power, Throughput, and Massive Storage Capability

TerraPoint projects generate huge amounts of data. "For a 20-mile-long flightline, you're probably talking about a gigabyte of data," says TerraPoint Data Processing Manager Dan Phillips. "We can have 16GB per day coming into our shop for processing. We need a very high-end processing system to handle these big data sets," says TerraPoint President Dan Cotter, who became used to big geographic databases during his 11 years with the U.S. Federal Emergency Management Agency. "Draw times are important. The labor cost of visual quality control becomes significant if you have to wait 10 or 15 minutes. So rendering speed is one reason why we acquired SGI Origin family technology." TerraPoint currently operates two SGI Origin 200 servers and uses 12 Octane workstations and 12 Silicon Graphics O2 workstations for image processing. The company has just acquired an additional 4TB SGI Fibre Channel RAID for a total of 6TB of storage, which is deployed in a storage area network environment to provide high-speed data access between the servers and the workstations.

The Future

TerraPoint selected SGI technology for its throughput, massive storage capability, and compute power. But SGI will play another critical role in TerraPoint's growth. Much of TerraPoint's future business will come from customers in overseas markets who need orthorectification data to enhance satellite and other aerial imagery. And because all of the Space Imaging ground stations around the world use SGI technology, TerraPoint will be working closely with SGI during its expansion.

For more information:
www.sgi.com/federal/solutions.html

All images provided by TerraPoint, LLC.

Virtual *Battlefield*

In 1994, when the U.S. Navy, Marine Corps, and Air Force, together with allies, announced their Joint Strike Fighter (JSF) program to develop and field an affordable next-generation strike fighter aircraft, the team at Lockheed Martin Corporation competing for the contract knew that extensive modeling and simulation technology would be required.

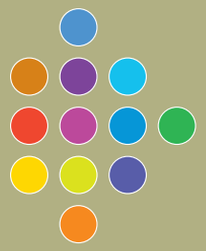

Affordability through commonality was a key criterion outlined in the mission of the program. The military customer wanted three different designs that would share key, high-cost components—e.g., engines, avionics, and many of the high-cost structural components. Common components would not only save money in the manufacturing, but also in common depot maintenance and service interoperability. The JSF program projected that the new family of aircraft would have commonality in the range of 70-90%, with emphasis on commonality in higher priced parts.

The demands for this new family of aircraft were complex. While sharing common characteristics—such as common fuselage lines, structure, systems, and software; a single-seat cockpit; a side weapons bay; and a blended swept wing—each of the services had also outlined specific characteristics for the JSF. The Navy wanted the JSF to be a multirole stealthy strike aircraft, the Air Force needed a multirole (primary-air-to-ground) fighter, and the Marine Corps required a multirole, short take-off vertical landing (STOVL) strike fighter. The United Kingdom Royal Navy and Royal Air Force specified a supersonic STOVL craft.To reduce costs in the development and testing of the JSF family of aircraft, Lockheed Martin's Aeronautics Company (LMAero) turned to SGI for its powerful, advanced suite of computers to simulate complex aircraft systems in every stage of design and testing.

Using SGI computing technology, Lockheed Martin has developed piloted simulators at the LMAero Flight Simulation Lab (FSL) in Fort Worth, Texas and Palmdale, California, which support design and evaluation of the flight controls and aerodynamics as well as the integrated tactical performance of the JSF.

Three flight control simulators support design and evaluation of the JSF advanced flight controls and aerodynamics. The Vehicle Management System (VMS) Control Law simulator in the FSL permits pilot assessment of the handling qualities of the designed aircraft, using simulations of the flight control and actuation systems. The VMS Hot Bench Simulator, also in the FSL, permits evaluations of the actual flight control computers with embedded airborne software. The VMS Integration Facility at the Palmdale facility integrates the flight controls, cockpit, and actuator hardware into a manned simulator located on a motion base platform. These simulators permit assessment of aircraft flight characteristics over the full spectrum of flight conditions, from taxi, takeoff, aerial maneuvering, and air-to-air refueling to landing on carrier decks and runways.

The Tactical System Simulator (TSS) facility supports assessment of the JSF tactical combat capabilities and includes integrated models of all primary on-board systems, including propulsion, aerodynamics, flight controls, weapons management, sensors, communications, pilot-vehicle interface, and avionics. In the TSS, pilots climb into the JSF cockpit and fly the aircraft in highly realistic simulated battlefield situations, including targets, threats, terrain, weather, and friendly forces. This allows pilots to evaluate system design in a realistic tactical environment and recommend changes early in the design process.

The heart of the TSS facility is the pair of full field-of-view displays, which surround the simulator cockpit with high-resolution imagery. The combination of Silicon Graphics Onyx2 computers and MultiGen-Paradigm image computation software presents the pilot with an all-aspect visual display suitable for both air-to-air and air-to-ground engagements.

The use of simulation to support aircraft design at the Ft. Worth facility began with the General Dynamics B58 Hustler, says Matt Landry, manager, simulation programs at the LMAero Flight Simulation Lab. But since then, the complexity of tactical simulation has increased to the degree of sophistication where pilots who fly the TSS say they cannot "game" the simulation due to the high degree of realism. "They said, 'The workload in the simulator is comparable to that in the battlefield,'" Landry said.

The ultimate test for virtual reality is actual reality. The Navy version of the Lockheed Martin JSF demonstrator, X-35C, had its first flight test on December 16, 2000. It was the second Lockheed Martin JSF Concept Demonstration Aircraft to undergo flight-testing. The X-35A, designed for U.S. Air Force use, successfully completed its flight-test program on November 22. A third JSF variant, the lift fan-equipped X-35B, is scheduled to begin vertical-flight testing later in 2001.

The JSF simulators have passed the reality test. "When we flew our first aircraft, the X35-A, the senior test pilot said it flew just like the simulator," Landry said proudly. "We are seeing a much stronger correlation between the simulator and the real world." The X-35A also set several flight test records, one of which was directly related to the simulator. The entire test program successfully concluded without the need to change any of the flight control system software. This was a result of the intense, detailed testing of the software in the simulator, in which all flight modes, failure modes, transitions, and interfaces were thoroughly tested before going to flight.

"All of LMAero's simulators use SGI computer horsepower," Landry said. "With recent improvements in processing power, SGI is now able to do computations that used to take a Cray system to run. This permits us to use system models of unprecedented complexity, which significantly increases the fidelity of the simulation."

Landry said one of the obvious advantages of using SGI machines is that more of the simulation tasks can be run on a single computer. Older simulators relied on special-purpose compute engines for computation of the complex, wide field-of-view display; nonlinear aerodynamics; in-cockpit graphics, and sensor imagery. With the new generation of SGI technology, all these computations can now be done on a single computer, reducing the cost of ownership, necessity for multiple development environments, and maintenance costs. A specific advantage of SGI, according to Landry, is its application to the full field-of-view display used in the TSS. This display immerses the pilot in a wrap-around visual scene consisting of eight back-projected images. Our fully immersive simulation environments require extremely powerful hardware to do distortion correction and image computation in real time using open-format databases. This is an ideal application of the SGI and MultiGen solution.

Simulations such as the TSS are beneficial to Lockheed Martin because they allow the company to assess the tactical effectiveness of its weapon systems. Before building a physical prototype costing millions of dollars, simulations allow Lockheed to test its designs in virtual space. "You want to quantify and demonstrate that the vehicle is controllable, lethal, and stealthy while assessing its tactical effectiveness," Landry said. "These models are high enough in fidelity to satisfy our goals. We are meeting our own high standards."

The current simulation technology is at an unprecedented level of fidelity, Landry said. A major area left for improvement is the out-the-window visual system, where the transport delays and resolution still need improvement. According to Landry, however, this is a narrow niche and he doesn't expect innovations to develop quickly. Another area where Landry sees improvement needed is the speed of computer interfaces.

Lockheed Martin's simulation initiatives are paying off. In November 1996, the JSF program entered the concept demonstration phase and selected two contractors, Boeing and Lockheed Martin, to build and fly concept demonstration aircraft. The successes realized in the first phase of flight test of the X-35A are in part due to the investment in the simulators. The final selection of one contractor for engineering and manufacturing development is scheduled for late 2001, after flight testing concludes.

For more information:
www.sgi.com/manufacturing/aerospace/
www.lmtas.com/

Images courtesy of Lockheed Martin Aeronautics Co.

Sometimes revolutions creep up on you. All the signs are there but you haven't seen them. Suddenly, there's a critical mass and then everything changes.

The Digital Content *Revolution*

In the world of information technology and communications something very much like that is happening. The problem is that we have all been so entranced with the detail we haven't seen the big picture.

Each of the parts – whether it's the Internet or wireless technology – seems so revolutionary in its own right, it demands our total attention. So we focus on one thing and we don't notice that it's the connections between many of these technologies that are really important.

That's just what's been happening with digital media. But now those connections are becoming so critical that it's imperative we drag our eyes away from the detail and understand the significance of that big picture. The problem is knowing where to start.

How about Florida, 1994? Time Warner, one of a new breed of global media giants formed by fusing the book and magazine publishing interests of Time Inc. with the movie and TV businesses of Warner Bros., and SGI, a company with a pedigree in high-performance computing and advanced graphics, were experimenting with interactive television. Two hundred bemused families had been presented with experimental TV sets incorporating SGI technology that allowed data streaming—the ability to download digital sound and pictures at will.

The experiment demonstrated that the TV set in the corner of the sitting room could become a window on a wide world of information and entertainment. But, in 1994, the technology was just too expensive to be appealing to people with other demands on their dollars.

In the few years since the Florida experiment, three big developments have made digital media the key technology revolution we all have to take seriously. First, costs have fallen as they always do with new technologies. Buying access to interactive TV is no longer a big deal for people in Florida or other parts of the developed world.

Second, broadband communication technologies are spreading like the ganglia of a giant brain. In Britain, five million homes—20% of the total—are projected to have broadband connections by 2005. In the U.S., broadband connections are predicted to grow from just over four million in 2000 to around 15 million by 2003. Business use of broadband is also on a steep upward curve—approaching 200 million subscribers worldwide by 2003. Broadband is the key to harnessing the full power of digital media, whose applications invariably use fatter data so they need more bandwidth.

Third, there's been an explosion in the amount of digital content available. The Forrester Report, an IT market research authority, says content distribution deals are set to quintuple by 2003. The size of the World Wide Web has burgeoned so that, in 2000, it's estimated there were 1.2 billion pages of information available. And, significantly, content that people would previously have expected to access in other formats—such as music, photographs, and video—is increasingly available digitally.

Digital media has changed the landscape so rapidly that many organizations are now struggling hard to keep up with both the business and the technological challenges. One critical issue is the need to repurpose content so that it can be reused in more than one format. For example, an ad agency developing a multimedia campaign might find that it has to repurpose the same material for TV broadcast, Internet sites, and CD-ROMs as well as conventional print formats (although even these may be produced using digital technology).

Repurposing presents new opportunities to exploit material for profit, but it also poses a raft of technological problems. It means that people who have been dealing in one form of media—with a limited palette of technologies—suddenly have to master many more. It's a formidable challenge, not least because most managers already have their hands full finding ways to grab the new business opportunities that digital media dangle tantalizingly before them.

Welcome to the world of Media Commerce, where digital content is created, managed, delivered, and transacted across networks. It's the world of the future, full of rich business opportunities. It's transforming marketplaces and making new products and new ways of doing business possible. And it's posing some mighty challenges. So just what is the extent of these challenges?

Let's start with content creation. Whether it's writing a book, taking a newspaper photograph, or producing a cartoon video, digital technology has revolutionized the way you do it. The newspaper photographer's three key pieces of equipment are now a ladder, a laptop, and a digital camera. Today, artists work with graphics software that helps them not only create the background scenes but even animate the characters within them.

The second essential activity if you want to exploit digital content effectively is content management. Any organization that is making plans to exploit its digital content seriously will soon find it's developing a wider range of content in a bewildering range of formats, so it needs to develop policies to rationalise formats and store content—in other words, an asset management strategy.

But its biggest challenge is to generate revenue from the content—after all, that's the business payoff. That means using increasingly sophisticated indexing software. More companies will find they have opportunities to grow revenue from intellectual property in the years ahead. But much of that intellectual property will exist in digital form. Content distribution — the third of the activities—is also important. There is little point in storing and managing large amounts of digital content if it can't be accessed by customers who want to use and (hopefully) pay for it. Digital content distribution is already creating its own range of problems, as digital content is often extremely easy to copy. So increasingly, content owners will be looking to develop distribution strategies and mechanisms that prevent illegal copying yet provide easy access for legitimate users.

The final piece of Media Commerce is the transaction. This is where the content is transferred from the owner to the user and, where appropriate, the payment is made. Broadly speaking, digital content transactions are going to be made in one of three ways—over the Internet, through a broadcast medium (such as interactive TV) or by telephony. There will be growth in all of these, but some of the most far-reaching developments in the near future are likely to be with mobile telephony, especially using the Universal Mobile Telephone System. For example, if you're entering a shopping center, content providers can transmit information about special offers available at stores within that center directly to your handset.

What's all this going to mean for companies that want to develop and exploit their own digital content? At a strategic level, most companies will need to do some deep thinking about the kind of business model that will best enable them to exploit their own digital content. That won't be easy, simply because we are stepping into new territory here and there are no rules to guide us. So one watchword must be flexibility.

Another key issue is scalability. Companies need to be able to scale their operations—and their IT delivery infrastructure—as demand expands. Often, scalability is the key to capturing an emerging market ahead of competitors. It's impossible to summarize the kinds of new services the digital content revolution will make possible simply because there are so many. But consider two examples where SGI has played a role in helping companies deliver innovative new services and win business benefit.

ShareAPhoto is a company that provides a service that allows people to convert their video and photographic material footage to show on the Internet. Previously, anybody who wanted to load photographs or video onto the Internet faced a complex technical task to convert the video to a different format. ShareAPhoto is providing an easier way to upload video directly from digital cameras into online albums using a Web browser. (See page 27 for more information).

Meanwhile, Denmark Radio's (Denmark's public sector TV company) conversion to digital production illustrates some of the key lessons about content management and distribution. The TV station is creating a system that not only enables it to produce, store, and distribute news broadcasts but also helps it manage and archive material afterwards. As a result, Denmark Radio benefits by developing long-term digital storage for its radio, television, and multimedia assets.

Whatever kind of problem you have to solve, there are a lot of critical issues to tackle in developing a business strategy for digital content. As with all business, risk management will be important, and that means finding a technology partner with the broad experience across all areas of these fast changing technologies, a role SGI expects to play for an ever-expanding cast of companies.

Faxback Form

We hope you have enjoyed reading this issue of SYNERGI, the SGI magazine of imagination, innovation, and insight. For further information on solutions provided by SGI, please fax back the form to the number below.

Please continue to send SYNERGI to me:

Name _____ Title _____

Company _____

Address _____

_____

City _____ State _____ ZIP _____

Phone _____ Fax _____ E-mail _____

(651) 683-7115

Attention: R. Fraser

1. What is your company's primary industry ? (Please circle one only.)

☐ Aerospace ☐ Education ☐ Medical
☐ Arch/constr ☐ Entertainment ☐ Petroleum/energy
☐ Automotive ☐ Gov't defense ☐ S/W development
☐ Chemical ☐ Gov't research ☐ Other

2. What applications best describe your interest areas? (Please circle all that apply.)

☐ AEC ☐ Earth resources ☐ Publishing
☐ Animation ☐ ECAD ☐ S/W development
☐ Chem/bio, chem/tech ☐ MCAD ☐ Virtual reality
☐ Medical imaging ☐ Visual simulation ☐ Other

3. Are you currently in the process of evaluating workstations or servers for purchase now or in the foreseeable future?

☐ Yes ☐ No ☐ Don't know

4. If yes, when do you expect to make a purchasing decision?

☐ 1–30 days ☐ 30–60 days ☐ 60–90 days ☐ 90+ days

5. What type of systems or services are you evaluating?

☐ UNIX® workstations ☐ UNIX servers
☐ Linux workstations ☐ Linux servers
☐ Windows NT workstations ☐ Windows NT servers
☐ Storage ☐ Flat panel displays
☐ Professional services ☐ Education services
☐ Support services ☐ Productivity services

6. Has a budget been approved for this time frame?

☐ Yes ☐ No ☐ Don't know

7. What is your IT budget for your current project?

☐ $50,000 ☐ $50,000–100,000 ☐ $100–200,000
☐ $200,000–300,000 ☐ $300,000–500,000 ☐ $500,000+
☐ n/a

8. What is your role in the decision-making process?

☐ Final decision maker ☐ Input only/recommender
☐ Purchaser/buyer ☐ Information gatherer
☐ End user

9. Please have an SGI representative contact me:

☐ ASAP ☐ 30–60 days ☐ 60–90 days ☐ 90+ days

10. Please contact me via e-mail regarding:

☐ Linux
☐ SGI Internet Server, SGI Media Server
☐ SGI customer services
☐ High-performance computing
☐ SGI Reality Center
☐ CFD
☐ CAD
☐ Silicon Graphics Onyx2
☐ Business intelligence, data mining
☐ Clustering
Other_____

11. Would you like to receive information via e-mail in the future?

☐ Yes ☐ No

Check out the latest SGI news and products at *www.sgi.com/*

At Toronto's Hospital for Sick Children, lovingly known as Sick Kids, genetic science leaves the lab and enters the real world: a world where some children are running out of time. Using high-performance data management, computing and visualization solutions designed by SGI, life scientists move quickly through mountains of genomic data to discover possible cures–and create hope. Dr. A. Jamie Cuticchia, head of bioinformatics at Sick Kids, stresses, "The identification and cure of genetic diseases is essential to the Sick Kids mission. Without supercomputer capabilities, our scientists would spend hours waiting." SGI gives research hospitals the power to find more answers, more quickly. And stay ahead. To learn more about SGI™ solutions, or for information on our services, consulting and support, visit our Web site.

www.sgi.com/ahead

sgi™

One step ahead

SGI in the U.S./Canada

U.S. Central Region:
39001 West Twelve Mile Road
Farmington Hills, MI 48331-2903
Phone: (248) 848-4400
Fax: (248) 848-5600

U.S. Western Region:
6200 LBJ Freeway—Suite 100
Dallas, TX 75240
Phone: (972) 788-4122
Fax: (972) 788-0957

Canada:
5800 Explorer Drive
Mississauga, Ontario L4W 5K8
Phone: (905) 625-4747
Fax: (905) 625-4476

U.S. Eastern Region:
890 Winter Street
Waltham, MA 02451
Phone: (781) 839-2100
Fax: (781) 672-5810

SGI Federal:
12200-G Plum Orchard Drive
Silver Spring, MD 20904
Phone: (301) 572-1980
Fax: (301) 572-3280

www.sgi.com

SGI Worldwide

Corporate Headquarters
SGI
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
USA
Phone: (650) 960-1980

North America	(800) 800-7441
Canada	(905) 625-4747
Latin America	(52) 5267-1387
Europe	(44) 118-925-75-00
Asia Pacific	(65) 771-0290
Japan	(81) 3-5488-1811

3082 (05/2001)