Success Story

Strathclyde University

"At the time, HP's Linux systems didn't scale beyond 16 processors, whereas SGI already was supporting systems with as many as 128 processors. And we had long experience with SGI, always receiving very good support and completely reliable products. Our Origin system, for instance, has never crashed in over three years of operation."

– Gian-Luca Oppo,
Professor,
Computational and Non-Linear Physics,
Strathclyde University

Strathclyde University Pushes Beyond Physics with SGI® Altix® and InfiniteStorage

For some researchers, conducting advanced research in the simulation of complex physical systems once was a nail-biting experience that dragged on for months. Without experience in parallel processing systems, many physics researchers had little choice but to run large numerical experiments on single-processor servers. Some problems ground on for half a year in an attempt to identify behavior of, for example, a beam of light for a tiny fraction of a second.

Researchers at the University of Strathclyde knew the problem well. "Several research groups were wasting so much time loading large problems onto single-CPU systems and waiting forever to see results," recalls Gian-Luca Oppo, professor of computational and non-linear physics at the 14,000-student university in Glasgow, Scotland. "You could wait two, three, or even six months to get any kind of feedback on your problem. Meanwhile, what happens if the system crashes?"

Strathclyde and in particular Oppo's virtual photonics research group opted for high-performance computing (HPC)-class resources in 2001. A 16-processor SGI® Origin® 300 system had served as a powerful platform for multi-processor physics applications since then. But more and more, Strathclyde's administration found that researchers in other scientific fields of study could benefit from the kind of performance boost that the Origin system offered to physicists. Chemistry, mathematics, biosciences and architecture students accustomed to working on PCs or Sun workstations all expressed interest in sharing a powerful resource with physics researchers – so long as that resource was user-friendly.

"We're committed to open-source technologies, and we were very impressed with the performance and scalability of [HP and SGI's] Itanium2."

– Gian-Luca Oppo, Professor, Computational and Non-Linear Physics, Strathclyde University

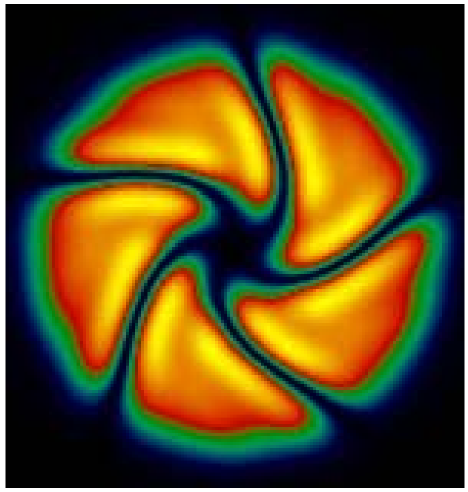

Optical sprinkler

More than a 'white box' cluster

Clearly, Strathclyde needed a new answer for the growing number of researchers with big-data problems. But few students outside the physics department had much contact with parallel operating systems. And the inherent limitations of UNIX® as a broadly shared computing environment suggested to the scientists and the administrators that they should pursue a new direction.

Traditional "white box" clusters wouldn't work for many of the applications targeted at the new system. "Clusters are fine for a lot of tasks, but they lack flexibility," says Oppo. "If you have a very good idea about what the code is going to do, you can configure a cluster to maximize your output. But if you have different numerical methods, memory demands and other factors, clusters can get exceptionally difficult."

Sharing a computational facility at Strathclyde meant that a Linux® operating environment was a crucial element, as was a system driven by Intel® Itanium® 2 processors. "We're committed to open-source technologies," Oppo recalls, "and we were very impressed with the performance and scalability of Itanium2." Given those criteria, Oppo and his colleagues evaluated systems from HP and SGI.

"At the time, HP's Linux systems didn't scale beyond 16 processors, whereas SGI already was supporting systems with as many as 128 processors," says Oppo. "And we had long experience with SGI, always receiving very good support and completely reliable products. Our Origin system, for instance, has never crashed in over three years of operation."

Number crunching and data management

To handle the stream next-generation problems they anticipated from throughout the university, Strathclyde installed a 28-processor SGI® Altix® 3000 system in January of 2004, and an upgrade to 32 processors soon followed. Since users wanted to run large problems without having to partition data sets, Strathclyde installed 32GB of globally addressable memory in its new Altix system.

"Some codes—like matrix calculations that are common in virtual photonics and spectroscopy—generate a large amount of data in the first part of the code. Once it's in your temp file, the second part of the code picks it up and starts to process it. And once that code is processed, you simply don't need that big chunk of data at all. So now we have a system that does an exceptional job of handling big input and output files.""

– Gian-Luca Oppo,
Professor,
Computational and Non-Linear Physics,
Strathclyde University

The Altix architecture allows global addressing of all memory in the system. For the first time, more complex data sets and complete workflows can be driven entirely out of memory, enabling productivity breakthroughs that traditional Linux clusters or repurposed e-commerce UNIX servers can't tackle. Altix systems feature a fully supported, standard 64-bit Linux operating system and advanced development environment specifically optimized for technical applications.

Strathclyde's new Altix server is backed by a 1 terabyte SGI® InfiniteStorage solution, and the SGI InfiniteStorage Shared Filesystem CFXS™ provides researchers with a high-availability, easily managed Storage Area Network (SAN) environment. "One of the issues we had was speed of access," says Oppo. "We needed a storage solution that gave lots of users fast access to a huge scratch disk—basically providing them with a giant temp file."

With SGI® InfiniteStorage Filesystems, Strathclyde users can transparently share data anywhere without replication. CXFS provides instant data access and sharing among systems on the SAN without network mounts or copies. The filesystem eliminates network data overhead, latencies and copies, speeding workflow and reducing complexity, backup time, and necessary disk space.

For long-term storage, Strathclyde users rely on their own local solutions. This, says Oppo, allows up to 75 percent of the SGI disk array to remain empty when no jobs are running on the Altix. "Data sets in these problems can be huge," he says. "Some codes—like matrix calculations that are common in virtual photonics and spectroscopy—generate a large amount of data in the first part of the code. Once it's in your temp file, the second part of the code picks it up and starts to process it. And once that code is processed, you simply don't need that big chunk of data at all. So now we have a system that does an exceptional job of handling big input and output files."

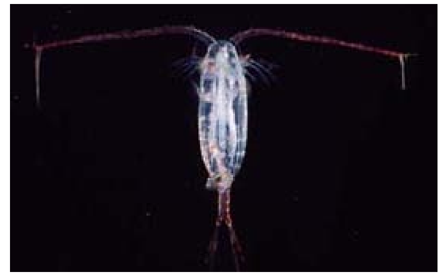

Copepods (crustaceans)

Results in hours, not months

The new SGI solution will help Strathclyde chart new courses for a broad array of computational research. In virtual photonics—the study of lasers and optical fibers which are the foundation of high-speed communication—researchers are answering an urgent need to simulate the performance of new photonics devices even before they're built. "Imagine the compute power required to accurately simulate a laser pulse lasting a million times less than a nanosecond," says Oppo. "Now imagine what it takes to simulate that performance over a period of hours. This is the work we are doing."

Parallel numerical simulations on Altix allow researchers to "leapfrog" those miniscule time steps to conduct detailed analysis of the output of these devices via new numerical techniques developed at Strathclyde. With the Altix system, Oppo says he and his colleagues can "transform exceedingly long simulations in large memory, multi-processor tasks that can provide accurate answers in hours instead of several months of computation time."

"[Oppo and his colleagues can] transform exceedingly long simulations in large memory, multi-processor tasks that can provide accurate answers in hours instead of several months of computation time."

– Gian-Luca Oppo, Professor, Computational and Non-Linear Physics, Strathclyde University

By hosting workshops to help port other codes to an HPC parallel environment like Altix and InfiniteStorage, the physics group has been joined by other scientists in utilizing this powerful compute and data management solution. Statisticians at Strathclyde, for instance, are studying the relationships between time and development stages to understand how to minimize the impact of heavy fishing on copepod plankton—a dominant foodstuff of aquatic creatures in sub-arctic waters of the North Atlantic. The Altix system, says William Gurney, professor of statistics and modeling science, allows researchers to employ a novel discrete-space-time approach to model ocean circulation and determine plankton reproduction patterns and influences. Using Altix combined with Strathclyde's unique methodology, says Gurney, has accelerated the research, making it "many orders of magnitude faster than conventional methods."

Other researchers are leveraging Altix in their search for a long-term alternative to oil and nuclear power as a sustainable and clean generation of energy. Strathclyde researchers are characterizing "heavy species" – atomic elements that are heavier than zinc – to determine how they might damage future energy sources based on controlled fusion. Other heavy species studies are focused on investigating the next-generation of etching techniques for the semiconductor industry. For this computationally intensive work, a symmetric multiprocessing system with large and fast memory caches is invaluable.

Oppo says he hopes to be able to scale Strathclyde's Altix system even beyond 32 processors, as funding allows. The vision of a powerful resource for broad research appears to dovetail nicely with the university's evolving profile as a research institution: In early 2005, Strathclyde physics joined forces with several universities in Scotland to attract funding that will help create "super teams" of researchers to pursue ground-breaking studies in such areas as space physics, new materials and, naturally, photonics. And while the university's Altix system today is available only to Strathclyde researchers, it may one day become a national resource.

"People come to the Altix system when they have good ideas but face simulation problems," Oppo says. "That's why it's good to have a shared system where solutions are exchanged. This in turn creates new ideas – and often even bigger simulation problems."

sgi

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000

3789 [04.2005] J14866