INDUSTRY BRIEF

Faster Results, More Discovery

Large Memory Systems for Life Science Research

Research environments today require significant amounts of computing power to analyze and simulate complex interactions of molecular and experimental data. As chemists, biologists and researchers push the boundaries of research, powerful systems are required that combine not only fast processing but very fast data access. The SGI Altix UV high performance computing system is designed specifically for large, data-intensive problems, to support Computational Biologists/Chemists and Bio- and Chem-Informaticists and others to more rapidly simulate processes and correlate these simulations with large amounts of experimental data. In addition, assembling raw sequence fragments from genome sequencers (in order to reconstruct the original sequence), can require massive amounts of memory in order to perform well.

The Altix UV system is a shared memory server, which means that all processing cores can directly address all memory in the system. Combining up to 256 Intel® Xeon® 7500 series processors (2048 cores) in a single OS instance, the Altix UV system provides applications the ability to address up to 16TB of memory to support unprecedented simulation complexity and resolution. In addition, Altix UV runs standard off-the-shelf Linux®, so that application developers can seamlessly move their scalable software from a Linux based laptop to a full Altix UV 1000 system and realize tremendous speedups. The maximum memory that the Altix UV system can accommodate is 8000 times more than a typical laptop and over 100 times more than a standard 2-socket server.

Recently, the Smith-Waterman algorithm for performing local protein sequence alignment analyzed 10 million combinations of protein sequences in 81.1 seconds. The algorithm leveraged the Pervasive DataRush™ platform to scale across an SGI Altix UV 1000 with 384 cores to achieve a sustained throughput of 986 billion cell updates/second. Dramatically demonstrating Pervasive DataRush scalability, the same executable used on a four-core system dynamically scaled across all 384 cores on the Intel® Xeon® processor 7500 series based server. The performance on the Altix UV 1000 system shattered the previous performance record by 43 Percent.

Key to Altix UV performance and scalability is the SGI NUMAlink® interconnect fabric. In addition to supporting shared memory programs, SGI has implemented a Message Passing Interface (MPI) Offload Engine that accelerates even cluster applications that are by nature distributed and that have significant amounts of communication between cores. Altix UV supports many types of multi-core development models, including MPI, OpenMP, Shared Memory ™ (SHMEM), and Partitioned Global Address Space (PGAS) languages.

SGI believes that for scientists, the focus should be on science, not the difficulties of distributed programming APIs or the management of complex computer systems. The Altix architecture enables easy installation and administration of a single system, while SGI's data management solutions simplify the storage and archiving of large data sets. The Altix UV product line can be easily integrated into an existing datacenter environment as well, where it can serve as a "super node," supporting applications which are too demanding for less capable "thin node" servers.

The Altix UV system has been designed to work with the latest acceleration technologies for even faster simulations. For applications that would benefit from GPGPUs, SGI offers the NVIDIA Tesla products which can be easily integrated into the Altix UV environment.

Customers utilizing the Altix UV: ICR, Minnesota Supercomputing Institute

The Altix UV is available in 3 model options.

- **Altix UV 10** – In a 4U rackmount form factor, Altix UV 10 packs up to four processors (32 cores), 64 DIMM slots and ten I/O expansion slots. With up to 1TB of shared memory, Altix UV 10 is ideal for a diverse set of technical computing and Enterprise applications.

- **Altix UV 100** - Addresses the mid-range market based on an industry-standard 19" rackmount 3U form factor. Altix UV 100 scales to 96 sockets (768 cores) and 6TB of shared memory in two racks. Altix UV 100 delivers up to 6.9 teraflops of compute power in a single system image.

- **Altix UV 1000** - For maximum scalability, Altix UV 1000 ships as a fully integrated cabinet-level solution with up to 256 sockets (2,048 cores) and 16TB of shared memory in four racks. Altix UV 1000 delivers up to 18.6 teraflops of compute power in a single system image.

For more information, visit www.sgi.com/altixuv.

Shared memory systems are ideal for a wide range of scientific applications where large amounts of data need to be accessed quickly and without the complexity of programming for a distributed environment. By utilizing memory which has much faster access than rotating disk drives (3 orders of magnitude), entire workflows can be accelerated and new science and discoveries can be made. Increased physics, model resolution and data gathering are all simplified on a system with a large shared memory.

INDUSTRY BRIEF

Customer Focus: Institute for Cancer Research

The Institute for Cancer Research is one of the world's leading cancer research organizations, internationally known for the quality of its science and has been named UK's top research center. ICR's focus is to relieve human suffering by pursuing excellence in the fight against cancer through:

- Research into the causes, prevention, diagnosis and methods of cancer treatment
- Education and advanced training of medical and scientific staff
- Treatment and patient care of the highest quality

> *"Systems biology demands massive integration of extremely large data sets. Large shared memory should enable us to handle such data at a much higher speed and with a greater focus on the biological questions at hand," said Peter Rigby, chief executive officer at ICR. "Altix UV should significantly help our work in this new, exciting area of cancer research."*

The ICR joins the growing list of globally significant high performance computing (HPC) facilities embracing Altix UV as the future of open, high performance, big-memory super-computing. Altix UV is providing the ICR with a massively scalable shared memory system to process its growing data requirements, including hundreds of terabytes of data for biological networks, MRI imaging, mass-spectrometry, phenotyping, genetics and deep-sequencing information across thousands of CPUs.

The solution that the ICR chose was the SGI Altix UV 1000 supercomputer, which supports up to 16TB of global shared memory in a single system image. It remains highly efficient at scale for applications ranging from in-memory databases to a diverse set of data and compute-intensive HPC applications; Altix UV is the only hardware solution equipped to meet the vast data processing requirements of the ICR.

The SGI® Altix® UV 1000 system that the ICR recently purchased from SGI includes:

- 512 cores of Intel® Xeon® processor 7500 series
- 4TB memory
- Water cooled
- SGI® Infinite Storage IS5000

SGI leads as an "applications-driven" organization, with strong expertise and evidence in helping customers achieve real science, research collaborations and develop best of breed technologies SGI became trusted advisor early to the ICR, working with the researchers to define the best system for their unique environment.

For more information about the Institute for Cancer Research visit: http://www.icr.ac.uk/

> *"The Altix UV supercomputer will allow extremely large, diverse data sets to be processed quickly, enabling our researchers to correlate medical and biological data on an unprecedented scale," said Dr. Rune Linding, cellular and molecular logic team leader at the ICR. "Eventually, this will lead to network-based cancer models that will be used to streamline the process of drug development."*

INDUSTRY BRIEF

Customer Focus: Minnesota Supercomputing Institute

The University of Minnesota Supercomputing Institute for Advanced Computational Research (MSI) features an SGI® Altix® UV 1000 supercomputer, to accelerate its research program. This new system, named Koronis, is made possible by a National Institutes of Health (NIH) grant awarded to the University of Minnesota in June 2010.

Assembling a team of more than 30 University of Minnesota faculty members in the life sciences community with similar specialized compute needs, including four major user groups in multi-scale modeling, chemical dynamics, bioinformatics and computational biology, and biomedical imaging. Professor Darrin York from the Department of Chemistry, with the assistance of HPC specialists at both SGI and MSI, submitted the grant proposal to the NIH in May 2009. The proposal was for a complete High End Computing (HEC) solution that would accommodate the needs of a broad spectrum of NIH-supported researchers. This HEC solution will facilitate breakthroughs in biomedical research and significantly impact human health.

The MSI Altix UV 1000 System:

- 1152 cores of Intel® Xeon® processor 7500 series
- 3.1TB memory
- SGI® InfiniteStorage IS4600
- High Performance visualization systems

MSI named the system Koronis after a Minnesota lake; Koronis will also include high-performance visualization workstations to handle large-size biomedical data and high-performance and -fidelity data storage to ensure adequate data analysis and processing.

These high-performance visualization workstations will enable researchers to develop next-generation sequencing, which can help provide a more complete picture of genomes, producing thousands of gigabytes of data. With next-generation sequencing, along with other large-sized data, Koronis will facilitate breakthroughs in biomedical research.

Professor York will develop new theoretical tools for biocatalysis and multi-scale quantum models for Ribozyme Catalysis, using computational methods to simulate complex chemical reactions that are catalyzed by biological molecules and, in particular, RNA. These molecular simulations, using so-called “multi-scale quantum models,” provide potentially powerful tools for studying catalytic mechanisms as well as detailed characterization of the transition state ensemble that are difficult, if not currently impossible, to directly observe by experiment. The simulations involve tens to hundreds of thousands of atoms, and require sampling tens to hundreds of millions of configurations along the reaction path. To carry out these large-memory, data-intensive computational simulations, MSI required a supercomputer that can handle an enormous capacity of data, which they found in SGI’s Altix UV 1000.

For more information about the Minnesota Supercomputing Institute, visit: https://www.msi.umn.edu/

Corporate Office
46600 Landing Parkway
Fremont, CA 94538
tel 510.999.8300
fax 408.321.0293
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5185.2860
Europe +44 118.912.7500
Asia Pacific +61 2.9448.1463