

A History of High-Performance Computing

Originating Technology/NASA Contribution

Faster than most speedy computers. More powerful than its NASA data-processing predecessors.

Able to leap large, mission-related computational problems in a single bound. Clearly, it's neither a bird nor a plane, nor does it need to don a red cape, because it's super in its own way. It's Columbia, NASA's newest supercomputer and one of the world's most powerful production/processing units.

Named Columbia to honor the STS-107 Space Shuttle Columbia crewmembers, the new supercomputer is making it possible for NASA to achieve breakthroughs in science and engineering, fulfilling the Agency's missions, and, ultimately, the Vision for Space Exploration.

Shortly after being built in 2004, Columbia achieved a benchmark rating of 51.9 teraflop/s on 10,240 processors, making it the world's fastest operational computer at the time of completion. Putting this speed into perspective, 20 years ago, the most powerful computer at NASA's Ames Research Center—home of the NASA Advanced Supercomputing Division (NAS)—ran at a speed of about 1 gigaflop (one billion calculations per second). The Columbia supercomputer is 50,000 times faster than this computer and offers a tenfold increase in capacity over the prior system housed at Ames. What's more, Columbia is considered the world's largest Linux-based, shared-memory system.

The system is offering immeasurable benefits to society and is the zenith of years of NASA/private industry collaboration that has spawned new generations of commercial, high-speed computing systems.

Partnership

To construct Columbia, NASA tapped into years of supercomputing experience, dating as far back as the early 1980s, when computational fluid dynamics (CFD) computer codes originated, and as recent as 2004, when the Agency adopted novel immersive visualization technologies

to safely pilot the Spirit and Opportunity Mars Exploration Rovers. In addition, NASA looked to Silicon Valley for some extra support and found a friend it had helped back in the heyday of early microprocessing technology.

In the first few years of the 1980s, Ames scientists and engineers assisted Mountain View, California-based Silicon Graphics, Inc. (SGI), by providing technical input to improve the company's high-performance workstation product line. NASA had purchased 18 of SGI's IRIS workstations and helped make them commercially viable

with several improvements. By 1984, NASA was SGI's biggest customer.

"NASA was a huge help to us as a young company, not only by being our biggest customer at a time when a lack of sales would have been disastrous, but they were one of our best customers in the sense that the engineers there gave us all sorts of valuable feedback on how to improve our product. Many of the improvements to the original workstations are still part of our most modern products," according to Tom Davis, former principal scientist and a founding member of SGI.



A bird's-eye view of the 10,240-processor SGI Altix "Columbia" supercomputer located at Ames Research Center. While Columbia is helping NASA achieve breakthroughs and solve complex problems in support of its missions, it has also been made available to a broader national science and engineering community.

SGI's payback to NASA was helping to build the behemoth Columbia supercluster. Santa Clara, California-based Intel Corporation, the world's largest computer chip maker and a leading manufacturer of computer, networking, and communications products, also assisted in the effort. Through extraordinary dedication and uncompromising commitment, the Columbia project team achieved what many in the supercomputing community considered impossible: conceiving, planning, and constructing the world's largest Linux-based, shared-memory system in just over 4 months.

The resulting system is an SGI Altix supercomputer system, based on SGI's NUMAflex shared-memory architecture for high productivity. It is comprised of 20 SGI Altix integrated superclusters, each with 512 processors; 1 terabyte of memory per 512 processors, with 20 terabytes total memory; 440 terabytes of online storage; and 10 petabytes of archive storage capacity (1 petabyte is equal to 1,024 terabytes, and 1 terabyte is equal to 1,024 gigabytes).

"NASA's indomitable spirit of exploration has led us to the Moon, to the surface of Mars, and even to the rings of Saturn," said Bob Bishop, vice chairman of SGI. "With Project Columbia, NASA will not only carry mankind further into space, but into new worlds of knowledge and understanding. After 2 decades of collaboration, NASA and SGI are on the cusp of a new age of scientific method and scientific discovery."

Product Outcome

A portion of the Columbia system has been made available on a broad basis to ensure the Nation's entire science and engineering community has access to the highly advanced supercomputer architecture. For example, throughout the 2004 hurricane season, the finite-volume General Circulation Model (fvGCM) running on Columbia had cranked out valuable, real-time numerical weather-prediction data targeted at improving storm tracking and intensity forecasts. A team at Goddard Space

Flight Center is utilizing the data to predict landfall up to 5 days in advance.

Additionally, Jet Propulsion Laboratory, Massachusetts Institute of Technology, and Scripps Institution of Oceanography scientists from a consortium called Estimating the Climate and Circulation of the Ocean (ECCO) teamed with the NAS Division to use the supercomputer to dramatically accelerate the development of a highly accurate analysis of global-ocean and sea-ice circulations. The ECCO team produces time-evolving, three-dimensional estimates of the state of the ocean and of sea-ice. These estimates, obtained by incorporating into a numerical model vast amounts of data gathered from instruments in the ocean and from space satellites—such as sea level, current speed, surface temperature, and salinity—serve as a practical tool to better understand how the ocean currents affect Earth's climate, to study the role of the ocean in the Earth's uptake of carbon dioxide, and to more accurately predict events like El Niño and global warming.

Meanwhile, NASA continues to lend technical advice to support the advancement of SGI's products. The lessons learned while SGI provides NASA with engineering prototype systems is helping to improve the scalability and reliability of the machines. When SGI developed a 256-processor high-performance system for Ames, the experience directly benefited the company's commercial 128-processor machines. When NASA doubled its 256-processor system to operate on a 512-processor system, SGI made the 256-processor system commercially available. Ames again doubled up (prior to having Columbia) by moving to a 1,024-processor system, leading SGI to make the 512-processor system an official commercial product.

"The main product outcome has been the development of larger and larger general purpose, single-system image machines that are practical and usable, not just curiosities," said Bron Nelson, a software engineer with SGI. "This is driven by Ames and SGI's belief that these




Image courtesy of Silicon Graphics, Inc., and Stony Brook University.

Using an SGI Altix system to successfully model how the HIV protease molecule works across time, researchers hope to determine how best to target it with drugs that could stop it from doing its job and thus prevent the HIV virus from developing altogether.

large, single-system image machines help to improve programmer productivity and ease-of-use, as well as ease of system administration."

Whether it is sharing images to aid in brain surgery, finding oil more efficiently, enabling the transition from analog to digital broadcasting, helping to model Formula 1 race cars and Ford Motor Company vehicles, or providing technologies for homeland security and defense, SGI has committed itself to working with NASA to ensure that it is putting out the best product possible and committed its resources to addressing the next class of challenges for scientific, engineering, and creative uses. ♦

Linux® is a registered trademark of Linus Torvalds.

Silicon Graphics®, SGI®, and Altix® are registered trademarks of Silicon Graphics, Inc.

NUMAflex™ is a trademark of Silicon Graphics, Inc.