

July 15, 2003

Teaching Computers to Work in Unison

By STEVE LOHR

Computers do wondrous things, but computer science itself is largely a discipline of step-by-step progress as a steady stream of innovations in hardware, software and networking pile up. It is an engineering science whose frontiers are pushed ahead by people building new tools rendered in silicon and programming code rather than the breathtaking epiphanies and grand unifying theories of mathematics or physics.


Yet computer science does have its revelatory moments, typically when several advances come together to create a new computing experience. One of those memorable episodes took place in December 1995 at a supercomputing conference in San Diego. For three days, a prototype project, called I-Way, linked more than a dozen big computer centers in the United States to work as if a single machine on computationally daunting simulations, like the collision of neutron stars and the movement of cloud patterns around the globe.

There were glitches and bugs. Only about half of the 60 scientific computer simulations over the I-Way worked. But the participants recall those few days as the first glimpse of what many computer scientists now regard as the next big evolutionary step in the development of the Internet, known as grid computing.

"It was the Woodstock of the grid — everyone not sleeping for three days, running around and engaged in a kind of scientific performance art," said Dr. Larry Smarr, director of the California Institute for Telecommunications and Information Technology, who was the program chairman for the conference.

The idea of lashing computers together to tackle computing chores for users who tap in as needed — almost as if a utility — has been around since the 1960's. But to move the concept of distributed computing utilities, or grids, toward practical reality has taken years of continuous improvement in computer processing speeds, data storage and network capacity. Perhaps the biggest challenge, however, has been to design software able to juggle and link all the computing resources across far-flung sites, and deliver them on demand.

The creation of this basic software — the DNA of grid computing — has been led by Dr. Ian Foster, a senior scientist at the Argonne National Laboratory and a professor of computer science at the University of Chicago, and Dr. Carl Kesselman, director of the center for grid technologies at the University of Southern California's Information Sciences Institute.



Dr. Ian Foster has helped create the basic software for grid computing.
(Photo by Peter Kiar.)

They have worked together for more than a decade and, a year after the San Diego supercomputing conference, they founded the Globus Project to develop grid software. It is supported mainly by the government, with financing from the Department of Energy, the National Science Foundation, NASA and the Defense Advanced Research Projects Agency.

There has been a flurry of grid projects in the last few years in the United States, Europe and Japan, most of them collaborations among scientific researchers at national laboratories and universities on projects like climate modeling, high-energy physics, genetic research, earthquake simulations and brain research. More recently, computer companies including IBM, Platform Computing, Sun Microsystems, Hewlett-Packard and Microsoft have become increasingly interested in grid technology, and some of the early commercial applications include financial risk analysis, oil exploration and drug research.

This month, grid computing moved further toward the commercial mainstream when the Globus Project released new software tools that blend the grid standards with a programming technology called Web services, developed mainly in corporate labs, for automated computer-to-computer communications.

Enthusiasm for grid computing is also broadening among scientists. A report this year by a National Science Foundation panel, "Revolutionizing Science and Engineering Through Cyberinfrastructure," called for new financing of \$1 billion a year to make grid-style computing a routine tool of research.

The long-term grid vision is that anyone with a desktop machine or hand-held computer can have the power of a supercomputer at his or her fingertips. And small groups with shared interests could find answers to computationally complex problems as never before.

Imagine, for example, a handful of concerned citizens running their own simulation of the environmental impact of a proposed real-estate development in their community. They wouldn't need their own data center or consultants. They would describe what they want, and intelligent software would find the relevant data and summon the computing resources needed for the simulation.

"The ultimate goal is a fundamental shift in how we go about solving human problems, and a new way of interacting with technology," Dr. Kesselman said.

That grand vision, however, is years away, perhaps a decade or more. Dr. Smarr is the former director of the National Center for Supercomputing Applications at the University of Illinois, where Web browsing software later used by both Netscape Communications and Microsoft was developed in the 1990's.

He compares the state of grid computing now to the Web in 1994, when groundbreaking work in a new technology had come from the elite science labs in the United States and Europe but before commercial investment had gathered momentum.

The grid is widely regarded as the next stage for the Internet after the World Wide Web. The Web is the Internet's multimedia retrieval system, providing access to text, images, music and video. The promise of the grid is to add a problem-solving system.

Computer scientists say the contribution of Dr. Foster and Dr. Kesselman to grid computing is

roughly similar to that made by Tim Berners-Lee to the development of the Web. Mr. Berners-Lee, who is now the director of the World Wide Web Consortium at the Massachusetts Institute of Technology, came up with the software standards for addressing, linking and sharing documents over the Web: U.R.L.'s (uniform resource locators), HTTP (hypertext transfer protocol) and HTML (hypertext mark-up language).

The heart of the grid problem is managing and linking computing resources. Dr. Foster and Dr. Kesselman, assisted by another software designer at the Argonne lab, Steve Tuecke, have devised basic grid standards with their own acronyms: GRAM (Globus resource allocation manager), M.D.S. (monitoring and discovery service), G.S.I. (grid security infrastructure) and GridFTP (grid file transfer protocol).

The wisdom of their work, according to computer scientists, lies in its farsighted simplicity, designing a set of minimalist standards that others can build upon. It is the same design philosophy, they note, found in the original Internet and the Web.

"If you look at the history of computer science, the people who have had the biggest impact are the ones who envisioned big systems and then came up with simple but smart mechanisms for building those systems," said Dr. Ken Kennedy, a computer science professor at Rice University. "That's what Ian Foster and Carl Kesselman have done."

In 1998, after they developed some early working software, the Globus leaders had to decide the best way to proceed. After long discussions, they chose not to make Globus commercial. Instead, they opted for the open-source model, in which computer code is openly shared, allowing programmers to modify, improve and fix the software. The decision, Dr. Foster recalled, was both practical and ethical.

"Our belief was that open source was the best way to maximize adoption," he said. "Globus is an infrastructure technology, and it is only going to be successful if everyone uses it. And if you're doing something that is primarily funded by the government, sharing the software seemed the most appropriate thing to do."

Grid computing is a far bigger challenge than simpler forms of distributed computing. Today, most grid projects remain the province of supercomputing centers and university labs. The research centers are linked by network connections about 20 times as fast as the standard high-speed connections and are equipped with storage systems able to handle vast data files and high-performance computers.

The Biomedical Informatics Research Network, begun in 2001 and supported by the National Institutes of Health, is a grid created to help scientists gain a better understanding of the way the brain works.

One project, called the Brain Morphometry BIRN, involves pooling and processing magnetic resonance imaging data to look for early anatomical and functional precursors of Alzheimer's disease. That knowledge may then be used to tailor drugs to inhibit the onset of the disease.

Researchers from Harvard, Duke, the University of North Carolina, Johns Hopkins, the University of California at Los Angeles, the University of California at San Diego, Massachusetts General Hospital and Brigham and Women's Hospital participate in the study. From their desktop computer,

any of them can tap into data anywhere on the brain project's grid.

"BIRN is a leading example of how you use this cyberinfrastructure to make team science happen to achieve stretch goals in research," said Dr. Mark Ellisman, the neuroscientist at the University of California at San Diego who led the design of the BIRN data grid.

The collaborative computing tools for data sharing are also fostering a new style of research. "We're helping a scientific community to understand that it does more good to make information more generally accessible than squirreling it away," Dr. Ellisman said.

Dr. Foster, the software tool maker, is encouraged by the applications built on his group's underlying technology.

"Like nearly everything in computer science, the work we've done on the Globus software is incremental," he said. "But it is having an impact. There are thousands of people doing collaborative, computing-intensive work in a variety of fields that they could not do before."

[Copyright 2003 The New York Times Company](#)