

BY M. MITCHELL WALDROP

ILLUSTRATION BY HOLLY LINDEM

COULD PUT THE PLANET'S INFORMATION-PROCESSING POWER ON TAP.

Grid Computing

Is Internet history about to repeat itself?

Maybe. Back in the 1980s, the National Science Foundation created the NSFnet: a communications network intended to give scientific researchers easy access to its new supercomputer centers. Very quickly, one smaller network after another linked in—and the result was the Internet as we now know it. The scientists whose needs the NSFnet originally served are barely remembered by the online masses.

Fast-forward to 2002. This summer, the National Science Foundation will begin to install the hardware for the TeraGrid, a transcontinental supercomputer that should do for computing power what the Internet did for documents. First, clusters of high-end microcomputers will be set up at four sites: the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign; the U.S. Department of Energy's Argonne National Laboratory outside Chicago; Caltech in Pasadena, CA; and the San Diego Supercomputer Center at the University of California, San Diego. Then, by early next year, those four clusters will be networked together so tightly that they will behave as a single entity.

This virtual computer will rip through problems at up to 13.6 trillion floating-point operations per second, or teraflops—eight times faster than the most powerful academic supercomputer available today. Such speed will enable scientists to tackle some of the most computationally intensive tasks on the research docket—from problems in protein folding that will form the basis for new drug designs to climate modeling to deducing the content and behavior of the cosmos from astronomical data.

But more than that, the TeraGrid will be a prime example of what has come to be known as “grid computing”—the massive integration of computer systems to offer performance unattainable by any single machine. The integration of these systems will be so transparent that users will no more notice they are on a network than motorists pay attention to which cylinder is firing at any given moment. To people logging onto the TeraGrid, the system will look like just another set of pro-

grams running on their office computers. But that look will be deceptive: what appear to be applications that reside on the local desktop machine might actually be data analysis tools running on the cluster at San Diego, or visualization software crunching bits at Argonne. The “files” TeraGrid users are working on might consist of databases scattered all over the country, containing thousands of gigabytes—a.k.a. terabytes.

Grid computing visionaries hope that this will be only the beginning—that the \$53 million TeraGrid will catalyze a new era of grid computing for the masses, much as the NSFnet broke down barriers that led to the blossoming of the Internet. Just within the past year or two, dozens of such projects have been announced in Europe, Asia and the United States, with more likely to come. And the developers of grid computing are now settling on a single standard—called the Globus Toolkit—that will help grid projects under development all around the world coalesce into a worldwide network of tappable computer power.

“Completely transformational” is how Larry Smarr, director of the California Institute for Telecommunications and Information Technology, sums up grid computing. Smarr, renowned for his role in developing the communications system that evolved into the Internet's backbone, says the technology is what the Internet has been building toward for the past three decades. “In the first phase,” he explains, “we got the wires up and hooked in all the computers. Then with the World Wide Web, we started hooking in all the online documents.” Now, he says, with grid computing, we'll be hooking in everything else (*see “Planet Internet,” TR March 2002*).

This means that users will begin to experience the Internet as a seamless computational universe. Software applications, databases, sensors, video and audio streams—all will be reborn as services that live in cyberspace, assembling and reassembling themselves on the fly to meet the tasks at hand. Once plugged into the grid, a desktop machine will draw computational horsepower from all the other computers on the grid. “What we're seeing,” says Smarr, “is the

emergence of a new infrastructure upon which first science, and then the whole economy, will be built.”

COMPUTING AS UTILITY

That's a tall order. But it certainly describes the hope at IBM, which is the prime contractor for the TeraGrid, as well as for similar national grids in Europe. David Turek, vice president of emerging technologies for IBM's server group, compares grid computing to the familiar grid of electrical power: “To use a hair dryer, you just plug it into a wall socket,” he says. “You don't have to worry about how the turbine is designed up in Niagara Falls, or the physics of power transmission.” That's exactly how Turek wants people to think about computing power. “In our vision of the future, if you're a customer who occasionally needs 10 teraflops, for example, don't buy a machine that's underutilized most of the time; buy it from the grid. So grid computing will play into our vision of computing as a utility.”

While companies like IBM would build the large-scale grids, Turek says that many users will want to set up grids of their own. “You might see 10 to 20 departments coming together to create a campuswide or companywide grid, each contributing some of the computer power they control,” he says. In another scenario, several independent companies, such as defense contractors, might do much the same thing to create “virtual organizations”—ad hoc grids that would allow them to use one another's proprietary data and software to prepare, say, a proposal for a new military aircraft. “That's why we're not going to espouse the grid as something that can be done only with IBM technology,” Turek explains. After all, he says, “if you get five companies wanting to come together on a grid, the likelihood of all five having the same servers is pretty slim.”

And that, Turek adds, is the beauty of the Globus Toolkit: a set of open-source software tools that is fast emerging as the de facto standard for grid computing, in much the same way that the hypertext transfer protocol, or HTTP, is the standard for linking documents on the Web.

Indeed, the growing acceptance of Globus is largely responsible for today's wave of grid computing excitement.

"The idea is to let the network provide the basic mechanisms for moving data around, while Globus provides mechanisms for resource sharing," explains Carl Kesselman of the University of Southern California's Information Sciences Institute. Kesselman has been developing the Globus Toolkit over the past five years in collaboration with Ian Foster—a University of Chicago computer scientist who heads Argonne's distributed-systems laboratory.

The mechanisms that Globus provides are as essential to the computing grid's operation as stoplights are to city traffic. One set of Globus software tools, for example, automatically roots out where on the grid a required database or program can be found. Other tools allow one-time login, so that the user isn't constantly being asked for passwords for site after site after site. Still others divide a computational job into multiple sub-tasks and parcel them out among the various systems on the grid. And most important, Globus provides tools to

implement security—assuring, for instance, that an outside program trying to interact with your machine is serving a legitimate purpose and hasn't been sent by some malicious hacker.

Of course, none of this is entirely new: "It's worth remembering," notes Kesselman, "that ARPAnet [the military-built ancestor of the Internet] was built in the 1960s to give users on one campus shared access to resources on a different campus." Likewise, he points out, methods for breaking computational jobs into smaller pieces for multiple machines were a perennial research topic throughout the 1970s and 1980s.

But it was only in the 1990s, Kesselman says, that the rapidly increasing power of computers and networks brought this trend, known as distributed computing, out of the laboratories. One result was a flurry of experiments in what is now known as "peer-to-peer" computing, all devoted in one way or another to harnessing the computing power and storage capacity of idle desktop machines. Among the best known of these efforts are Napster, the MP3 music file-sharing system, and SETI@home, in which radio


telescope data from the search-for-extraterrestrial-intelligence project are distributed to PCs across the Internet.

At the same time, however, the high-performance-computer community began a series of less publicized but much more ambitious experiments in "meta-computing." The idea was to make many distributed computers function like one giant computer. The metamachine's keyboard and display would be sitting on someone's desktop, as usual. But its central processor might actually be a super-computer in Illinois, say, while its graphics processor might be an immersive-virtual-reality facility in California. It worked, says Kesselman—the only problem being that experimenters had to reinvent the wheel every time. "There was still no standard software for distributed computing," he says, "no infrastructure to support it."

The technology's watershed event came in 1995, at a supercomputing conference sponsored by the Institute of Electrical and Electronics Engineers and the Association for Computing Machinery. There, 11 separate high-speed networks were briefly connected into one

Grid Work

Many computers are better than one—especially for solving tough scientific problems



giant metacomputer in a demonstration called I-Way. Attendees thronging the San Diego Convention Center could play with an interactive model of the Chesapeake Bay ecosystem, or a high-resolution simulation of colliding spiral galaxies—some 60 applications in all. Foster, who led the team that created some of the system's underlying software, was especially impressed by I-Way's potential use in collaborative design. In one demonstration, he recalls, researchers at Argonne teamed up with those at an industrial group, Nalco Fuel Tech, to make a virtual-reality simulation for designing incinerators. "Users at different sites could fly together through the incinerator, place injectors in it at various points and jointly study the effect on its output," he recalls.

The demonstration had its intended effect. "I-Way convinced people that grid computing had great potential," says Foster. One important payoff was that in October 1996, the U.S. Defense Advanced

Research Projects Agency funded Kesselman and Foster's Globus project to provide a solid foundation for grid computing. At the 1997 supercomputer conference, Foster and Kesselman demonstrated a grid with some 80 sites worldwide running Globus software—another feat that, in Foster's view, "convinced people that grid computing was worthwhile and real." At that point, moreover, Foster and Kesselman had even started to *call* it "grid computing," playing on the analogy to the electrical grid.

PHYSICS AND BEYOND

Once the concept was introduced, grid computing suddenly seemed to fill a need of scientists all over the world. In Geneva, for example, the high-energy physics lab of the European Organization for Nuclear Research (known by the acronym CERN) was already planning its next-generation particle accelerator, the Large Hadron

Collider—an effort promising to generate an overwhelming amount of data. "We estimated that when the collider started running in 2006 it would produce eight to 10 petabytes of particle collision data per year," says Fabrizio Gagliardi, director of CERN's annual seminar on computing for physicists. That's *petabytes*—millions of gigabytes.

Portions of this immense data load would have to be distributed to the institutions all over the world that participate in CERN experiments. And since the most interesting physics tends to be found in the rarest events, Gagliardi explains, scientists "would be processing every bit of that data in multiple ways"—looking for hints of the theoretically predicted but elusive Higgs boson, say, or particles that possess the mysterious quality known as supersymmetry. In short, the collider portended an enormous data management problem for which existing computer systems seemed inadequate. "We

The Growing Grid

A sampling of grid computing projects sprouting up all over the world

PROJECT	LAUNCHED	SPONSOR	MAIN PURPOSE
Access Grid www-fp.mcs.anl.gov/fl/accessgrid	1999	U.S. Department of Energy, National Science Foundation	Internet-based collaboration—including lectures and meetings—among scientists at facilities around the world
European Data Grid www.eu-datagrid.org	2001	European Union	Data analysis in high-energy physics, environmental science and bioinformatics
Grid Physics Network (GriPhyN) www.griphyn.org	2000	NSF	Data analysis for four physics projects: two particle detectors at CERN's Large Hadron Collider, the Laser Interferometer Gravitational Wave Observatory, and the Sloan Digital Sky Survey
Information Power Grid www.ipg.nasa.gov	1999	NASA	Computational support for aerospace development, planetary science and other NASA research
International Virtual DataGrid Laboratory (iVDGL) www.ivdgl.org	2002	NSF and counterparts in Europe, Australia, Japan	World's first truly global grid: will link high-performance-computer centers in Europe, Australia, Japan and the United States
Network for Earthquake Engineering and Simulation (NEESgrid) www.neesgrid.org	2001	NSF	Integrated computing environment for 20 earthquake engineering labs
TeraGrid www.teragrid.org	2002	NSF	General-purpose infrastructure for U.S. science: will link four sites at 40 gigabits per second and compute at up to 13.6 teraflops
U.K. National Grid www.grid-support.ac.uk	2001	U.K. Office of Science and Technology	Support for grid projects within Britain
Unicore www.unicore.de	2000	German Federal Ministry for Education and Research	A seamless interface to high-performance-computer centers at nine government, industry and academic labs

defined a computational architecture for what we would need,” Gagliardi recalls. “Then we went shopping for a system of tools to build it—and discovered that the computer scientists had already come up with solutions.”

Several solutions, actually. At the University of Virginia, computer scientist Andrew Grimshaw had been working since 1993 on an attractive and well-thought-out set of grid computing protocols known as Legion. (Legion is now being marketed by Avaki of Cambridge, MA, which Grimshaw founded.) But Globus had the advantage of being “open”: in the interests of getting it

Foundation. It focuses on the vast amount of physical data generated by four different sources: two specialized particle detectors housed at the Large Hadron Collider; the Laser Interferometer Gravitational Wave Observatory, a Caltech-MIT collaboration that will detect gravitational waves from pulsars and the like; and the Sloan Digital Sky Survey, an international effort to map the faintest possible stars and galaxies—more than 100 million celestial bodies in all. More recent initiatives include the NSF’s Network for Earthquake Engineering Simulation grid, an effort to integrate observations and computer

made from this technology? “If computing is a utility,” Foster says, “who’s going to pay for the infrastructure? What kind of services are people prepared to pay for?” In particular, where is the killer app, the must-have application that will drive the growth of grid computing the way the spreadsheet did personal computing? Most current grid projects have barely moved past the if-we-build-it-they-will-come stage.

On the other hand, says Foster, “we do have some ideas.” One notable example is the Access Grid, an Argonne-developed system—based, like so much else in grid computing, on Globus—that

THE COMING EXPLOSION OF ACTIVITY COULD CREATE A WORLD OF INTERLINKED COMPUTER GRIDS—A DEVELOPMENT DWARFING THE INTERNET BOOM OF THE 1990s.

adopted as widely and as rapidly as possible, Foster and Kesselman had decided to emulate the developers of the now famous Linux operating system and make the Globus source code available to any users who wanted it, so that they could study it, experiment with it and suggest improvements.

The result was that Globus became the foundation for the European DataGrid, a three-year demonstration and software development project that launched on January 1, 2001, with a commitment of 13.5 million euros (roughly \$12 million) from the European Union. By the beginning of 2002, the DataGrid had deployed more than 100 computers—20 at CERN, the others at sites around the continent, according to Gagliardi, now the DataGrid’s director. The project has also expanded beyond particle physics to include two other scientific disciplines that face similarly daunting data-crunching and processing challenges: earth observation and biology.

Meanwhile, grid computing has been finding an even warmer welcome among scientists in the United States—with Globus again being the choice of virtually every large project. One of the first to get going was the Grid Physics Network. Organized by Foster and University of Florida physicist Paul Avery, this effort was launched in September 2000 with \$11.9 million from the National Science

simulations now scattered among some 20 different labs, with the goal of producing more effective designs for earthquake-resistant structures.

And now, of course, there’s the TeraGrid—the “put-your-money-where-your-mouth-is grid,” as Argonne’s Charles Catlett calls it. “We’ve been talking for years,” says Catlett, the project’s executive director. But for the TeraGrid to achieve what it promises, the high-powered microcomputer clusters located at its four physical sites will have to be tied together by a dedicated network running at 40 gigabits per second, which will be right on the ragged edge of the state of the art. “This will show us a lot about how the software really works in a production environment,” says Catlett. He’s talking about the Globus software, the Internet protocols, the Linux operating system—all of it.

On the technical side, Catlett says, one of the big challenges is making sure that Globus can successfully scale up. It is critical, he notes, to make sure that Globus’s services and protocols “can deal with hundreds or thousands of times more devices than they handle now.” “Obviously,” agrees Foster, “there is lots that still needs to be done.”

Then there’s the business side. Here, grid computing runs into the same question that sank so many of the over-optimistic dot coms: how will money be

supports large-scale, multisite meetings over the Internet, as well as lectures and collaborative work sessions. It already links more than 80 academic and industry sites around the globe. Furthermore, says Foster, as more and more big scientific projects like the TeraGrid and the DataGrid come on line, there’s every reason to think that they will serve as laboratories for new grid applications that will then make their way into the commercial world, with huge impact. After all, the Internet’s killer app, the World Wide Web, didn’t come out of a corporate lab. It came out of CERN.

GRID UNLOCKED

While the Web may be a tough act to follow, grid computing advocates have been paving the way for the technology’s hoped-for commercialization by focusing on such nitty-gritty issues as standards-setting. “Remember how much we’ve gained from the fact that every computer runs the Internet Protocol,” says Foster. To achieve the same universality for grid computing, the U.S. grid community has merged with those of Europe and Asia to form the Global Grid Forum—an organization patterned after the Internet’s standards-setting body, the Internet Engineering Task Force. The forum’s goal is to make sure that Globus, Legion and any other grid protocols can

interoperate seamlessly. “If every computer uses standard methods for managing authentication, authorization, describing resource capabilities and negotiating access for resources,” says Foster, “that’s a big win.”

The grid pioneers are likewise building alliances with their counterparts in commercial peer-to-peer computing. In practice, however, peer-to-peer efforts appear to be most effective for problems that can easily be broken into myriad small, independent pieces—a category that does *not* usually include, say, the complex physics simulations and virtual-immersion applications

Microsystems and Veridian in the United States, together with Fujitsu, Hitachi and NEC in Japan—announced that they would implement the Globus Toolkit on their machines as a standard platform for grid computing. Then early this year, Microsoft completed a contract with Argonne to translate the existing Globus Toolkit to Windows XP, according to Todd Needham, manager of the software giant’s University Research Programs group.

If nothing else, Microsoft’s move should hasten the day when home and office computers will be able to join the grid by the millions, just by plugging in. But perhaps just as significantly, it also

dwarf even the Internet boom of the 1990s. In the future envisioned by Smarr, grids of every size will be interlinked. The “super-nodes,” like TeraGrid, will be networked clusters of supercomputers serving users on a national or international scale. The more numerous mid-sized nodes will use software such as Entropia to harness the power of multiple desktop and laptop PCs. If the TeraGrid and other supernodes are like central electric power stations, Smarr explains, these smaller nodes will be like solar energy collectors that capture a diffuse yet enormous resource.

Still more numerous will be the millions of individual nodes: personal

GRID USERS WILL EXPERIENCE THE NET AS A SEAMLESS COMPUTATIONAL UNIVERSE. SOFTWARE, DATABASES AND SENSORS WILL BE REBORN AS SERVICES ASSEMBLED ON THE FLY.

where grid computing really shines. Nonetheless, Foster says, the potential for synergy is clear. That’s why the Globus protocols have already been integrated into such industrial-strength peer-to-peer systems as the Condor protocols developed at the University of Wisconsin-Madison and the Entropia platform from Entropia of San Diego, both of which are designed to capture the unused capacity of an organization’s networked workstations.

The payoff for such efforts is that the computer industry now seems to be taking grid computing very seriously indeed—with the most notable example being IBM. Last August, at the same time it won the contract to build national grids in the United Kingdom and the Netherlands, as well as TeraGrid in the United States, Big Blue announced that it would “grid-enable” many of its server systems. This initiative, which would mean that servers in many institutions and organizations could be plugged into grid networks quickly and easily, was said to be as big or bigger than IBM’s commitment to Linux, which already stood at roughly \$1 billion. (Indeed, IBM had already used Globus to link its own R&D labs in the United States, Israel, Switzerland and Japan.)

Yet IBM is hardly alone. Last November, eight other computer makers—Compaq, Cray, Silicon Graphics, Sun

symbolizes the fast-developing alliance between grid computing and “Web services,” a similar technology that has emerged independently over the past few years and has been embraced in slightly different forms by Microsoft, IBM and Sun, among others. Like grid computing, the Web services idea revolves around future software applications that are created on the fly out of programs and data that live on the Internet, not the user’s machine. The main difference between this idea and grid computing is that Web services software tends to be much more closely tied to the World Wide Web protocols, as well as to Web-based standards such as XML.

Once again, however, as Microsoft and IBM’s embrace of Globus suggests, the potential for synergy is obvious. In January, Foster, Kesselman, IBM’s Jeffrey Nick and Argonne’s Steven Tuecke proposed an Open Grid Services Architecture that would integrate the two approaches, and announced that this framework would be implemented as version 3.0 of the Globus Toolkit. IBM, Microsoft, Platform Computing, Entropia and Avaki announced their support of the new architecture, with other companies to follow.

And in the future? History is indeed about to repeat itself, declares grid computing advocate Smarr—except that the explosion of grid activity may very well

machines that users plug into the grid to tap its power as needed. If, say, the members of a citizen’s group were worried about a proposed development project, they could use the grid to run the same simulations that the developers and government officials involved used. That way, they could easily see the effect of the development on everything from ground water to traffic patterns to employment. By using grid-based tele-immersion technologies, the citizens could even walk through the simulated project and get a realistic sense of what it would feel like to *be* there.

And thanks to the wireless revolution, “micronodes” will be everywhere. “Because of the miniaturization of components,” says Smarr, “we’ll have billions of endpoints that are sensors, actuators and embedded processors. They’ll be in everything, monitoring stress in bridges, monitoring the environment—ultimately, they’ll even be in our bodies, monitoring our hearts.”

And that, he emphasizes, is why we have to lay a solid foundation for the grid now, building in security and all the rest from the start. “We can’t do it as an afterthought,” he says. “The planet is assembling the grid infrastructure that it will live on for the rest of 21st century.” ■

To join an online discussion of this article, visit www.technologyreview.com/forums/grid