

# Meet



# The Grid

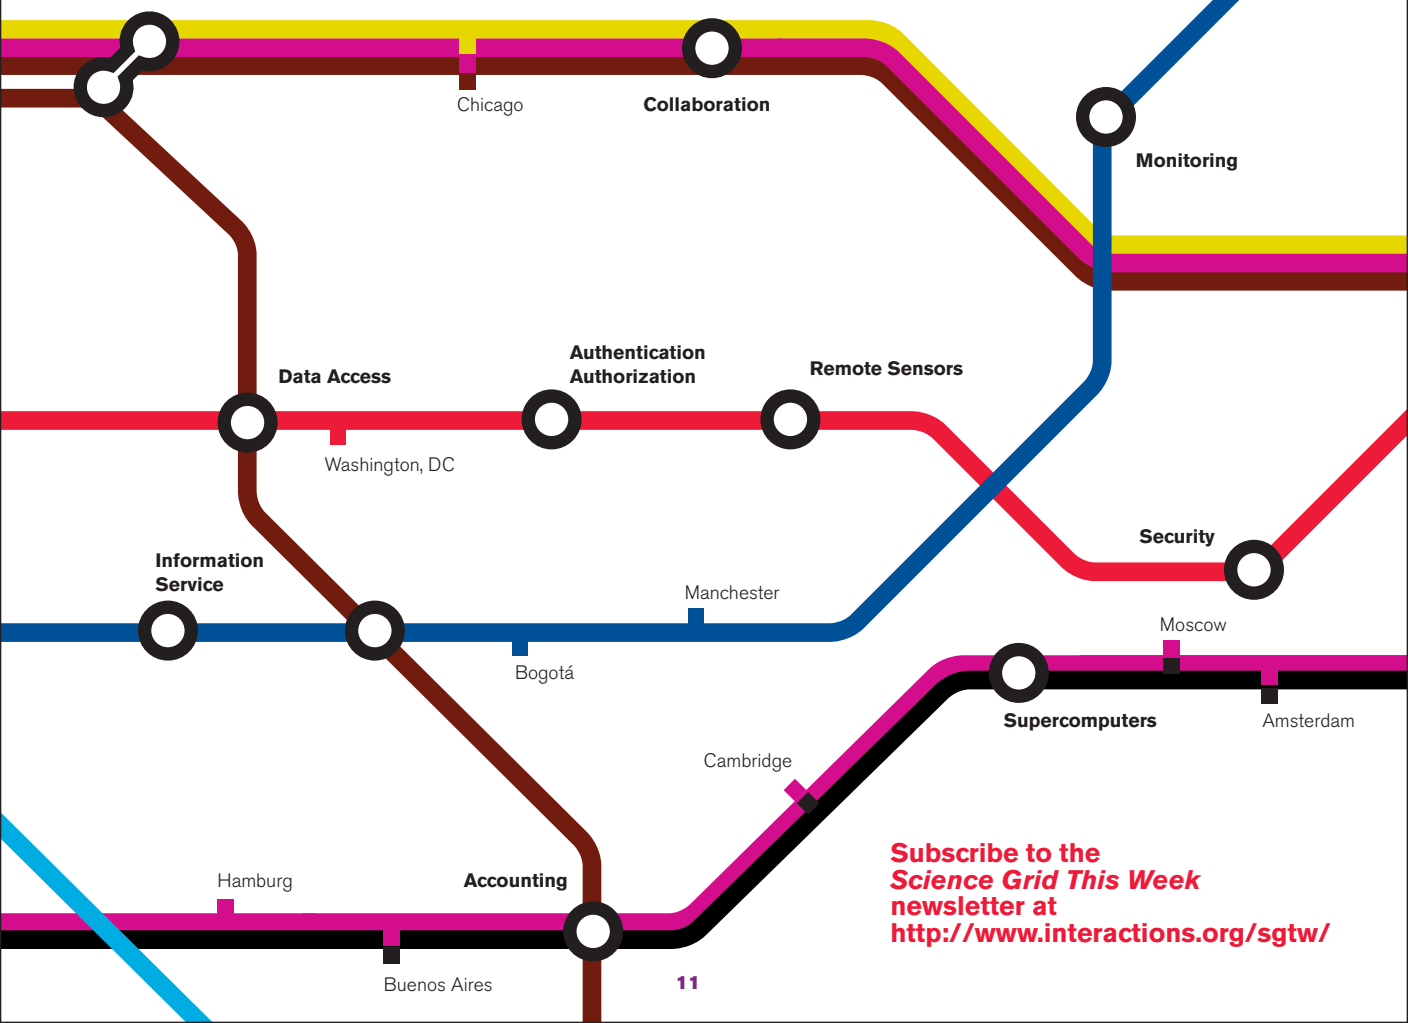
Today's cutting-edge scientific projects are larger, more complex, and more expensive than ever. Grid computing provides the resources that allow researchers to share knowledge, data, and computer processing power across boundaries. **By Katie Yurkewicz**

Oliver Gutsche sits in a quiet warren of cubicles in Fermilab's Wilson Hall, concentrating on his computer screen, ignoring the panoramic view from his 11th-floor window. He's working feverishly toward a deadline less than two years away, when over 5000 scientists will participate in the largest and most international grid computing experiment ever conducted.

Gutsche is a member of the Compact Muon Solenoid (CMS) particle physics experiment, one of four experiments being built at the Large Hadron Collider at CERN in Switzerland. When the LHC, which will be the world's highest-energy particle accelerator, begins operating in 2007, vast amounts of data will be collected by its experiments. Scientists worldwide will need to sift through the mountain of data to find elusive evidence of new particles and forces.

"The Compact Muon Solenoid experiment will take 225 megabytes of data each second for a period equivalent to 115 days in 2008," says Gutsche. "That means each year we'll collect over two petabytes of data."

One petabyte is a lot of data (you'd need over 1.4 million CDs to hold it), and the LHC experiments will collect petabytes of data for many years. Any single institution would be hard-pressed to store all that data in one place and provide enough computing power to support thousands of eager scientists needing daily access. Thus the LHC experiments and other scientific collaborations count on a new way to securely share resources: grid computing.



Subscribe to the  
**Science Grid This Week**  
newsletter at  
<http://www.interactions.org/sgtw/>

## The grid vision

The term grid arose in the late 1990s to describe a computing infrastructure that allows dynamic, distributed collaborations to share resources. Its pioneers envisioned a future where users would access computing resources as needed without worrying about where they came from, much like a person at home now accesses the electric power grid.

One such pioneer, Carl Kesselman from the Information Sciences Institute at the University of Southern California, says, "In today's society, scientists more often than not operate within an organizational structure that spans many institutes, laboratories, and countries. The grid is about building an information technology infrastructure for such virtual organizations."

Virtual organizations (VOs) are typically collaborations that span institutional and regional boundaries, change membership frequently, and are governed by a set of rules that define what resources are shared, who is allowed to share them, and the conditions under which sharing occurs. Grids provide the hardware and software that allow VOs to get things done. As cutting-edge scientific tools become larger, more complex, and more expensive, researchers increasingly need access to instruments, data, and collaborators not in their home institution.

"It's no exaggeration to say that we'll collect more data in the next five years than we have in all of human history," says Microsoft's Tony Hey, former director of the UK e-Science project. "Grid computing and e-Science will allow new, exciting, better-quality science to be done with this deluge of data."

## Applications in many fields

In fields such as biology and geology, grids will enable scientists to bring together vastly different types of data, tools, and research methods to enable scientific breakthroughs.

"For example, there is a group of researchers in England studying the effects of proteins on heart cells," explains Hey, "and another group in New Zealand with a mechanical model of a heart responding to an electrical stimulus. If the two groups access each other's data, modeling programs, and computing resources, someday they might be able to determine exactly how a certain protein produced by a genetic defect induces an unusual electrical signal that leads to a heart attack."

Particle physicists like Gutsche belong to some of the largest scientific collaborations using grids today. Gutsche's work takes place within the Open Science Grid (OSG), a grid computing project that will provide the framework for US physicists to access LHC data from the ATLAS and CMS experiments. His project is to get a data analysis program running on the OSG, so that the 500 US CMS collaborators can create data sets tailored to their individual research interests. First, however, Gutsche needs to understand the inner workings of grids.

Like many particle physicists, Gutsche's career requires him to be a part-time computer scientist. Computing expertise, a history of international collaborations, and data-intensive science have led physicists to be founders or early adopters of many distributed computing technologies such as the Internet, the Web, and now grid computing, an application of distributed computing.

"The banking system was an early example of distributed computing," explains Ian Foster from the University of Chicago and Argonne National Laboratory, who, with Kesselman, published *The Grid: Blueprint for a New Computing Infrastructure*, in 1998. "The system is very focused on moving information around for a very specific purpose using *proprietary* protocols. In grid computing, there is an emphasis on bringing together distributed resources for a variety of purposes using *open* protocols."

## How grids work

Grids will enable scientists, and the public, to use resources, access information, and connect to people in ways that aren't possible now. A unified grid-computing

## Where's my grid?

Many members of large science collaborations already have specialized grids available to advance their research. Those not so fortunate may well have access to shared resources through one of the many multidisciplinary projects sprouting up worldwide. These grids range from "test beds" of only a few computers to fully-fledged projects sharing vast resources across continents.

University researchers and students use campus grids, such as the Grid Laboratory of Wisconsin (GLOW) at the University of Wisconsin-Madison, or the Nanyang Campus Grid in Singapore, to share computing resources from different departments. Universities that haven't caught the grid bug yet can join state-wide or regional projects such as the

North Carolina Statewide Grid, which will benefit business, academia, and government when completed, or SEE-Grid, which includes 10 countries in south-eastern Europe.

National grid computing projects abound. In Japan, researchers in academia and business use NAREGI, and over 25 European countries have access to either a national grid or to the European Union-funded Enabling Grids in E-science (EGEE) infrastructure. In the United States, collaborations can share and access resources using the Open Science Grid (OSG) infrastructure, or apply for time on the TeraGrid, which links high-end computing, storage, and visualization resources through a dedicated network.

system that links people with resources is made up of four layers of resources and software stacked on top of each other. Each layer of this grid architecture depends on those below it. The network layer is at the base. It connects all of the grid's resources, which make up the second layer. On top of the resources sits the middleware, the software that makes the grid work and hides its complexity from the grid user. Most people will eventually only interact with the uppermost software layer, the applications, which is the most diverse layer. It includes any program someone wants to run using grid resources.

Using the grid, a scientist could sit down at her computer and request, for example, a climate prediction for the next 10 years. She would open the appropriate grid-adapted application and provide the geographic location as well as the time range for the prediction. The application and middleware will do the rest: make sure she's a member of a VO that allows her to access climate resources; locate the necessary historical data; run a climate prediction program on available resources; and return the results to her local computer.

The steps that have been taken so far toward the seamless grid vision have been made possible by a sudden increase in network connectivity over the last decade.

"Companies had installed a lot of optical fiber, thinking they'd be able to sell it at a good profit," says Stanford Linear Accelerator Center's Les Cottrell. "When the bubble burst, instead of selling it to other companies at a fire-sale rate, many of the businesses were willing to negotiate deals with academic and research organizations. As a result, for example, the SLAC connection to the Internet backbone has increased over 60 times since 2000, from 155 megabits per second to 10 gigabits per second, and we can send ten times as much data across transatlantic lines in the same amount of time."

The availability of better network hardware, software, and management tools allows scientists to use the grid to share more than just processing power or files. The resource layer of the grid, connected by high-speed networks, also includes data storage, databases, software repositories, and even potentially sensors like telescopes, microscopes, and weather balloons.

### The brains of a grid

Making many different networks and resources look like a unified resource is the job of the middleware—the "brains" of a grid. The types of middleware used by a grid depend on the project's purpose.

"If you're building an information grid, where you're collecting information from 15,000 radio antennas, you focus on information services," explains Olle Mulmo from the Royal Institute of Technology in Sweden. "If you're making lots of CPUs

available, you focus on resource management services. You'll need data services if you're compiling a lot of data and making it available to others. The fourth main category is security, which works across all other categories."

Enforcing a VO's rules about who can access which resources, making sure that access is secure, and keeping track of who is doing what on a grid creates some of the most challenging problems for grid developers. The middleware must provide the solutions.

"Without security—authorization, authentication, and accounting—there is no grid," explains Fabrizio Gagliardi, project director for the Enabling Grids for E-science project funded by the European Union.

The stringent requirements for security and accounting differentiate grid computing from other distributed computing applications. Advances in security already allow academic researchers to use grid technology, but many obstacles remain to grids' commercial use.

"In particle physics, once you've verified that you belong to a certain VO, it doesn't really matter which of the VO's resources you use," adds Gagliardi. "But if you're submitting grid jobs to a business that's charging for resources, there must be a strict accounting of which resources you've used and how much you've used them."

## A myriad of grids

Gutsche is adapting a CMS analysis application that already runs on the LHC Computing Grid (LCG), the infrastructure that supports European LHC physicists, to run on the US-based OSG. Due to differences in middleware, the same application doesn't automatically run on both the LCG and the OSG. This is common in today's grid world, which includes small, large, single-focus, and multi-disciplinary grids.

Right now, a grid expert will spend weeks or months adapting a certain application to interface with one flavor of grid middleware, and then repeat the process to use another grid. Once standards have been adopted by the grid community, making applications work on one or more grids will be much easier.

"The grid community is now big enough and experienced enough that users may need to talk to several grid infrastructures," says Mulmo. "To do that today, they often have to have several middlewares installed. But now we start to see the international communities and collaborations forcing us—the middleware developers—to cooperate on common interfaces and interoperability."

Grid visionaries believe that, eventually, there will be one worldwide "Grid" made up of many smaller grids that operate seamlessly.

"It's always going to be a bit like the Internet," says Foster, "with common protocols, and a lot of common software, but lots of different networks designed to deliver a different quality of service to different communities. Some large communities will want or need their own dedicated infrastructure, but the small group of archaeologists will need to use general-purpose infrastructures."

## First signs of success

Gutsche sits at his computer and prepares to test the data analysis application on the grid. He initializes his grid interface and goes through the steps to identify himself to the OSG and the LCG. The program will extract statistics from a particle physics data set. He selects the data set he wants from a Web site, and types the names of both the data set and the program into a small file. He starts the user analysis submission application, which reads the file. The computer takes it from here. It splits the request into many smaller jobs and submits them to a central computer known as the resource broker.

The resource broker finds the data set, makes sure it is complete, and checks the computer farm where the data is located to verify that there are enough resources to run all of Gutsche's jobs. If so, his requests are scheduled on the farm. If all goes well the jobs will run successfully and Gutsche can collect the results from the resource broker.

Gutsche's first attempt on this day is aborted, and a second attempt meets a similar fate—evidence that the grid is still a job in progress. Eventually, his requests run successfully on an LCG site in Italy and on an OSG site at Fermilab.

"To run on the grid right now you still sometimes need an expert who knows specific information about each grid site," says Gutsche. "By 2007, we hope that all of this will be hidden from the user. We won't know where our job is running, or that two sites are using different middleware. All we'll know is what data we want, and where to find our results."

Despite significant development still to be done, many scientists have already used grids to advance their research. Physicists, biologists, chemists, nanoscientists, and researchers from other fields have used dozens of grid projects to successfully access computing resources greater than those available at their home institutions, and to share and use different types of data from different scientific specialties. Multidisciplinary grids like the OSG, the TeraGrid in the United States, and the Enabling Grids in E-science project in Europe continue to add more applications at the cutting edge of science.

Just before stopping for lunch, Gutsche receives an instant message from a CMS researcher who wants help analyzing data on the grid. This is one of the first requests Gutsche has received, and he knows that it won't be the last one. Grids are catching on, and Gutsche looks forward to helping more and more people get started, watching grids grow into standard tools for great discoveries.

## Grid growth will drive supercomputing capacities

The growth of grid computing could super-size the future of supercomputers, if Horst Simon is on the mark. "There are some who think that supercomputing could be replaced by the grid," says Simon, of the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory. "For example, using many PCs, SETI@home is a clear success, and it can do a significant job.

"But I use the analogy of the electric power grid," Simon continues. "It would be great to have a lot of windmills and solar panels producing electricity across the whole country, but now, we still can't run the power grid without the large power plants. In the same way, supercomputers are truly the power plants to drive the computing that cannot be distributed across the grid.

"So the grid will drive even more demand for supercomputing, which can produce more and more data to effectively utilize the capacity of grid computing. Think of the LHC and the Tevatron as supercomputers—grid computing will support CMS, ATLAS, all the particle physics detectors as the accelerator-supercomputers produce more and more data."

Simon and colleagues Hans Meuer, of the University of Mannheim, Germany; Erich Strohmaier, of NERSC; and Jack Dongarra, of the University of Tennessee-Knoxville, produce the twice-yearly TOP500 list of the world's fastest supercomputers. Their next list will be published at SC|05 in

Seattle, November 12-18, the annual international conference on high performance computing, networking and storage. On the current list, which was released at the ISC2005 conference in Heidelberg, Germany, on June 22, the top spot went to the BlueGene/L System, a joint development of IBM and DOE's National Nuclear Security Administration (NNSA) and installed at DOE's Lawrence Livermore National Laboratory. BlueGene hit 136.8 teraflops, or trillions of calculations, per second, although Simon granted that judging a supercomputer solely on speed is like judging a particle accelerator solely on brightness.

The definition of a supercomputer has evolved since the TOP500 list originated in 1993, due to the competitive performance of massively parallel technologies in the 1990s. Now, supercomputers are defined as the largest systems available at any given time for solving the most important problems in science and engineering, although they are primarily used for science applications. Simon sees a growth spiral for supercomputing-grid synergy. "Supercomputing and the grid are complementary—mutually reinforcing," he says. "The grid provides the right tools in the middleware to get more science out of the data."

Meanwhile, supercomputing will continue its own growth spiral. "Supercomputing follows Moore's Law of exponential development," Simon observes. "The standard laptop of 2005 would have made the first TOP500 list of 1993."

**Mike Perricone**