Eur. Phys. J. C 34, 67–75 (2004)
Digital Object Identifier (DOI) 10.1140/epjc/s2004-01769-5

THE EUROPEAN PHYSICAL JOURNAL C

Challenges of the LHC: the computing challenge

P. Messina

Argonne National Laboratory, MCS Division Bldg. 221, 9700 South cass Avenue, Argonne, IL 60439, USA

Received: 15 December 2003 /
Published online: 4 May 2004 –

1 LHC's computing needs

The LHC will generate unprecedented volumes of data, hence meeting the LHC computing needs will require innovative approaches that involve linking storage and computing resources that are distributed worldwide. The success of this strategy will depend on advancing the state of the art in a number of technologies, primarily in the software realm. This paper deals entirely with the LHC *off-line* computing needs, from raw data to the physics plots (calibration, reconstruction, simulation, analysis).

The current estimates are that the major LHC experiments will store data onto permanent storage at a raw recording rate of 0.1–1 GigaBytes/sec (GB/s). A single copy of the archive is estimated to grow at a rate of 5–8 PetaBytes (PB)/year and at any time 10 PetaBytes of data will reside on disk. (A Petabyte is 10^{15} bytes. In more familiar terms, it takes more than one million CDs to store one Petabyte.) Each of the four LHC experiments will store between 3 and 10 PB on tape. The total data volume will be tens of Petabytes by 2007–8 and an Exabyte (10^{18} bytes) five to seven years later.

The analysis of these data will require tens of thousands of processors (the high-end commodity processors of 2008, not today's), perhaps as many as 100,000 such processors. Thus the sheer scale of the data and the corresponding analysis poses challenges. If one believes the rule of thumb that when something increases by one order of magnitude, it changes in nature, then the LHC computing task must truly require different approaches, since it is several orders of magnitude greater than previous scientific data investigations.

Paul Messina

2 Distributed, "grid computing" approach chosen

Given the very large requirements for LHC data analysis, it was not considered feasible to put all of the resources at CERN. MONARC (MOdels of Networked Analysis at Regional Centers for LHC experiments), a collaborative effort of all four experiments, has developed a strategy to meet the LHC needs that uses computing and storage resources at physics research centers (including laboratories and universities) worldwide to tackle the analysis [1]. The LHC community contains more than 5000 physicists, residing in about 300 institutes in about 50 countries. The MONARC approach was endorsed as the appropriate one after a comprehensive review of the LHC computing needs [2].

Under the MONARC model, while CERN will retain a copy of all the data, it will not have the computing capacity to satisfy the needs of the thousands of physicists who will undertake the analysis of the data. Copies of subsets of the data will be sent to the sites that will provide resources for LHC data analysis. Even if CERN were to have sufficient computing power, distributing the data and computing resources is desirable since it reduces the need for repeated transfer of data from a central site (CERN) to each user site.

LHC computing will be done on resources located at a large number of Regional Computing Centers in many different countries, interconnected by fast networks. In other words, the LHC computing services will be implemented as a geographically distributed Computational Data Grid. The participating sites will have varying levels of resources, organized hierarchically in Tiers. An important benefit of this approach is that it enables physicists all over the world to contribute intellectually, without requiring their physical presence at CERN.

To be more specific, a multi-tier hierarchical model similar to that developed by the MONARC project has been adopted as the key element of the LHC computing model. In this model, for each experiment, raw data storage and reconstruction will be carried out at the Tier0 centre, which will be at CERN. Analysis, data storage, some reconstruction, Monte-Carlo data generation and data distribution will mainly be the task of several Regional Tier1 centers, followed by a number of (national or infra-national) Tier2 centers, by institutional Tier3 centers or workgroup servers, and by end-user workstations (Tier4). The CERN-based Tier0 + Tier1 facility will support all LHC experiments whereas some Tier1 centers may be dedicated to a single experiment.

A rough estimate is that the sum of the resources at centers outside of CERN will be twice the resources at CERN and that the sum of the resources at all the Tier1 centers will be equal to the power of the resources at the Tier0 centre, as will be the sum of the resources at all the Tier2 centers.

It is worth noting that some Tier1 and Tier2 centers may well be part of a larger institutional computing facility that serves other user communities, in addition to physicists engaged in LHC experiments. This aspect of the distributed facilities poses some technical and managerial challenges, as will be described later in this article.

To be usable, this distributed, hierarchical set of computing and data storage resources must have software and policies of operation that provide to the user a fairly uniform interface and tools to facilitate the migration of data and analysis runs from one part of the tree to others.

The LHC Computing Grid project (LCG) led by CERN is developing and deploying the software, methodologies, and policies needed to create and operate this distributed, hierarchical computing environment [3].

3 A few words on the network infrastructure

A very capable network infrastructure will be required to support the anticipated data flows among the elements of the LHC global computing environment. The estimated bandwidth between Tier0 and the Tier1 centers is 1.5 to 3 Gbps for a single experiment. The traffic between other pairs of nodes in the distributed systems will be comparable, with lower numbers for the lower tiers.

Fortunately, such a network infrastructure is emerging and is certain to be available when LHC data analysis begins in earnest. The exponential use of the web by industry and the general population led commercial carriers to install a prodigious amount of optic fiber and related equipment, with far more capacity than the current demand. That excess capacity, coupled with advances in optical network technology (such as dense wave division multiplexing) have resulted in steeply declining network prices. Furthermore, largely due to the adoption of Grids by the global high-energy physics community, transoceanic networks for research are becoming much faster; in 2003 there was at least one transatlantic network running at 2 Gb/s, faster than most networks within continents. The European Union has created the GEANT network which provides a European "backbone" network for research and education with 10 Gb/s bandwidth presently and with firm plans for further upgrades in the near future. GEANT connects individual country high-speed research networks (such as RENATER/France, GRNET/Greece, GARR/Italy, FCCN/Portugal, REDIRIS/Spain, SuperJANET/United Kingdom and ACONET/Austria). In total, GEANT and the networks it connects reach almost four thousand institutes in 33 countries. Other world regions have or are putting in place high-speed research networks (e.g., in the United States the Teragrid network (40 Gb/s backbone, 30 Gb/s to individual sites), the Internet2 and Lightrail networks, the high-speed networks created by the CANARIE organization in Canada, CERNET in China, Academic and Research Network in Indonesia, Japan Gigabit Network and SInet in Japan, KOREN in Korea, Research Networks in Malaysia, PHNET and PREGINET in the Philippines, SingAREN in Singapore, Thailand and APAN, the Asia Pacific Advanced Network. Equally important, high-speed transoceanic links are bridging these networks so that there will soon be a global research network infrastructure fast enough and with sufficient connectivity to support LHC data transfer needs.

4 What is grid computing?

The Computing Grid (usually just called "the Grid") is a powerful concept that provides a unifying principle for many activities in – and infrastructure plans for – computational science and engineering. It is a premier example of applications-driven research and development that are inspired by the confluence of several technological trends: dramatic advances in network transport, storage devices, and computing power. (The Grid is also quite relevant for commercial applications and many are being pursued, but in this article we will limit ourselves to the world of science.)

5 The grid vision

The "Grid concept" is to enable *resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations,* and to do so without requiring central control or omniscience.

A quote from a description of a particular grid project, the Teragrid [http://www.teragrid.org] presents a vision from the perspective of science:

"An exciting prospect for the TeraGrid is that, by integrating *simulation and modeling* capabilities with *collection* and *analysis* of huge scientific databases, it will create a computing environment that unifies the research methodologies of theory, experiment, and simulation."

The name "Grid" or "computing Grid" was chosen based on an analogy with the electrical power grid. Part of the concept was to be able to obtain seemingly unlimited, ubiquitous distributed computing power and access to remote data and to do so transparently, just as one gets electrical power in the office or at home without having to know what generating plant produced the electricity. Of course the computing grid is much more complex because it must provide transparent access to a variety of information technology resources, such as:

- distributed data collections and data bases,
- computers, of many different types,
- instruments with digital output, and
- telecollaboration tools.

Grid Computing has been identified as an important new technology by a remarkable spectrum of scientific and engineering fields as well as by many commercial and industrial enterprises. See for example [4–13].

The widespread adoption of the grid computing paradigm has taken place very rapidly, even faster than was the case for the web. In only a decade since the formulation of the first concepts that led to the Grid [14], there are scores of grid computing projects underway or in the planning stages in dozens of countries and there are even some production grids for both research and commercial applications. What makes grid computing such a compelling concept?

Grid Computing enables or facilitates the conduct of virtual organizations – geographically and institutionally distributed projects – and such organizations have become essential for tackling many projects in commerce and research. With grid computing one can readily bring to bear the most appropriate and effective human, information, and computing resources for tackling highly complex and multidisciplinary projects.

In commerce, grids will facilitate the integration of efforts across large enterprises as well as the contributions of contractors for projects of finite duration. For instance, new services may be provided in health care as well as medical research.

It is becoming apparent that the use of Grids will be an enabler for major advances and new ways of doing science. Grids have the potential to integrate as never before the triad of scientific methods – theory, experiment, and computation – and to do so on a global scale. This integration can be accomplished by providing a unified environment in which one can execute simulations using models based on theory, access relevant experimental data, perhaps obtain instrument data in real time under control of the simulation, and compare the computational and experimental results. Grids also provide a way to greatly increase the number of individuals who analyze observational data, to facilitate telecollaboration, and to provide broader access to unique experimental or computational facilities.

6 A brief history of grids

A brief history of grids may help explain their nature. While distributed computing began several decades ago (and the Grid can be thought of as a form of distributed computing), the essence of the technologies and methodologies that we now refer to as "the Grid" can be traced to the Gigabit Testbed project initiated by Robert Kahn in the late 1980s [15]. The five testbeds in that project (which was funded by DARPA and the US National Science Foundation) dealt with the issues of using high-speed networks to link geographically distant computers, visualization facilities, and data collections. The testbed teams developed hardware, software, and protocols for supporting the very fast networks and interfacing them to computers. In addition, much effort was focused on creating software that would facilitate the *dynamic and simultaneous use of those resources* to support applications such as interactive exploration of multi-sensor data [16], cancer radiation treatment planning, and climate simulation. So we see that from the outset, Grid technologies (often called *metacomputing* in those early days) were driven both by applications and by infrastructure technologies, the latter including fast wide-area networks, large data archives, software and hardware interfaces, and visualization technologies.

The use of grids for high-end scientific computing, while no longer the prevalent use of grid technologies, is by no means dead. A recently formed activity in Europe was formed to do just that. Distributed European Infrastructure for Supercomputing Applications (DEISA) [17] is a consortium of leading national supercomputing centers in Europe aiming to jointly build and operate a distributed terascale supercomputing facility.

By the mid 1990s a confluence of trends and research advances enabled large-scale demonstrations of Grids. The I-Way experiment of 1995 [18] showed that over a dozen systems on multiple wide-area networks could be linked through common software and that many applications could be executed on the ensemble of resources thus created. Soon after, telecollaboration [19] became an additional focus as it was recognized that the Grid would provide good support for many aspects of distributed research collaborations as such approaches become more prevalent. Instruments and sensors were also added to the scope of resources managed by Grid software, thus providing real-time or near real-time access to data from those sources.

In the same time-frame the Web became an everyday tool for many millions of people around the globe. This phenomenon had several effects on the Grid. One was that most people became familiar with accessing remote resources; typically the resources accessed are static documents but some are dynamic. Consequently, the idea of harnessing major remote computational and data resources was no longer quite so foreign. Second, most institutions installed higher speed connections to the internet as demand increased and prices fell. Third, researchers began to put more and more data collections on line and accessible to others, facilitated by the additional trend of rapidly decreasing data storage costs [20].

Projects were formed to conduct research and develop software tools to enable grid computing, notably Condor [21], Globus [22], UNICORE [23], Legion [24], and their products form the majority of the software technology in use today.

By the late 1990s, the confluence of these trends and advances led to the initiation of projects that could only be done on the Grid or ones that reap major benefits from Grid approaches. Among those are the European Data grid [26], and the Digital Sky [26], which led to Virtual Observatory projects such as [9] and the Astrophysical Virtual Laboratory [27]. Work on the software components that implement the Grid concept – by then usually called *middleware* – accelerated as the new applications became operational and revealed shortcomings or missing functionality.

As is often the case in computing trends, technology advances in several fields inspired and enabled grid computing. Dramatic improvements in the cost-performance and reliability of disks have enabled even small research groups to keep many terabytes of data on-line. Sensor technology has advanced as well and scientists are gathering more and more data. A major motivator for the use of Grids is the access they provide to the huge data collections that are being assembled, maintained, and made available electronically by many disciplines. Unlike computing power, such data archives are not so readily replicated at each user site, hence they must be accessed remotely. Furthermore, multidisciplinary investigations often require the simultaneous access of several data collections, each of which is in a different location. Finally, the analysis of the data can require powerful computer systems that are in another location and the visualization of the results of the analysis might require the use of a system at yet another site.

Computer science research projects worldwide are gradually identifying needed functionalities and ways to provide them, as well as creating a body of software and methodologies that include more and more functionality and provide better support for applications. There are also many application-oriented Grid projects, some of which are operational, including some with an international span, that focus on addressing the challenges their application domain poses for the Grid infrastructure.

An indication of the magnitude of the trend towards adopting the Grid as the computing environment for science and engineering is the existence of support for widely used middleware such as Globus [22] by a number of computer hardware and software companies and *commercial* software efforts for systems such as Legion and for supporting Grid applications. Every major computer manufacturer has internal grid projects, some already have commercial offerings, and there are nearly fifty commercial sponsors of the Global Grid Forum, a grid middleware standards body [28]. A few articles and books that provide useful introductions to grids as they are evolving currently can be found in [29–34].

7 Benefits of grid computing

In just a decade, the potential benefits of Grids have become recognized to the extent that some government agencies and commercial companies have adopted them for production use. Grids are seen as a way to greatly increase the number of scientists who will analyze observational data, to federate data bases to enable the study of complex, multidisciplinary issues, to facilitate telecollaboration, and to provide broader access to unique experimental or computational facilities. Many believe that the use of Grids is likely to be an enabler for major advances and new ways of doing science. Certainly Grids will integrate as never before the triad of scientific methods: theory, experiment, and computation.

8 Benefits of grid computing for LHC

As has already been alluded to, by using Grid Computing, as adapted in the MONARC model, should provide a number of benefits, such as:

- empowering more universities and individual scientists to do research on LHC data, and without having to be at CERN,
- sharing LHC computing resources dynamically,
- handling peak loads better,
- providing capacity "on-demand",
- enabling opportunistic use of non-LHC computing resources,
- avoiding duplicating calculations already carried out by others, through the use of Virtual Data (see [10, 35] for a description of Virtual Data).

By using widely deployed grid software as much as possible and by connecting to facilities that serve other technical communities, additional potential benefits might accrue, such as sharing with other communities the effort of maintaining and enhancing the grid middleware, network and grid monitoring tools, and security mechanisms. One is reminded of Metcalfe's Law:

"The usefulness, or utility, of a network equals the square of the number of users"

One wonders whether Metcalfe's Law should be modified to apply to grids, perhaps:

"The usefulness, or utility, of Computational Grids equals the cube of the sum of the number of users, disciplines, and different resources that participate."

The LHC Computing Grid [3] was formed in 2002 to create a new computing environment that will support the LHC computing requirements. The LCG builds upon relevant efforts of other projects, including two pioneering projects led by CERN and funded by the European Union: the European DataGrid [25] and Enabling Grids for E-science and industry in Europe (EGEE) [36].

The EDG project focused on enabling next generation scientific exploration that requires intensive computation and analysis of shared large-scale databases, millions of Gigabytes, across widely distributed scientific communities. It is a three year project that began in 2001. In many ways the EGEE is a natural, larger-scale follow-on to the EDG.

The EGEE project aims to integrate current national, regional and thematic Grid efforts to create a European Grid infrastructure for the support of the European Research Area. This infrastructure will be built on the EU Research Network GEANT and will exploit Grid expertise that has been generated by projects such as the EU DataGrid project, other EU supported Grid projects and national Grid initiatives such as UK e-Science, INFN Grid, Nordugrid and the US Trillium (a cluster of projects). The EGEE project will begin operation in the spring of 2004.

Grid environments are still in the early stages, so perforce the LCG has to adapt and deploy technologies that are still under development. Fortunately, other projects have similar needs and are engaged in developing many of the needed components. See for example [7–11]. Many of these projects involve distributed access and analysis of scientific, medical, or engineering data, and – while not at the same scale as the LHC – require rather similar functionality as the LCG.

9 Challenges

Despite the existence of many Grid projects that support real applications, there is still much to be done. Some of the existing software has adequate functionality but is not yet robust or easy to install. Fundamental issues such as security and fault tolerance require more work. In a number of cases, it is not just the middleware that needs to evolve to provide the required functionality. Operating systems, data archiving systems, and network software need to be enhanced to support co-scheduling, deadline scheduling, global name spaces, and bandwidth reservation, for example. Better interfaces to database systems are also badly needed.

While standardization of Grid middleware will accelerate the rate of progress, the pace of standardization must take into account the limited experience we have with existing approaches and software: better ideas will surely emerge but we need to facilitate the deployment of Grids in order to determine what works well and what needs to be improved. The Global Grid Forum is a community-led standardization effort that is struggling with these issues.

The current shortcomings and difficulties are not unusual in new fields. Given the great strides already taken by early grid projects, the intense interest by applications communities, and the potential benefits of Grid environments, Grid technologies and applications are exciting fields to pursue.

The previous remarks allude to challenges that all grid projects face. The challenges that LHC computing faces can be categorized into three types: technical, research, and managerial challenges.

9.1 Technical challenges

There are many difficult technical challenges due to the scale, heterogeneity, physical distribution, and dynamic variation of the resources and analysis tasks. To get a feeling for the scale, consider the data points in Table 1.

Table 1. Comparison of parameters related to the handling of one Terabyte and one Petabyte of data

	TERABYTE	PETABYTE
RAM time to move	15 minutes	2 months
1 Gb WAN move time	10 hours	14 months
Disk cost	7 disks	6800 disks
	= $ 5000	= $ 7 million
Disk power	100 Watts	100 KW
Disk weight	5.6 kg	33 tons
Disk footprint	Inside machine	60 m^2

Therefore, storing the data is certainly feasible – at a cost – but requires attention to facilities and ways to cope with frequent hardware failures. If one has to keep 10 PB on disk, nearly 70,000 units will be required. If the mean time to failure is 100,000 hours, a disk will fail every hour or two.

As was mentioned previously, obtaining adequate network speed is not expected to be a challenge. There are wide area networks already in operation at tens of Gigabits per second. Cost has become affordable, so by the time LHC is operational it is likely that all major network links will be of the order of 10 Gb/s per experiment or better.

With such large volumes of data and many millions of individual files, ways have to be developed to reduce the difficulty of:

– sending copies of subsets to many sites and keeping track of what site has which files and replica management,
– storing the data in a safe way, and especially,
– finding the desired files for a given analysis in the context of "dauntingly complex metadata."

Hence a comprehensive data management effort is needed to design and develop a consistent and complete mechanism for tools to manage storage access, data transfer, replica management, and file access from jobs.

An area that is even less mature is workflow management, to allow jobs to move across grids, run on various resources, access data, and receive status and output at a user specified location.

Another technical challenge arises from heterogeneity, which makes interoperability much more difficult. The LHC computing environment will have heterogeneity in essentially everything, including policies:

– computing resources,
– storage resources,
– applications,
– network speeds,
– management domains,
– policies, especially security mechanisms and policies.

Because the resources and the users are distributed, a number of technical issues arise, most of which do not yet have robust solutions:

- identifying the best resources available for the task at hand, in real time,
- global access and global management of massive and complex data,
- monitoring, scheduling, and optimization of job execution on a heterogeneous grid of computing facilities and networks,
- end-to-end networking performance.

Furthermore, the resource requirements will be highly dynamic. In the traditional physics data processing model, the tasks can be categorized as follows:

- event simulation,
- detector calibration,
- reconstruction,
- physics analysis.

Resource access patterns are less predictable than for the other three (jobs are initiated from almost any HEP site in the world; large variation in the patterns of data access).

These tasks involve intimate combinations of data and computation, with unpredictable (autonomous) development of both. In other words, the dynamic, sometimes chaotic nature of the computing load is inherent in the LHC computing requirements.

Thus the nature of physics computing raises the need for:

- on demand computing,
- real-time resource identification,
- fault tolerance,
- virtual data support (retrieve instead of recomputing, unless it will cost less to recompute) – Book-keeping of what has been computed and what has not, in a global environment.

The technical challenges sketched out above are challenging and in general production-quality solutions are not available. However, in most cases there are prototypical implementations, and solutions will emerge with sufficient time, level of effort, and careful development guided by experience with early applications.

9.2 Research challenges

Some topics probably require more than refinement and professional implementation. These are research challenges, some examples of which are:

- Integration of workflow and data base access, co-optimized. For example, if a job has been loaded into computer memory for execution but the devices that store the data needed for that job are busy or unavailable, user time and computer resources will be wasted;
- Performing distributed queries on a global scale. This will be necessary since the data will be distributed among various sites;
- Dealing with dynamic variability in authorization of access for a given user, what resources are operational and available to take on the work, data and schema, and performance of the computers, the networks, and the data servers;
- Dealing with chaotic resource demands due to the thousands of physicists who may submit jobs on the distributed resources;
- Generating metadata automatically for discovery, automation of tasks. Without such metadata, it may be hopelessly time consuming to find the desired files or database records;
- Data provenance tracking, so that one can determine exactly what computations were performed to derive the data products.

The distributed nature of the computing environment raises the need for:

- Flexible and extensible user interfaces that hide most of the complexity of the environment;
- Ways to identify the best resources available for the task at hand;
- Global access and global management of massive and complex data collections;
- Monitoring, simulation, scheduling and optimization of job execution on a heterogeneous grid of computing facilities and networks;
- Achieving and monitoring end-to-end networking performance, application integration;
- Technologies and services for security, privacy, accounting.

In short, informatics research advances will be required to devise mechanisms for implementing some of the functionality that the LHC computing community would like to have.

9.3 Managerial challenges

The managerial challenges are perhaps the thorniest because they involve political and cultural considerations that are sometimes in conflict with the concept of a computing facility that encompasses resources from many different institutions and that requires using software designed and produced by others.

One major challenge is how to effect the transition from research prototype software to production software. The transition requires much more money and different types of people.

Much of the software required to create the LHC computing environment is still immature. Although a good bit of it is in use at a number of grid projects worldwide, it is still far from being easy to install, well documented, professionally implemented, robust, reliable, and interoperable with other software components. To make the transition requires spending a great deal more effort/money than was needed to develop the initial version. A rule of thumb is that it takes 10 to 100 times the effort to develop production-quality software as it took to develop the initial prototype. In addition, the people who are needed to develop the production software need to have different skill sets and motivations from the people whose research created the prototype. Until now, much of the maintenance of the software has been carried out by the group that did the research on which it is based. It is admirable that they were willing to do so, but it is not a sustainable arrangement. On the other hand, there has to be an excellent relationship and lines of communication between the groups who develop new software and the groups charged with supporting an operational grid. Often newcomers to the grid world question the validity of the approach taken by a piece of software and may set out to design a new way of providing the functionality, only to run into major difficulties that had been identified by the original creators of the software. One wants to avoid such experiences, since they waste both time and money.

Having the managerial will to use the right types of professionals for each task is not always easy in research institutions that are accustomed to inventing their own solutions for much of their research. In addition, it is easy to convince oneself that one's needs are unique and therefore unique solutions have to be developed. However, the LHC computing needs turn out not to be very different from those of other scientific and engineering disciplines or even those of some commercial grids. Management will have to consider carefully when it is essential to develop LHC/HEP specific solutions versus when "community" or commercial software is used. In general, the latter choice should be taken if at all possible. By using widely deployed and used software, the cost of maintenance will be much lower (others will be working to ensure that the software runs on new versions of operating systems and new hardware) and its interoperability with other components of the environment is much more likely. The standardized software may not be as esthetically pleasing and users may need to learn new user interfaces, but the long term benefits are substantial.

Policy issues raise many impediments to creating the LHC computing environment. Many of the resources that will comprise the LHC computing facilities will be at different institutions, funded by different governments, and often serving the computing needs of other communities in addition to the LHC physicists. Therefore, mechanisms have to be established for sharing resources that are funded in part for other applications. In addition, security policies vary greatly yet LHC users will frequently need to carry out their computing on systems in several different administrative domains. Even policies on the use of disk and archival storage vary widely and those differences can cause tremendous difficulties in running jobs that use the distributed data.

Fortunately, the LHC community does not have to do it all. As has been mentioned, there are other projects and efforts that have identical or similar needs and goals and with which LHC can collaborate, share the cost, and obtain a better end product. One can readily identify projects and initiatives with which collaboration would be mutually beneficial – and in many cases is taking place to some extent – including, to name just a few:

- European Union sponsored projects: EDG, EGEE, GridLab [37], many other grid projects, and the GEANT network;
- UK e-Science: Core Programme, GridPP [38], Astrogrid [39];
- United States projects: NASA's Information Power Grid (IPG) [40], the Extensible TeraGrid Facility [41], Grid Physics Network (GriPhyN), iVDGL, National Virtual Observatory, NSF Middleware Initiative (NMI) [42], and the Cyberinfrastructure initiative;
- Japan's National Grid Research Initiative, Naregi [43].

Many other countries and regions are also implementing grids for science and engineering, the previous list represents a small fraction of the efforts worldwide.

There is also strong commercial interest in grids. Cisco, HP, IBM, Intel, Microsoft, Oracle, SGI, Sun, Qwest and many other major companies are investing in grid computing technologies and services. Sponsors of the Global Grid Forum include nearly 50 companies [44]. Those companies invest in grid computing because they anticipate a large commercial customer base. Indeed, a few industrial end-user companies are already developing grids for supporting their applications, some of which have global span as well.

The commercial interest in grid computing is good news, because their investments will hopefully produce much useful, interoperable and supported grid software. But there is also bad news: commercial interests do not always match science needs. For example, some commercial suppliers believe that businesses want to build grids that operate only within their company. This is in contrast with science grids that usually span administrative domains and thus have to face many issues that intra-company grids avoid. Also, harvesting of idle computer cycles to reduce costs is often cited as the target commercial application, but science grid applications usually involve retrieval and analysis of vast, distributed data collections, not just "cycle sharing." Finally, if the commercial grid software is proprietary, source code is not available, and community standards are not followed, it may not be suitable for use by science grids.

Despite those potentially negative aspects of commercial grid software, the LHC computing community – and science grid projects in general – should see the difficulties as a challenge. They need to find ways to steer commercial investments to address science needs as much as possible. The Global Grid Forum provides a setting for this, since many of the participants in GGF working groups are from industry. Another mechanism is to form joint projects with commercial companies – users as well as producers – especially ones that demonstrate that some business applications require grid functionality similar to science grids. The UK e-Science Programme has been particularly successful at mounting joint projects with industry [45, 46].

10 The OMII concept as one way to address some of these challenges

A concept referred to as the Open Middleware Infrastructure Institute (OMII) has been proposed as one possible mechanism to address some of the challenges associated with creating and maintaining production-quality software. The OMII, if implemented, would be an international organization (an Institute) sponsored by governments and industry, whose mission would be to produce and maintain open source, standard-conforming and interoperable middleware, building on existing efforts.

Its goal would be to ensure that Grid middleware becomes production-quality and acquires sufficient functionality quickly enough to meet the expectations of the emerging grid user communities.

Implementation of the institute would be a distributed/virtual organization. It could have software development/production centers on each continent, for example. Its constituents would include/involve developers, producers, and integrators of production-quality grid middleware.

Software development could be done by university groups, research laboratory groups, and industrial concerns. There are examples of excellent software products from all three types of institutions. Selection of developers/maintainers would be based on a competitive proposal process.

To be more specific, the OMII would:

- Produce open source software that
  a) could be installed by user organizations to provide grid functionality, and
  b) computer and software companies could adopt and give added value by supporting it, porting to new platforms, optimizing performance on particular platforms, etc., such as was done with MPICH for MPI message-passing libraries, for example;
- Maintain and support the software it produces;
- Follow GGF and other relevant standards (e.g., become a member of W3C);
- Take reference implementations developed by others and turn them into production-quality software. Possibly develop early reference implementation of emerging GGF standards;
- Offer a "GGF standard compliance certification" function for producers of software who want to verify that their product complies with one or more GGF standards.

The UK e-Science Core Programme is the first to try to implement the OMII concept. It has funded the UK component of OMII, which will begin operation in early 2004.

11 Summary and conclusions

The grid approach to meeting LHC computing needs will require substantial technical, research, and managerial efforts.

LHC computing requires grid computing yet Grid technologies are not yet mature. There are many open issues to be addressed and missing functionality to be developed and more gaps will emerge as uses of computing grids proliferate. However, there are grounds for optimism that grid computing will evolve to be the highly useful technology that it promises to be. The commercial and research applications that are driving the grid are also providing the intellectual and financial resources that will lead to more and more production applications of grid computing. Another positive sign is the growing interest in the computer science community in research related to grid computing. Unlike traditional scientific computing, creation and use of grids involve a number of mainstream computer science topics and issues, such as database technology, digital libraries, cybersecurity, ontologies, semantic webs, and web services. Therefore, there is reason to believe that the LHC computing challenges will be met successfully over the next few years.

If the LHC computing challenges are met through grid computing, all scientific fields will have gained a flexible, powerful computing environment in which additional resources of all types can be added readily and accessed easily, including new algorithms and software, which are at least as important as the hardware. The interoperability mechanisms that will have been developed will enable these broader benefits.

The grid approach is most likely to be successful for LHC computing if the LHC community recognizes that many of its needs are shared by other sciences and commerce. While LCG may well lead the way – and should influence what is developed – in the long run it will benefit the most if it can adopt widely deployed and maintained grid software and standards. Once again, the physics community will be a key motivator and early adopter of an important new technology, but it must collaborate with other communities to get the best results in the long run.

Acknowledgement. This work was supported in part by the Mathematical, Information, and Computational Science Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy, under Contract W-31-109-ENG-38.

References

1. Models of Networked Analysis at Regional Centres for LHC Experiments.
   http://monarc.web.cern.ch/MONARC/,
   MONARC Phase 2 report CERN/LCB 2000-001, March 2000,
   http://monarc.web.cern.ch/MONARC/docs/phase2report/Phase2Report.pdf
2. Report of the Steering Group of the LHC Computing Review. CERN/LHCC/2001-004, CERN/RRB-D 2001-3, 22 February 2001.
   http://lhc-computing-review-public.web.cern.ch/lhc-computing-review-public/Public/

3. The LHC Computing Grid project, http://lcg.web.cern.ch/LCG/ and http://cern.ch/Hans.Hoffmann/C-RRB-Oct02-Plenary.ppt
4. UK Research Councils E-science Program, http://www.research-councils.ac.uk/escience/
5. European Commission Sixth Framework Research Program, http://europa.eu.int/comm/ research/fp6/index_en.html
6. Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure. January 2003, http://www.cise.nsf.gov/sci/reports/atkins.pdf
7. George E. Brown, Jr. Network for Earthquake Engineering Simulation, http://www.nees.org/
8. The National Ecological Observatory Network (NEON), http://www.nsf.gov/bio/neon/start.htm
9. The National Virtual Observatory, http://www.us-vo.org/
10. The Biomedical Informatics Research Network (BIRN), http://www.nbirn.net/
11. The Grid Physics Network (GriPhyN), http://www.griphyn.org/index.php
12. The Space Physics and Aeronomy Research Collaboratory (SPARC), http://intel.si.umich.edu/sparc/ and http://www.crew.umich.edu/
13. DOE Scientific Discovery Through Advanced Computing (SciDAC), http://www.osti.gov/scidac/
14. P. Messina, CASA Gigabit Network Testbed, in The Concurrent Supercomputing Consortium: Scientific and Engineering Applications, Caltech Concurrent Supercomputing Facilities Technical Report CCSF-1-91, Pasadena, CA, May 1991
15. Gigabit Testbed projects, http://www.cnri.reston.va.us/gigafr/index.html
16. P. Messina, Distributed Supercomputing Applications, in The Grid: Blueprint for a New Computing Infrastructure, I. Foster, C. Kesselman, eds., Morgan Kaufman, chapter 3, 1999, ISBN 1-55860-475-8
17. Distributed European Infrastructure for Supercomputing Applications (DEISA), www.deisa.org
18. I-Way, http://archive.ncsa.uiuc.edu/General/Training/SC95/I-WAY.nextgen.html
19. DOE National Collaboratories Program, http://doecollaboratory.pnl.gov/
20. R. Williams, P. Messina, F. Gagliardi, J. Darlington, G. Aloisio, Report of the European Union – United States joint workshop on Large Scientific Databases, Annapolis, Maryland, USA, 1999 September 8–10, CACR – 179 (October 1999), http://www.cacr.caltech.edu/euus/
21. Condor, http://www.cs.wisc.edu/condor/
22. Globus project, http://www.globus.org
23. UNICORE, http://www.unicore.de
24. Legion, http://legion.virginia.edu/
25. European Data Grid, http://eu-datagrid.web.cern.ch/eu-datagrid/
26. Digital Sky[http://www.npaci.edu/envision/v15.3/digitalsky.html]
27. Astrophysical Virtual Laboratory, http://www.eso.org/projects/avo/
28. Global Grid Forum, http://www.gridforum.org/
29. I. Foster, C. Kesselman (eds.), *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufman, 1999, ISBN 1-55860-475-8
30. I. Foster, C. Kesselman, S. Tuecke, The Anatomy of the Grid: Enabling Scalable Virtual Organizations. Int. J. Supercomput. Appl. **15**(3) (2001)
31. The Grid: A New Infrastructure for 21st Century Science. I. Foster. Physics Today **55**(2), 42-47 (2002)
32. I. Foster, C. Kesselman, J. Nick, S. Tuecke, The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration. Open Grid Service Infrastructure WG, Global Grid Forum, June 22, 2002
33. F. Berman, G. Fox, T. Hey (eds.) *Grid Computing: Making the Global Infrastructure a Reality*. Wiley, 2003, ISBN: 0-470-85319-0
34. I. Foster, C. Kesselman Price, *The Grid 2: Blueprint for a New Computing Infrastructure*, Morgan Kaufman, 2003, ISBN: 1-55860-933-4
35. International Virtual Data laboratory, http://www.ivdgl.org/
36. The EGEE Project, Enagling grids for E-science in Europe, http://public.eu-egee.org/
37. GridLab, http://www.gridlab.org/
38. GridPP, The grid for UK particle physics, http://www.gridpp.ac.uk/
39. Astrogrid, http://www.astrogrid.org/
40. NASA IPG, http://www.ipg.nasa.gov/
41. The TeraGrid Project, www.teragrid.org
42. The NSF Middleware Initiative, http://www.nsf-middleware.org/
43. Japan's National Grid Research Initiative, http://www.naregi.org/index_e.html
44. GGF sponsors, http://www.ggf.org/L_Involved_Sponsors/2003_spons.htm
45. UK e-Science industrial outreach, http://www.gridoutreach.org.uk/
46. UK e-Science industrial projects, http://www.nesc.ac.uk/projects/industrial_current.html