Editor: Geoffrey Fox, gcf@indiana.edu

INTEGRATING COMPUTING AND INFORMATION ON GRIDS

By Geoffrey Fox

IN EARLIER ARTICLES, WE DISCUSSED MANY ASPECTS OF GRIDS. THIS TIME, WE LOOK IN MORE DEPTH AT SOME OF THE DIFFERENT CAPABILITIES GRIDS MUST SUPPORT. HOWEVER, because there are several definitions of Grids, we will adopt two views of them: those that support e-Science (representing increasing global collaborations and shared resources) and those that provide an infrastructure with the ability to dynamically link managed resources (to support the execution of large-scale, resource-intensive, distributed applications).[1]

These descriptions emphasize a high-level user requirement and the system capability, respectively. We are integrating resources, managing them, and using them to support distributed collaborative engineering and science.

Figure 1 shows the rich variety of resources you might wish to use in a typical e-Science scenario: raw data from small and large sensors and satellites, curated processed data, computing resources both as massively parallel machines and as a pool of independent workstations, and, finally, a user interface with visualization and analysis capabilities. Grids must embrace such multiple heterogeneous resources and then integrate them. In this article, we will discuss how this requires the meeting of three worlds: computers, databases, and sensors.

Grid Information Resources

Globus: www.globus.org and Chapter 5 of www.grid2002.org.

Condor: www.cs.wisc.edu/condor/ and Chapter 11 of www.grid2002.org.

Grid Service Interface to Information Grids: www.ogsadai.org and Chapter 14 of www.grid2002.org.

OGSA and GT3: www.globus.org/ogsa and Chapter 8 of www.grid2002.org.

OGSI Open Grid services Infrastructure: www.gridforum.org/ogsi-wg/.

e-Science and Data Deluge: www.escience-grid.org.uk/ and Chapter 36 of www.grid2002.org.

Virtual Observatories: www.us-vo.org, www.astrogrid.org, and Chapter 38 of www.grid2002.org.

Particle Physics Grids: http://eu-datagrid.web.cern.ch/eu-datagrid/, www.griphyn.org, and Chapter 39 of www.grid2002.org.

Bioinformatics Grids: www.mygrid.info, www.discovery-on-the.net/, and Chapter 40 of www.grid2002.org.

Collaborating

It is difficult to create definitive lines between different ways of doing e-Science and different ways of looking at data. Data can be streamed in from sensors, stored in large files, produced as a visualization or science output of a program, or stored in databases. Furthermore, we have a conventional pipeline sensor, data, information, and knowledge, describing increasing refinement as bits are passed through an analysis process. Here, we focus on two extreme models. *Run-a-job* Grids execute linked programs that typically fetch and store data in files and run on multiple remote computers that could be high performance (massively parallel). *Information* Grids let you access a set of databases holding either metadata or raw information.

At their simplest, run-a-job Grids support remote "shells" that give command-line interfaces to executing jobs on remote computers. Sophisticated systems like Globus and Condor allow the jobs and their associated files to be on geographically dispersed machines; furthermore, they support the scheduling of many simultaneous jobs. The challenge of this type of Grid is clearly seen in the initial stages of the analysis of data from particle-physics experiments. The international science collaborations formed for experiments at the CERN Large Hadron Collider accelerator will process the raw data from their detectors on a mammoth worldwide network of computers. This computational network will need to support tens of thousands of simultaneous jobs running reliably 24 hours a day and reading the tens of petabytes of data produced each year.

Typical information Grids handle applications such as virtual observatories and bioinformatics where the typical service provided is accessing a database. In each case, the database holds either metadata or highly processed data from many experiments in astronomy and biology. Bioinformatics has de-

 Copublished by the IEEE CS and the AIP 1521-9615/03/$17.00

veloped this model over many years, with several sites across the world responsible for ensuring the data is properly entered and maintained in the database. This data curation activity seems likely to grow in importance in many fields and become an essential feature and driver of Grids.

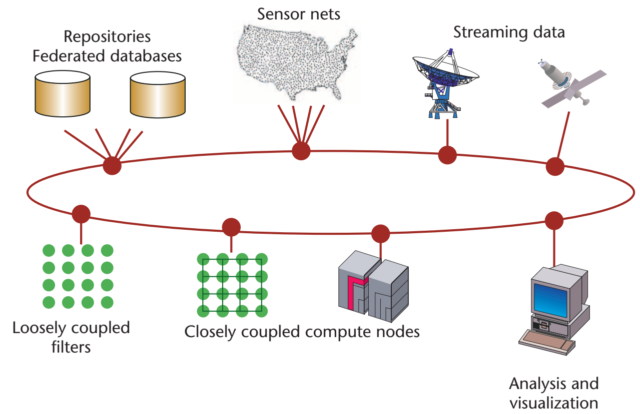

Figure 1. A Grid integrates a wide variety of geographically distributed resources. Various styles of computers are linked together as filters or simulations.

Grid Models

Today, Information Grids start with a service model, where information is extracted by sending messages to a database service. A service model's distributed-object model and XML Schema specification differentiate it from the classic Unix shell run-a-job model with a traditional command-line interface and classic concepts such as Fortran jobs and opaque but efficient data files. Figure 2 depicts a typical Grid multitier service architecture view, with program, service, and client levels overlaid. We usually discuss information Grids at the service level, which has richer functionality but typically lower bandwidth than the program view, where the run-a-job Grids operate.

In earlier articles, we discussed the Open Grid Service Architecture (OGSA) approach that will unify Grids by giving it a service view to all capabilities.[2] Today, most information Grids adopt a pure Web service approach, and most run-a-job Grids use Globus; these will be unified as OGSA is refined and standardized in the Global Grid Forum. Globus Toolkit 3 (GT3) implements these ideas and preliminary implementations already are available. We discussed Grid computing environments in the March/April issue. There, we saw many successful service models built on top of Globus GT2, the original non-Web service version; so we can expect OGSA's more systematic approach to be at least as successful.

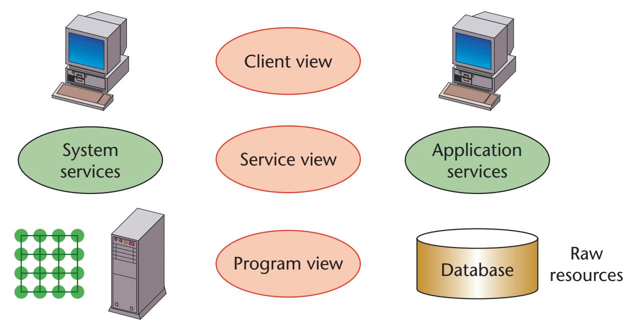

Figure 2. Service architecture Grid program (or resource), service, and client levels (or views).

Building Blocks

A critical development for information Grids is in the OGSA-DAI (Data Access and Integration) software produced by the UK's e-Science program. This project is designing and implementing a Grid service interface for important databases (including both relational and XML). As well as providing a uniform database interface, OGSA-DAI will support distributed queries across multiple databases and the integration of filters that provide customized views of a data source, as shown in Figure 3. As OGSA-DAI matures, expect it to be a critical information Grid building block.

These filtered views would be described initially as a Web service with their input and output interfaces specified by Web Service Definition Language (WSDL). This WSDL specification is easily upgradeable to be Grid service-compatible by adding the extra Open Grid Services Infrastructure (OGSI) features. So, I see a reasonably clear model for building new information Grids with Web service-based filters linked to OGSA-DAI wrapped databases.

Complexity Grids

Finally, we note that many applications need to combine the features of information and run-a-job Grids and, so, we need to discuss "complexity" Grids, so-called because they support the emerging fields of geocomplexity and biocomplexity. Figure 4 shows one way to look at the problem; as generalized data assimilation where the Grid manages distributed data sources that could be databases (biology), sensor nets (environment), or streams from multiple satellites (Earth science).

This data can be preprocessed in a distributed fashion as it is pruned and transformed for use in large-scale simulations. Note that the anticipated "data deluge" is expected to produce so much data that substantial filtering could be needed to project the data onto those components of greatest value to guide the simulation. Actually, you could consider particle physics to have the figure's general structure. The data corresponds to the raw events from the accelerator, the filters correspond to the initial processing to produce data summary tapes, and the data assimilation "simulation" corresponds to the physics analysis phase. In the latter case, you typically need a large par-

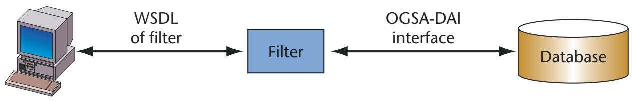

Figure 3. The pipeline of data from a back-end Grid data source. The database interface is defined by OGSA-DAI, and the filter result is defined by its own WSDL specification.

allel machine for partial-differential-equation-dominated fields like climate, ocean, and weather simulations. In the particle physics case, largely uncoupled machines are necessary to support Monte Carlo simulations and data analysis.

Complexity Grids represent one hybrid of run-a-job and information Grids and emphasize the integration of experimental data with simulations and analysis tools. An organization's computing environment forms another such hybrid. Expect a growing interest in enterprise Grids, with, as a special case, campus Grids, supporting university communities.

This article has been necessarily vague as we are only now building the core interfaces and infrastructure. However, we understand enough today that we can start to build these Grids with some confidence that our work will not require drastic changes as the field evolves.

Some might dismiss the Grid as "yet another effort to redo distributed computing" done by a bunch of parallel computing guys who don't even know all the nifty things Corba discovered. We disagree because we have a set of new compelling science and engineering application drivers and a wonderful technology base from the Internet (such as Web services). We will see a lot of important standards work from the Global Grid Forum, some aspects of which we mentioned here. We also will see several Grid technologies and systems with the GT3 reference implementation creating the early excitement. However there will be other commercial and academic offerings that will hopefully interoperate in some fashion. Personally, I am interested in developing Grids that offer key qualities: Autonomic (fault tolerant and self-healing), lightweight and peer-to-peer (to satisfy general users), and with a rich metadata model (the Semantic Grid).

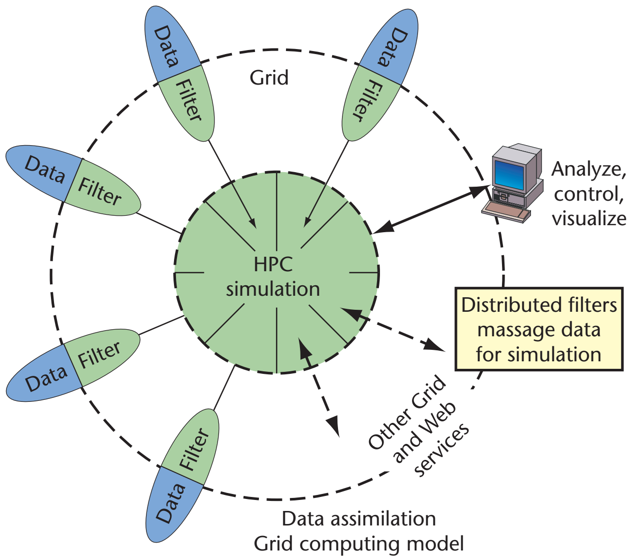

Figure 4. The "queen bee" diagram for data assimilation on the Grid. Other services such as computational steering and visualization are also part of this Grid.

References

1. F. Berman, G. Fox, and T. Hey, eds., *Grid Computing: Making the Global Infrastructure a Reality*, John Wiley & Sons, 2003, www.grid2002.org.
2. G. Fox, "Grid Computing Environments," *Computing in Science & Eng.*, vol. 5, no. 2, 2003, p. 68.

Geoffrey Fox is director of the Community Grids Lab at Indiana University. He has a PhD in theoretical physics from Cambridge University. Contact him at gcf@indiana.edu; wwwcommunitygrids.iu.edu/IC2.html.

Submissions: Send two copies, one word-processed file and one PostScript file, of articles and proposals to Francis Sullivan, Editor in Chief, *CiSE*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; cise@computer.org. Submissions should not exceed 6,000 words and 10 references. All submissions are subject to editing for clarity, style, and space.

Editorial: Unless otherwise stated, bylined articles and departments, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *CiSE* does not necessarily constitute endorsement by the IEEE, the AIP, or the IEEE Computer Society.

Circulation: *Computing in Science & Engineering* (ISSN 1521-9615) is published bimonthly by the AIP and the IEEE Computer Society. IEEE Headquarters, Three Park Ave., 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314, phone +1 714 821 8380; IEEE Computer Society Headquarters, 1730 Massachusetts Ave. NW, Washington, DC 20036-1903; AIP Circulation and Fulfillment Department, 1NO1, 2 Huntington Quadrangle, Melville, NY 11747-4502. Annual subscription rates for 2001: $41 for Computer Society members (print only) and $52 for AIP member society members (print plus online). For more information on other subscription prices, see http://computer.org/subscribe or http://ojps.aip.org/cise/subscrib.html. Back issues cost $10 for members, $20 for non-members. This magazine is available on microfiche.

Postmaster: Send undelivered copies and address changes to Circulation Dept., *Computing in Science & Engineering*, PO Box 3014, Los Alamitos, CA 90720-1314. Periodicals postage paid at New York, NY, and at additional mailing offices. Canadian GST #125634188. Canada Post Publications Mail Agreement Number 0605298. Printed in the USA.