E-Science: The Grid and the Semantic Web

David De Roure, *University of Southampton*
James A. Hendler, *University of Maryland*

Over the past few years, researchers have been treated to two visions of the Internet's future. One is the Semantic Web, the next generation of World Wide Web technology. The second is grid computing, the next generation of internetworked processing promoted by Ian Foster and Carl Kesselman.[1] The Semantic Web is described as "an extension of the current Web in which information is given well-defined meaning, better enabling computers and people to work in cooperation."[2] Grid computing is defined as "flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources."[3] For several years, these visions have developed separately, in distinct communities, with only a few researchers exploring the interaction between the two.[4]

In this article, we discuss their differences and, more importantly, their similarities, and explore the work needed to bring the two together. We focus particularly on scientists' needs, an area in which the high-power computing made possible by grid computing and the large-scale, distributed information management enabled by Semantic Web technologies will need to be integrated. In particular, this will enable new approaches to interdisciplinary scientific endeavors made possible by these new technologies.[5]

E-science needs

"E-science" is the term often applied to the use of advanced computing technologies to support scientists. In the UK, e-Science is a major research program launched in 2001 with a focus on "global collaboration in key areas of science, and the next generation of infrastructure that will enable it."[6]

The UK program focuses primarily on enabling new science, rather than on the infrastructure per se, with computer scientists working closely with application scientists. The Grid was regarded as e-Science's infrastructure, but it was clear that information management—not simply high-performance computing or specialized scientific instrumentation—would be required to meet program goals. This contrasts somewhat with previous US cyber-infrastructure efforts and collaborations promoted by the US National Science Foundation and National Institutes of Health. These concentrated more on high-performance computing and, more recently, large databases for the former and sharing specialized equipment for the latter.

Information management would be required also because the need for interoperability is rarely more evident than when starting a new e-Science project. Researchers involved often wish to assemble the components, services, information, and knowledge from previous and concurrent projects. In fact, interoperability is key to all aspects of scale that characterize e-Science, such as scale of data, computation, and collaboration. For example, to predict the physical properties of a crystal, a chemist may wish to correlate a new molecular structure with existing structural databases. We need interoperable information in order to query across the multiple, diverse data sets, and an interoperable infrastructure to make use of existing services for doing this.

As the Semantic Web deals to a large extent with interoperability,[7] it became apparent that, from the e-Science perspective, both grid and Semantic Web technologies would be necessary. Neither technology on its own would enable us to achieve the full e-Science vision.[8] This integration—a *Semantic Grid*—would serve as an infrastructure for this vision.

This is a compelling story for the grid community, which was increasingly finding a need for metadata as it tackled integration issues. Consequently, the level of abstraction was moving from data to information, and the terms "information grid" and "knowledge grid" began to be used. However, with the exception of specific domains and projects (notably in the life sciences), the broader grid community was largely unfamiliar with W3C efforts above the XML layer. Hence, the Semantic Grid Research Group was created in the Global Grid Forum, its goal "to realise the added value of Semantic Web technologies for Grid users and developers" (www.semanticgrid.org/GGF). Its short-term intention was to bring established RDF technologies to Grid projects while also preparing for the later adoption of other emerging Semantic Web technologies,

1094-7167/04/$20.00 © 2004 IEEE
Published by the IEEE Computer Society

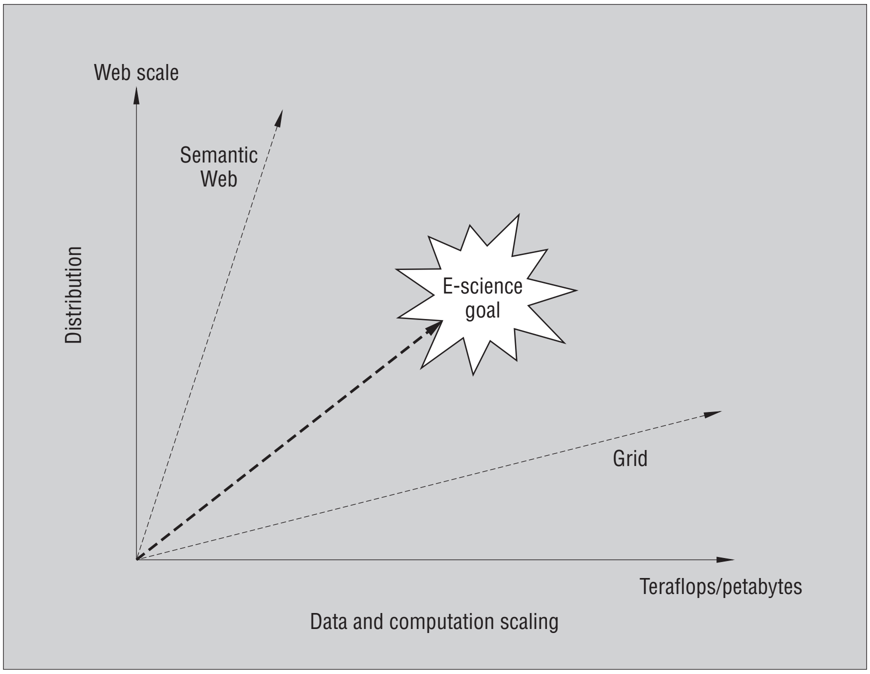

Figure 1. E-science needs aren't totally met by either the Grid or the Semantic Web. Rather, a Semantic Grid vision requires features of both.

particularly integrating Semantic Web and Web Service-based computing.

Semantics on and in the Grid

Some applications the Semantic Web community is exploring are immediately applicable to e-science. For example, ontologies, the backbone of many Semantic Web applications, already exist in certain domains, such as bioinformatics. Consequently, applying the Semantic Web on top of the Grid has made good progress in certain areas, as in the myGrid project, which is producing a virtual laboratory workbench to serve the life sciences community (www.mygrid.man.ac.uk). Projects are successfully using Semantic Web technology to enhance publishing and linking scientific materials.

However, to achieve the truly interoperable infrastructure e-science requires, we must go much further, applying Semantic Web technologies inside grid middleware. This work is made practical by the service-oriented approach known as the Open Grid Services Architecture, which the grid community has adopted to make it easy to create dynamic "virtual organizations" by service integration.[9]

OGSA was originally conceived as a Web Services enhancement to meet grid applications' requirements. Web Services have subsequently evolved to meet similar needs, OGSA implementations based on these new specifications are emerging, and the adoption of Semantic Web techniques within this grid infrastructure is now closely aligned to the growing research in Semantic Web Services.[10] Achieving interoperable infrastructure requires the development of common vocabularies and metadata frameworks as the basis for description, discovery, and integration of the services, together with the use of domain-specific knowledge for problem solving in order to compose services.[11,12]

This ongoing work is strong support for pursuing the application of Semantic Web technologies to e-science. The grid research community is now very much lined up behind a vision of computing that's inherently modular and distributed. And, the e-science community needs the specialized computing environments made possible by grid computing services. The Semantic Web provides a crucial missing piece for linking these services and providing the information management capabilities needed for e-science to flourish. Furthermore, the grid community provides a number of drivers that increase the promise of success:

- A broad range of applications that stand to gain real benefits from Semantic Web technologies, starting with established RDF solutions.
- An enthusiastic and coherent research community that's motivated by practical outcome.
- A community-based standards mechanism (the Global Grid Forum), which can support creating and adopting appropriate schema and ontologies that meet the needs of e-scientists.
- A vibrant developer community that has the potential to adapt to generating and working with metadata and new tools (in contrast to the business community, for example).
- In some cases, the Semantic Web will benefit from using grid infrastructure (for example, large volumes of data and processing shouldn't be an obstacle in this arena).
- The emergence of a service-based architecture, based on OGSA, that can provide an abstraction from the underlying grid implementation for Semantic Web-based applications.
- A growing awareness that, without some kind of computing infrastructure that can help enhance communication across applications and research areas, we can't bridge the disparate communities that must work together to develop modern interdisciplinary applications.

Despite these drivers, experience implementing this vision in several e-Science projects suggests many challenges for the e-science and Semantic Web communities if we are to realize this potentially revolutionary change in science.

The challenges

The Semantic Web vision is a network of information sources with rich metadata. It will provide truly Web-scale integration of many information sources aided by automated algorithms for search and discovery. The Grid vision is much more centered on the network of machines and the protocols for interoperability at the data level. E-science needs fall somewhere in between.

The scaling need not necessarily be to the millions of users and tens of billions of URIs (uniform resource identifiers) that Semantic Web visionaries have their hopes set on. Nor does e-science require the teraflop and petabyte computations that motivate many grid researchers. However, the distribution scale must be well above the number of users that many in the Grid world are exploring, and it must move more data for longer computations than many current Semantic Web efforts target (see Figure 1). The combination of these technologies must therefore overcome some obstacles to allow each community to realize some of the other's benefits.

Here we discuss five of these challenges. First we look at realizing the network effect, which underlies the power of the Semantic Web. We then consider moving beyond the centralized stores that are being built today, toward something necessarily more like the Web. The next challenge, automated service assembly, is currently attracting significant attention in the research community. Collaboration infrastructure underlies the notion of the Grid for supporting virtual communities. And our final challenge is the set of issues that arise when the Grid is embedded in the real world. In essence, these challenges help set a research agenda for those interested in a Semantic Grid that can help revolutionize the practice of science, and particularly the pursuit of interdisciplinary scientific breakthroughs.[5]

Realizing the network effect

Fundamentally, much of the Semantic Web's added value comes from accumulating descriptive information about e-Science's artifacts and resources. As different stages of the scientific process work with the same referents—perhaps a sample for analysis, a piece of equipment, a chemical compound, a person, or a publication—metadata can be recorded in various stores (in databases or on Web sites). This distributed metadata is effectively interlinked by the objects it describes. This, for example, enables us to ask new kinds of questions that draw on the aggregated knowledge. The current World Wide Web realizes its network effect (and thus has been able to grow to an amazing size) by allowing any page to point to any other page (or technically for resources to link to other resources). This allows one user to link to, comment on, or disagree with, any statements made on the Web by another. The scaling of the Semantic Web depends on a similar network effect being reached in "information space"—allowing the sharing and linking of machine readable content, and gaining power by linking to, extending, or even disagreeing with that specified in another Semantic Web document.

To achieve this effect for science, however, we need shared, unique URIs for the objects (real and virtual), and appropriate assertions of relationships—including equivalence—between them. Although this can be a challenge in the general case, it's achievable on the scale of a project or a community of users—for example, the Life Sciences Identifier.[13] In the grid community, the Storage Resource Broker (www.npaci.edu/dice/srb) implements persistent identifiers and a global namespace for a distributed filestore, which is a step in this direction. For Semantic Web technologies to take hold, more scientific communities must recognize the importance of linking their resources and make more of them nameable on the Web.

> Organizational issues define boundaries that will strongly influence the design of metadata stores' distributed architecture, particularly the issues of ownership and management.

Moving beyond centralized stores

We need the machinery that enables us to work with the distributed metadata. Although not the vision of the future, current best practice in the Semantic Web is using a centralized, persistent, scalable triplestore (database of RDF triples) to collect knowledge in one place and work with it there. This provides a single repository that's easy to manage and query. The knowledge can be collected from distributed stores via an RDF bus, perhaps on a publish-and-subscribe basis.

For example, CS AKTive Space[14] was designed in part to illustrate such a network effect. Through harvesting information from multiple Web sites and databases in the UK computer science community into a triplestore, a service is created that can answer questions that couldn't be answered before. The triplestore contains many millions of RDF assertions representing information about people, projects, publications, research topics, and so on that conform to a developed ontology for this domain (http://triplestore.aktors.org). To make this work, the equivalence of multiple references to the same objects (for example, different names for the same researcher) must be established, which proves challenging. Collecting data from diverse and heterogeneous sources is costly, especially during the bootstrap phase of a store. Once the network effect kicks in, the added value might provide incentives to help overcome the bootstrap and maintenance costs. Interestingly, the incentive of visibility in the CS AKTive Space service motivates information providers to take steps to comply with the system. So, the onus of interoperability shifts to the information providers.

If we're to realize Grid computing, we must explore mechanisms through which these large-scale knowledge repositories can communicate and share data. The University of Maryland's MINDSWAP Web site (www.mindswap.org), for example, is a different kind of Semantic Web portal, but it also has information about a number of computer scientists. However, having been developed in the US in a separate effort, it uses different terminology than CS AKTive Space. While some terms are easily mapped from one to another ("post doctoral research assistant" has essentially the same meaning in both ontologies), others are much harder to map. For example, in the US, a full professor, an associate professor, and an assistant professor are all referred to as "professor," whereas in the UK, these three ranks are described respectively as professor, reader or senior lecturer, and lecturer—only someone who has achieved the rank of professor carries that title.

Suppose we want the CS AKTive Space site and the MINDSWAP site to share information. One solution would be to agree on common terms for the mismatching concepts and then augment each system's central triplestore with the other site's information. This solution seems contrary to both Web and grid approaches. A better solution would be to develop a protocols-and-query mechanism so that the stores could dynamically query and display each other's information. Then, allow a third "mapping ontology" to be used to provide the linking, thus eliminating the need for all parties to agree to all mappings. Grid technology has some useful techniques and services for achieving distribution as data-store size grows, such as the Storage Resource Broker middleware for connecting to heterogeneous data resources over a network and for accessing replicated data sets. Perhaps Semantic Web researchers need to look to grid solutions for inspiration.

Although large persistent stores might turn out to scale well for now, the eventual distributed e-science infrastructure can't be achieved with massive triplestores and monolithic ontologies. Rather, distributed RDF servers will work with multiple ontologies in a world that's not centrally managed—this puts the "Web" into "Semantic Web," with the express goal of sharing some of the enablers that are responsible for the Web's growth. This will arise in e-Science as we scale up the collaborations beyond individual projects.

It's also necessary to support the scale and breadth of new scientific endeavors. Instead of one vast centralized triplestore for the planet, we'll have a vast distributed virtual triplestore—the Semantic Web—which will essentially self-organize, as the Web does. The shape and size aren't clear, but the communities need to pay attention to the challenges of the distributed architecture. Indeed, you might say there's not one Grid but multiple minigrids, for reasons that have as much to do with the social dynamics of organizations as with technology.

In science, however, organizational issues define boundaries that will strongly influence the design of the distributed architecture of metadata stores, particularly the issues of ownership and management. The metadata has digital-rights management issues, distinct from those of the data—in the same way that access to a library book has a different set of constraints than accessing the library catalog. In some domains, there's an established policy of open access to information sources; in others, digital-rights management underlies the business models.

The challenge comes as we introduce the Semantic Web inferencing enabled by the aggregation of knowledge. For example, people are often surprised at what they find in query results on the CS AKTive Space triplestore, even though the results were inferred from a combination of publicly available information sources. However, it's exactly such inferencing and added conclusions that motivate e-science. The ability to aggregate data resources and process them in formerly unforeseen ways lies at the heart of e-science's motivation of producing potential scientific breakthroughs through such aggregation. This highlights a distinction between rights to data and rights to value-added services based on that data. So, digital-rights management for knowledge services that aggregate and transform data is an important research issue for e-science success.

Automated assembly

A typical e-Science project today might involve a team of graduate-level researchers from computer science and an application discipline working to deliver e-science applications. A generically usable e-science infrastructure in five years shouldn't demand that degree of skilled and dedicated labor. Instead, we envision applications assembled with ease, meeting the requirements of scientists working in and across their traditional disciplines with minimal computer science support. This requires tools to facilitate creating and configuring the e-science environment. Designing those tools requires a focus on a new user: the e-science application-builder. In the fullness of time, we expect this task to move toward the e-scientists themselves; e-science configured by the participating scientists.

> Beyond composing and analyzing the workflow, a set of capabilities will be needed for execution monitoring in this kind of environment, which is a largely unexplored area.

The ability to describe, discover, and compose services provides the essential plumbing to create e-science applications. Whereas Web Services describe how to connect pipes, the semantic service descriptions being created in e-Science projects tell us about service functionality. So, we now understand how to do the mechanics, but we are only now beginning to appreciate how we can generate the applications that meet e-scientists' needs.[11,12] This is a case study in the general problem facing the communities: How do users engage with the Semantic Web and Semantic Web Services?

Current research in Web Service composition[15] shows a potential starting place. This research has developed a composer that lets a user create a customized workflow by composing Web Services. A key enabling technology is using OWL-S[16] for annotating the inputs and outputs of services against ontological terms. A simple reasoner helps determine which inputs and outputs can be linked directly and which can be linked using a translation from one format or ontology to another.

Using the composer, a user creates a workflow of services in a goal-driven way. The user starts the composition process by selecting one of the services registered to the composer and specifying some input to that process. For example, the user could choose "FIR filter" as a service and provide the input "a sensor service" (meaning the system would be free to choose one) or some specific sensor. Similarly, the user could specify a particular visualizer or analysis device and a specific data set or specify (using an ontology) any data set meeting certain characteristics. The system uses a filtering technique based on the service's nonfunctional attributes—that is, the ontological properties that aren't directly inputs and outputs to the service. In the case of a sensor, these would be features such as sensor location, type, deployment date, and sensitivity. The system is also extensible, which we believe will be an extremely important functionality for scientists' use. Any composition generated by the user and the system can be saved as a service itself, allowing it to be reused at a later time (with a different set of inputs) or used for composition with other services.

We believe such a semantics-based, workflow-composition system that's able to use Grid services would be a significant step toward the Grid's wider deployment and use in e-science. A good place to start can be using the OGSA coupled with Semantic Web mappings. OGSA extends WSDL with specific service description elements and ports that relate to making the services available on the Grid. So, researchers are starting to explore developing a version of DAML-S that can directly reason about OGSA services' properties and that can extend OGSA capabilities into new choreographies managed by a program like the composer. This would let a user generate a plan of how to achieve a goal and produce a workflow. The workflow would go to an analysis system, which would examine it for cost, service quality, efficiency, resource use, and so on.

Beyond composing and analyzing the workflow, a set of capabilities will be needed for execution monitoring in this kind of environment, which is a largely unexplored area. That is, when a workflow is run against a set of scientific data, many problems can arise. These can include computational problems (a needed resource is unavailable because of a crashed server or denied access) or problems relating to the scientific process itself (such as lack of provenance for the data, data range issues, or unexpected data set noise). The service composition and monitoring system could track the data derivation and sourcing to help with problems such as these, as well as more mundane scientific problems such as date of publication or authorship.

Creating these powerful composition and workflow-monitoring tools poses another challenge. Developing "orchestrations" of services is an ongoing business need, and many efforts to develop standards can be used in business-to-business e-commerce and related efforts. The Web Services Business Process Execution Language (www.oasis-open.org/committees/wsbpel) is being proposed as the key enabler, but it's essentially a procedural language—you can express arbitrarily complex compositions using a variety of programming-like constructs. Unfortunately WSBPEL, being a Turing-complete language, is not easily amenable to the analyses just described. Other groups are working on less expressive languages that are directly analyzable, but none has been shown to meet the scientific community's needs.

We believe that a process language similar to that used in OWL-S for composite processes extended with key OGSA features (particularly time and resource use) is necessary for e-science's computational requirements. Neither OWL-S nor OGSA alone contains all the features necessary to have semantics powerful enough for composing provably correct workflows and analyzing their resource requirements. A language that combines the two, however, would be a powerful foundation for creating the needed capabilities, and such a language would be an important enabler for e-science capabilities.

Collaboration tools

With both the Grid and the Semantic Web focusing on machine-processing—e-science's automated back office—are we neglecting users? In fact, we are doing this *for* users, to liberate them from the mechanics of e-science so that they can exercise their scientific expertise to generate new science enabled by the new infrastructure's power. In addition to scale of data and computation, new science involves scaling scientific collaboration—and again demands interoperability to achieve it. The Semantic Web promises us improved ways of helping people collaborate. For example, it helps determine "communities of practice" and provides powerful techniques to search for people with particular expertise or working in particular ways. The Semantic Web also can help communities work with specialized resources such as bibliographic archives and scientific data archives, recommending documents and locating resources for users who might not know the keywords in the jargon of another scientific community.

The full richness of the e-Science vision presents further challenges at the interface between humans and the Grid, and the Semantic Web plays a part in all of them.

The grid community has adopted a set of real-time collaboration tools that can directly support meetings and can be integrated with grid applications. So, e-scientists can work together with the data and tools (for example, enjoying collaborative visualization or working together to steer computations). The Semantic Web offers opportunities for extending these collaboratory tools. Capturing scientific discourse during such meetings could be brought into the e-science infrastructure, too, supported by knowledge technologies. For example, the CoAKTinG project (www.aktors.org/coakting) represents the discourse and argumentation structure of discussions explicitly in a graphical notation; this graph provides a hypertext perspective that directly links to the appropriate artifacts.[17] It also supports instant messaging and visualization of presence to enhance group peripheral awareness, planning and task support, and meeting capture and replay. These tools interoperate through the use of ontologies.

We believe this is an exciting opportunity in which using semantic technologies could positively affect the practice of science. What if we could extend these tools to allow the archiving, search, and retrieval of scientific meetings, the intermediate results (some of which don't pan out), the decision structure of research groups, and so on? This would present an opportunity for scientists not only to access the results of others but also to discover how those results were obtained. They could then avoid problems that might have been resolved in the past or look for opportunities that might not have been explored. In particular, often a line of inquiry that fails to pan out might have missed a crucial factor being explored elsewhere. Making it easier for scientists to identify opportunities for joint progress and collaborate with scientists outside their own research groups and areas is an exciting capability enabled by semantic technologies yet to be explored.

The Grid meets the physical world

The full richness of the e-Science vision presents further challenges at the interface between humans and the Grid, and the Semantic Web plays a part in all of them. We'll illustrate this in three aspects.

The Grid is a digital world, and the Semantic Grid is a richer, more powerful, and more autonomic one. As the digital realm increases in scale, we can no longer attend to each part manually but will increasingly delegate routine configuration, repair, and curation to the processes (or agents) running autonomously in it. It's a place where science occurs in a massively accelerated way, where processes that once took years now take minutes, and where virtual organizations of scientists and agents alike can be created and disbanded dynamically. Through all this, it's a place where intelligent systems make entirely new scientific research possible.

The digital world of the Grid must somehow be manifest in the physical world, and this is surely part of the e-science picture, too. Experimental equipment coupled with the Grid generates massive amounts of data, and the collaboration tools we've described couple in the scientists themselves. The laboratory environment becomes a place where data—and metadata—is captured exactly as it is created.[18] The interface to the Grid isn't a Web browser; rather, it's the scientist's working environment, and the scientist's actions become a source of data and annotation. Pervasive computing will increasingly provide the interface to the Grid, and the application of Semantic Web technologies inside grid middleware will reach out to the devices, too. The issues of services description, discovery, and composition are common to the pervasive world, and, in return, it helps with metadata capture and naming.

Clearly this vision also impacts the culture of scientific research. Scientific and scholarly processes that were once conducted on paper now occur in digital form. This change will challenge established practices—it demands a shift in the culture, and even the business models, of scholarly communication. The traditional model of the academic journal paper is superseded by the vision that scientists (or their agents) looking at a paper in an archive can navigate its rich e-science hypertext, which provides access to the original data and metadata, experimental procedure, and discussion. This is the notion of *publication at source*,[19] giving scientists access to original data and making explicit the provenance of derived results—all hugely important in the scientific context. It's a vision that will be realized through the Semantic Web, bringing together the richly interlinked layers of information flow that occur in e-science, capturing relationships between a rich array of content.

The e-science vision is a future research environment based on virtual organizations of people and agents—highly dynamic and with large-scale computation, data, and collaboration. We've focused on science, but this thinking extends more generally to the e-research environment. This includes both quantitative and qualitative research in the social sciences, as well as arts and humanities. Perhaps more importantly, the vision can help us bridge the gaps between disciplines, allowing exciting new combinations for mixed groups of scientists collaborating on ever-harder problems. The full richness of this vision requires us to bring grid computing and Semantic Web technologies together and presents an exciting research challenge for those of us willing to accept it.

Acknowledgments

This work is supported partly by the UK Engineering and Physical Sciences Research Council under the e-Science Core Program and the Advanced Technologies IRC, the European Commission IST Programme, UK Dept. of Trade and Industry, US Army Research Laboratory, DARPA, Fujitsu Labs of America, College Park, Lockheed Martin Advanced Technology Laboratories, the National Science Foundation, the National Institute of Standards and Technology, and NTT Corp.

David De Roure's and **James A. Hendler's** biographies appear on p. 25.

References

1. I. Foster and C. Kesselman, eds., *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann, 1999.
2. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol 279, no. 5, May 2001, pp. 34–43.
3. I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *Int'l J. Supercomputer Applications*, vol. 15, no. 3, fall 2001, pp. 200–222.
4. C.A. Goble et al., "Enhancing Services and Applications with Knowledge and Semantics," *The Grid: Blueprint for a New Computing Infrastructure*, 2nd ed., I. Foster and C. Kesselman, eds., Morgan Kaufmann, 2003, pp. 431–458.

IEEE Intelligent Systems

Advertiser/Product Index
January/February 2004

	Page No.
IRI 2004	Back Cover
MIT Press	7

NEXT ISSUE
March/April 2004
Data and Information Cleaning and Preprocessing

Advertising Sales Offices

Sandy Brown
10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; phone +1 714 821 8380; fax +1 714 821 4010; sbrown@computer.org.

Advertising Contact: Marian Anderson, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; phone +1 714 821 8380; fax +1 714 821 4010; manderson@computer.org.

For production information, and conference and classified advertising, contact Marian Anderson, *IEEE Intelligent Systems*, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314; phone +1 714 821 8380; fax +1 714 821 4010; manderson@computer.org; www.computer.org.

5. J. Hendler, "Science and the Semantic Web," *Science*, 24 Jan. 2003, pp. 520–521.

6 A.J.G. Hey and A.E. Trefethen, "The UK e-Science Core Program and the Grid," *Future Generation Computer Systems*, vol. 18, no. 8, Oct. 2002, pp. 1017–1031.

7. J. Hendler, T. Berners-Lee, and E. Miller, "Integrating Applications on the Semantic Web," *J. Inst. Electrical Engineers of Japan,* vol. 122, no. 10, Oct. 2002, pp. 676–680.

8. D. De Roure, N. Jennings, and N. Shadbolt, *Research Agenda for the Semantic Grid: A Future e-Science Infrastructure*, tech. report UKeS-2002-02, UK e-Science Technical Report Series, National e-Science Centre, Dec. 2001, www.semanticgrid.org/documents.

9. I. Foster, C. Kesselman, and S. Tuecke, "The Physiology of the Grid" *Grid Computing: Making the Global Infrastructure a Reality,*" F. Berman and A.J.G. Hey, eds., John Wiley & Sons, 2003.

10. S.A. McIlraith, T.C. Son, and H. Zeng, "Semantic Web Services," *IEEE Intelligent Systems*, vol. 16, no. 2, Mar./Apr. 2001, pp. 46–53.

11. C. Wroe et al., "A Suite of DAML+OIL Ontologies to Describe Bioinformatics Web Services and Data," *Int'l J. Cooperative Information Systems*, vol. 12, no. 2, Mar. 2003, pp. 197–224.

12. L. Chen et al., "Towards a Knowledge-Based Approach to Semantic Service Composition," *Proc. 2nd Int'l Web Conf.* (ISWC 2003), LNCS 2870, Springer-Verlag, 2003, pp. 319–334.

13. P. Werner et al., "URN Namespace for Life Science Identifiers," I3C, Mar. 2003; www.i3c.org/wgr/ta/resources/lsid/docs/LSIDSyntax9-20-02.htm.

14. N. Shadbolt et al., "CS AKTive Space: Representing Computer Science in the Semantic Web," submitted to *Proc. 2004 World Wide Web Conf.*, ACM Press, 2004; http://eprints.ecs.soton.ac.uk/archive/00008638.

15. R. Masuoka et al., "Ontology-Enabled Pervasive Computing Applications," *IEEE Intelligent Systems*, vol. 18, no. 5, Sept./Oct. 2003, pp. 68–72.

16. OWL Services Coalition, *OWL-S: Semantic Markup for Web Services*, white paper, Nov. 2003; www.daml.org/services/owl-s/1.0.

17. S. Buckingham Shum et al., "CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid," *Proc. 2nd Workshop Advanced Collaborative Environments*, Argonne Nat'l Laboratory, 2002; www.unix.mcs.anl.gov/fl/events/wace2002.

18. m.c. schraefel et al., "Breaking the Book: Translating the Chemistry Lab Book into a Pervasive Computing Lab Environment," to be published in *Proc. ACM SIGCHI Conf. Human Factors in Computing Systems* (CHI 2004), ACM Press, 2004.

19. J.G. Frey, D. De Roure, and L.A. Carr, "Publication at Source: Scientific Communication from a Publication Web to a Data Grid," *Proc. EuroWeb 2002*, British Computer Society, 2002; http://ewics.bcs.org/conferences/2002/euroweb.

PURPOSE The IEEE Computer Society is the world's largest association of computing professionals, and is the leading provider of technical information in the field.

MEMBERSHIP Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEB SITE
The IEEE Computer Society's Web site, at **www.computer.org**, offers information and samples from the society's publications and conferences, as well as a broad range of information about technical committees, standards, student activities, and more.

BOARD OF GOVERNORS
Term Expiring 2004: *Jean M. Bacon, Ricardo Baeza-Yates, Deborah M. Cooper, George V. Cybenko, Haruhisha Ichikawa, Thomas W. Williams, Yervant Zorian*
Term Expiring 2005: *Oscar N. Garcia, Mark A. Grant, Michel Israel, Stephen B. Seidman, Kathleen M. Swigger, Makoto Takizawa, Michael R. Williams*
Term Expiring 2006: *Mark Christensen, Alan Clements, Annie Combelles, Ann Gates, Susan Mengel, James W. Moore, Bill Schilit*
Next Board Meeting: *28 Feb. 2004, Savannah, Ga.*

IEEE OFFICERS
President: ARTHUR W. WINSTON
President-Elect: W. CLEON ANDERSON
Past President: MICHAEL S. ADLER
Executive Director: DANIEL J. SENESE
Secretary: MOHAMED EL-HAWARY
Treasurer: PEDRO A. RAY
VP, Educational Activities: JAMES M. TIEN
VP, Pub. Services & Products: MICHAEL R. LIGHTNER
VP, Regional Activities: MARC T. APTER
VP, Standards Association: JAMES T. CARLO
VP, Technical Activities: RALPH W. WYNDRUM JR.
IEEE Division V Director: GENE H. HOFFNAGLE
IEEE Division VIII Director: JAMES D. ISAAK
President, IEEE-USA: JOHN W. STEADMAN

IEEE COMPUTER SOCIETY

COMPUTER SOCIETY OFFICES
Headquarters Office
1730 Massachusetts Ave. NW
Washington, DC 20036-1992
Phone: +1 202 371 0101
Fax: +1 202 728 9614
E-mail: hq.ofc@computer.org

Publications Office
10662 Los Vaqueros Cir., PO Box 3014
Los Alamitos, CA 90720-1314
Phone: +1 714 821 8380
E-mail: help@computer.org
Membership and Publication Orders:
Phone: +1 800 272 6657
Fax: +1 714 821 4641
E-mail: help@computer.org

Asia/Pacific Office
Watanabe Building
1-4-2 Minami-Aoyama, Minato-ku
Tokyo 107-0062, Japan
Phone: +81 3 3408 3118
Fax: +81 3 3408 3553
E-mail: tokyo.ofc@computer.org

IEEE

EXECUTIVE COMMITTEE
President:
CARL K. CHANG*
Computer Science Dept.
Iowa State University
Ames, IA 50011-1040
Phone: +1 515 294 4377
Fax: +1 515 294 0258
c.chang@computer.org
President-Elect: GERALD L. ENGEL*
Past President: STEPHEN L. DIAMOND*
VP, Educational Activities: MURALI VARANASI*
VP, Electronic Products and Services:
LOWELL G. JOHNSON (1ST VP)*
VP, Conferences and Tutorials:
CHRISTINA SCHOBER*
VP, Chapters Activities:
RICHARD A. KEMMERER (2ND VP)†
VP, Publications: MICHAEL R. WILLIAMS†
VP, Standards Activities: JAMES W. MOORE†
VP, Technical Activities: YERVANT ZORIAN†
Secretary: OSCAR N. GARCIA*
Treasurer: RANGACHAR KASTURI†
2003–2004 IEEE Division V Director:
GENE H. HOFFNAGLE†
2003–2004 IEEE Division VIII Director:
JAMES D. ISAAK†
2004 IEEE Division VIII Director-Elect:
STEPHEN L. DIAMOND*
Computer Editor in Chief: DORIS L. CARVER†
Executive Director: DAVID W. HENNAGE†
* voting member of the Board of Governors
† nonvoting member of the Board of Governors

EXECUTIVE STAFF
Executive Director: DAVID W. HENNAGE
Assoc. Executive Director: ANNE MARIE KELLY
Publisher: ANGELA BURGESS
Assistant Publisher: DICK PRICE
Director, Finance & Administration:
VIOLET S. DOAN
Director, Information Technology & Services:
ROBERT CARE
Manager, Research & Planning: JOHN C. KEATON