Prototyping the Workspaces of the Future

The Access Grid creates collaborative scientific workspaces that challenge traditional desktop metaphors by integrating large-scale visualization displays and lab instruments.

Rick Stevens, Michael E. Papka, and Terry Disz
Argonne National Laboratory and University of Chicago

Collaborative, immersive virtual reality technology has been in use since the late 1980s. In the mid-1990s, people used these systems to investigate multiuser wide-area collaboration scenarios. While such efforts were pioneering in many respects, they proved less than suitable as everyday work environments. Researchers found that people tire easily when spending extended time in the dark spaces required for projector-based virtual reality or when immersed in synthetic worlds without access to high-resolution text displays or high-quality interaction devices.

Multimedia technology capabilities have grown considerably during the past decade, and all modern desktop systems can now handle video and audio with ease. Indeed, desktop PCs have exceeded many dedicated graphics systems in rendering performance and pixel bandwidth. Moreover, with near-ubiquitous high-speed multipoint networks and protocols now available on the Internet, new models for communication are emerging.

The Access Grid project (www.mcs.anl.gov/fl/research/ag/ and www.accessgrid.org) builds on and extends the use of these technologies — collaborative virtual reality, desktop multimedia, point-to-point remote graphics — in ways that are better suited for users in the 21st century.[1] The Access Grid is a first step toward room-based computing environments,[2] which we believe will challenge desktop metaphors, user environments, and perhaps even computer deployment in the decades ahead. Rather than simple, single-stream videoconferencing implemented with special-purpose desktop appliances, we envision entire rooms or laboratories instrumented for rich, full-time, multimodal communications between groups connected over inherently multipoint, high-speed networks. At Argonne Nation-

 1089-7801/03/$17.00©2003 IEEE Published by the IEEE Computer Society

Figure 1. Access Grid node in use. This group at Argonne National Laboratory used the Access Grid to participate with other groups in the US National Science Foundation's announcement of the Extended TeraGrid Facility.

al Laboratory, we are working with three scientific communities to turn this vision into reality.

The Access Grid's Evolution

The Access Grid is one of an important class of emerging grid applications designed to support wide-area, real-time, computer-mediated communications. The project grew out of a long-standing set of research activities at Argonne's Futures Laboratory.[3] In 1994, the Futures Lab began developing new types of collaboration environments based on aggressive assumptions about the future of networking and computing. As an outgrowth of this work, the Access Grid includes influences from the LabSpace project,[4] MOOs,[5] Jupiter,[6] ManyWorlds,[7] and CAVERNsoft projects.[6,8] Like these earlier efforts, the Access Grid relies on a strong spatial metaphor for resource organization — similar in spirit to that used in some text-based virtual reality environments.

Unlike its predecessors, however, the Access Grid is designed to exploit high-performance peer-to-peer (P2P) multimedia services (audio, video, and text) in creating multiperson, shared virtual workspaces. The goal is to develop computer-augmented environments that can support natural audio and video communication for distributed workers.

Beyond the Traditional Approach

While several research groups have focused on the concept of connecting users via desktop video and audio services, the Access Grid project takes two important steps away from this tradition. The first is to create environments for small groups rather than individual users. As Figure 1 illustrates, this means focusing on room-scale systems with large display surfaces and a dedicated computing, display, and multimedia infrastructure.

Large-scale grid-oriented scientific collaborations often involve dozens of institutions and hundreds of researchers. In part, the Grid's development was a response to the need to support highly distributed scientific projects. Many of these large-scale collaborations have emerged as virtual organizations — including US National Science Foundation (NSF) centers, NSF information technology research sites (ITRs), US National Institute of Health (NIH) research resource centers, and US Department of Education (DOE) "Scientific Discovery through Advanced Computing" (SciDAC) centers. Each of these organizations has a persistent management structure, a significant number of shared tasks undertaken by the group, and significant resources dedicated to achieving shared goals.

Within most large projects such as these, however, we can identify core institutions and core teams of researchers. The Access Grid's initial design point was thus to target groups of approximately six to eight users per site at about eight sites — the ideal grouping for the system. By focusing on small-group-oriented technology, the Access Grid project has been able to specify a high-end resource model with more capabilities — more shared bandwidth, higher-end audio equipment, dedicated computing systems, and so on — than systems designed for desktop deployment.

The second step was to build into the Access Grid framework the concept of *persistent virtual venues*. Such venues, or network-based virtual meeting points, provide an organizing framework to control the scope of access and resources available for collaborative sessions. The Web-based AG 1.0 virtual venue software manages multicast address allocation and media tool startup. Node operators locate virtual venues by name by following links from a given virtual venue Web site. For media tools, the Access Grid uses a version of the Video Conference Tool (VIC)[9] that is modified for more efficient rendering, the Robust Audio Tool (RAT) version of Visual Audio Tool (VAT; ftp://ee.lbl.gov/conferencing/vat/), and a text chat tool.[10]

Resource Assumptions

We based our efforts with the Access Grid on two important assumptions about resources in deployment scenarios:

- Bandwidth will soon be at least several orders of magnitude greater than it is today.
- Computing cycles will be inexpensive enough to embed into the infrastructure of meeting rooms.

These led to the belief that the Access Grid could support many parallel video and audio stream sources from each site; that we could make these streams available to all other sites; and that we could easily make dedicated computing resources available at each site for audio processing, video capture, screen control, and collaborative services. This combination of features provides a rich set of communications channels.

Our resource assumptions have affected not only the initial Access Grid design but also the philosophical drivers behind our development efforts. We have built the Access Grid around the idea of a multiuser, semipublic workspace that supports audio and video connections to other such spaces in which users can move freely, carry on local conversations and interactions, and enter and exit rooms in an ad hoc manner. The use of large display surfaces and hands-free full-duplex audio makes interacting with remote participants similar to interacting with those who are physically collocated. Especially important is the ability to conduct multiple conversations simultaneously. High-quality video and audio (including multiple views of video) are important to achieving this sense of presence in Access Grid sessions.

Researchers use more than 200 Access Grid nodes, deployed at private companies and research institutions worldwide, almost daily for distributed meetings, seminars, classes, and a host of other collaborative activities. As Figure 2 illustrates, more than 160 of these nodes have registered at the Argonne Access Grid Web site since June 1999 (www.accessgrid.org). Nodes are typically connected via the high-speed services of Internet2, ESnet, or other regional, national, or international research and academic networks. Figure 3 shows the number of meetings scheduled on the Access Grid per month since September 1999. The spike in October 2001 might reflect the large number of sites preparing for the SCGlobal event at the Supercomputing 2001 conference; the sharp rise in early 2002 could be attributed to world events from the previous fall.

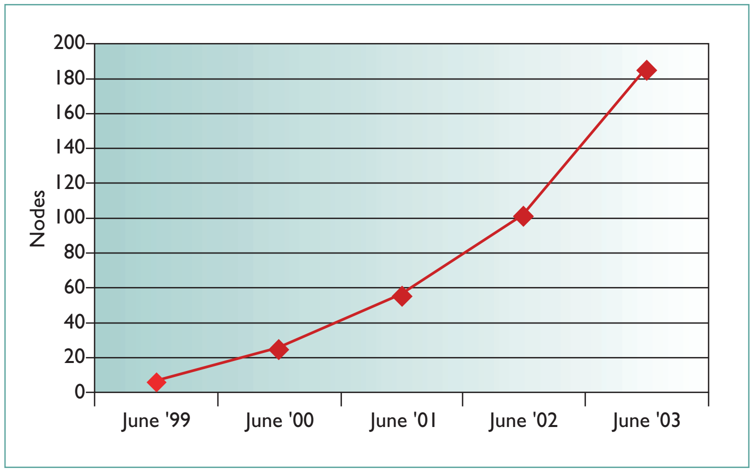

Figure 2. Registered Access Grid nodes. Since June 1999, more than 160 nodes at research institutions and companies worldwide have registered with the Argonne Access Grid Web site.

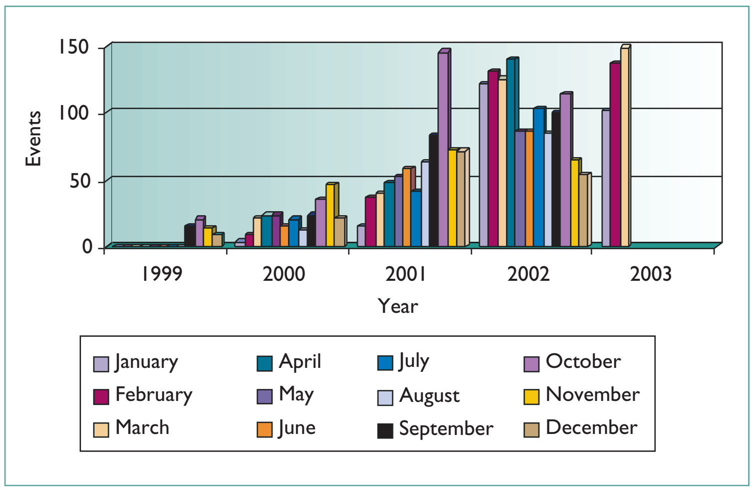

Figure 3. Scheduled events. This graph shows the number of meetings scheduled through the Access Grid scheduler since late 1999. It does not include the many ad hoc meetings in the more than 160 Access Grid meeting rooms.

Design Requirements

In developing the Access Grid, we have identified several design requirements. These requirements are based on our experiences with conventional collaboration systems such as MOOs, desktop videoconferencing, text messaging, and teleconferencing. They incorporate what we feel are the best aspects of these and, in some cases, are designed to avoid the worst aspects.

- *Presence*. One goal of any advanced collaboration system is to reproduce as many physical presence cues as possible to recreate the feeling of copresence among attendees, regardless of their actual location. To provide natural-style audio communications, the Access Grid specifies the use of multiple, always-on, full-duplex microphones, placed to provide maximum coverage within physical workspaces. Access Grid video presents nearly life-scale images, which further increases the feel of "being there."

- *High-quality audio*. The Access Grid implements high-quality audio using professional-level microphones, 16-KHz compressor–decompressors (rather than standard 8-KHz codecs), and commercial echo-cancel products that suppress unwanted noise and prevent audio feedback and echoing.
- *Context and continuity*. With only a single video stream, a typical videoconference can appear to each user as a shifting landscape of different views with little or no context. The Access Grid solves this problem by providing multiple camera views — three cameras on the participants deliver multiple overlapping perspectives of the audience, and one camera pointed at the display provides important contextual cues to remote sites about what the audience is seeing. These streams are delivered simultaneously to all sites, providing everyone with enough context to understand and participate in the meeting.
- *Consistent meeting-place design*. The Access Grid describes a set of best practices for room layout and architecture, including seating arrangements and where to place lighting, cameras, and microphones (see www.accessgrid.org). At Argonne, "life scale" means our video images are a little over 4 feet high, which, at the 4:3 horizontal:vertical aspect ratio of each of our three projectors, requires a display wall a little over 16 feet wide to accommodate the use of three projectors. These figures dictate everything else about the space, including how far away people can sit, how wide the viewing area can be, how many people the room can accommodate, how many microphones are necessary, and so on.
- *Internet connection*. To support these requirements, the Access Grid software infrastructure must provide multiple video and audio streams and manage their distribution. The H.261 video codec we currently use requires a robust Internet connection of at least 20 Mbits/sec for an average AG session. As new codecs are added, we see this requirement growing correspondingly.
- *Multicast infrastructure stability*. The AG video and audio streams are delivered using the IP multicast protocol. This creates a second-order requirement for a set of tools to manage multicast. The US National Laboratory for Applied Network Research has developed tools for debugging and measuring multicast performance. These tools center on a multicast beacon that periodically transmits and listens for signals from all other beacons on a given multicast address. They include a tool for analyzing and visualizing the results the deployed beacons deliver.

The fact that meeting participants want to collaborate from different geographic locations drives a second set of requirements.

- *Meeting management*. All meetings face logistical and management problems, but Access Grid meetings exacerbate this fact because participants are not all collocated. Among the technologies the Access Grid offers for managing remote meetings, a text chat application replaces private face-to-face meetings and whispering. An out-of-band audio back channel — outside normal communications — is useful when real-time command and control is essential to the meeting. Lastly, a variety of meeting schedulers is available for automatically scheduling spaces and document repositories and for notifying participants.
- *Navigation and interaction scope*. Real-life meetings are held in rooms that control the scope of interaction; people can find the rooms by name or number. The Access Grid provides a similar model with virtual venues, which users can locate by name rather than network address. Authorization mechanisms guard each venue against unwanted attendees.
- *Grid integration*. Targeting scientific communities to use fully collaborative meeting places requires that those places support scientific applications that increasingly require access to grid resources. Because it is based on a services model, the Access Grid automatically supports full integration of grid tools and resources within its virtual venues, making them available to authorized users who visit the venue.

With the exception of Grid integration, AG version 1.0 (first made available in early 1999) satisfied most of these requirements. The redesign of the AG software aims to make it more open and easier to adapt to the requirements of an increasing user base.

The Next Step

The ultimate driver for the Access Grid architecture (see Figure 4) is the need for workspaces that provide features and interactions are unavailable in other systems.

The next-generation AG 2.0 software, released in June 2003, promises to greatly broaden the audience of potential developers and service providers. The AG 1.0 virtual venue Web server provided minimal security, mapped multicast addresses onto virtual rooms, and launched media clients, but it had no mechanism for easily adding functionality. AG 2.0 is based on a secure Web services model that will give users and developers opportunities to create new services and add them to a virtual venue for the benefit of all.

We have identified the following elements as high priorities for development in the new release:

- *Scalable virtual venue service*. AG 1.0 is not scalable, and it provides persistence only as simple presence. AG 2.0 will implement a P2P venue service that operates much as the Web does: anyone can host a server (a virtual space), and anyone on the network can visit. The new software will scale to thousands of nodes, with no centralized services, and virtually anyone will be able to create new spaces trivially and link them into the P2P infrastructure.
- *Access Grid security*. In the past, AG users have generally taken a casual approach to security and privacy during sessions. Many applications need real security, however, and require robust access controls and privacy of shared media and applications. The AG 2.0 model leverages the large body of existing work in general Internet security and Grid security, in particular.[11] This model uses public-key infrastructure standards for identity certificates and Globus Toolkit extensions for proxy certificates to implement a secure authentication framework with support for single sign-on. The Globus Toolkit provides transport security as well, which in turn uses the well-understood secure sockets layer (SSL) protocol for point-to-point secure connections. We use these mechanisms to distribute keying information to the end points of the multicast-based media tools to support privacy for the multipoint media sessions.
- *Application sharing and dynamic workspace docking*. *Workspace docking* is analogous to docking a laptop into a network to gain access to local services. Our plans for workspace docking will let users construct ad hoc *space grids* where they can "share" a portion of their personal workspace (desktop applications and data) with other AG users, nodes, or sites, whether local or remote. The docking infrastructure involves migrating or launching one or more specific application clients (linked with multicast as needed) to the AG displays and attaching them to the user's server. This functionality should become available in the AG 2.x release, in late 2003.
- *Node management and user interfaces*. For AG 2.0, we have developed a software layer that will improve node operations through simplified user interfaces, automated node configuration, and node management functions. A key design goal is to enable teams to quickly integrate new types of displays, instruments, and specialized computing devices into a node's management domain and to enable shared applications to exploit these new capabilities.
- *Asynchronous collaboration capabilities*. Work is under way to extend the Voyager system[12] — the multistream, multimedia network-based recording and playback engine developed at Argonne over the past eight years — to include streaming data types, which are required for capturing the interactions and events that occur in the Access Grid's persistent spaces. These include streams of control information used for distributed slide shows or Web browsing; high-resolution lossless encodings of experimental data or simulation output; and streams of navigation information from distributed exploration of large data sets. This capability will be available in later AG 2.x releases next year.
- *Network services*. Later AG 2.x releases will also include an automatic network services module. Network services are resources[13] — video and audio transcoders or speech-to-text converters, for example — that AG sessions can access from

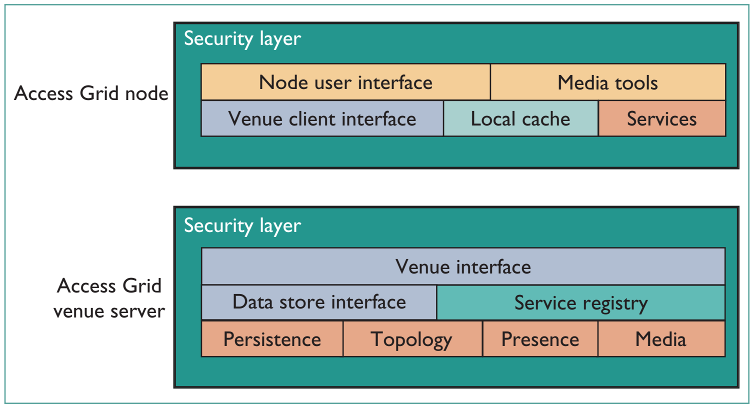

Figure 4. Access Grid services-model architecture. The layered architecture is divided between nodes and venues. All communications and storage operations are conducted through secure interfaces.

Figure 5. Personal Access Grid node. Users can interact with other researchers on things like climate-model visualizations via personal interfaces to the grid.

the network. The AG network services architecture concentrates the required functionality into a small number of key components: the *network services engine* will act as the principal contact point between users; the virtual venue will enable the brokering of access to the services; and the *resolution engines* will provide detailed support for determining the appropriate instances of services to be supplied.

The AG 2.0 release dramatically changes the infrastructure while maintaining the core functionality that has made the Access Grid so popular. Future releases will capitalize on this new infrastructure by providing new services and integrating application-specific requirements.

Collaboration Scenarios and Challenges

We have been working with two scientific communities during the past year as part of the US National Computational Science Alliance (NCSA) Expedition projects. An *expedition* is a narrowly focused project that brings together technology and applications partners to create a collaboration model for scientific problem solving. The Scientific Workspaces of the Future (SWoF) expedition is led by Argonne and includes participants from the University of Illinois at Chicago, the National Center for Supercomputing Applications, Boston University, Brown University, Los Alamos National Laboratory, and Oak Ridge National Laboratory. The SWoF expedition's goal is to develop virtual communities for the fields of atmospheric sciences and molecular biology using the Access Grid infrastructure. We plan to open communications with the high-energy physics community later this year.

Atmospheric Modeling and Simulation Community

One of the atmospheric sciences community's major goals is to deploy a common collaboration and visualization infrastructure throughout the member sites of the University Corporation for Atmospheric Research (UCAR) within the next three years. UCAR members include all US schools that grant doctorates in atmospheric, oceanic, and related sciences. These institutions are very interested in integrating their human and experimental resources, but no program is yet focused or funded to accomplish this task. The proposed NCSA SWoF expedition will directly address this important need.

By developing a standard set of tools that can be deployed inexpensively to all participating institutions, this expedition promises to help accelerate the transformation of an entire discipline. The challenge is to develop a standard way for multiple institutions to actively participate in formulating and analyzing large-scale computational models of the ocean and atmosphere (see Figure 5). These models produce far more output than most UCAR institutions have the ability to generate visualizations for, but much of the national expertise for such analysis is at these institutions. Therefore, remote high-performance visualization is critical. Moreover, graduate departments in atmospheric and oceanic studies tend to be small, and remote courses and seminars from other similar departments could greatly enhance their programs.

Computational Molecular Biology Community

Another SWoF expedition involves members of the computational biology community who have started numerous multi-institutional activities in the past several years. These include the first distributed national seminar and tutorials on computational biology and bioinformatics, which were done via the Access Grid, and a major effort to develop a national computational biology curriculum for delivery via the Access Grid. The SWoF Computational Biology expedition focuses on developing the software infrastructure necessary for building collaboratories to link the biomedical research community. The project goal is to deploy

Alliance SWoF expedition technology to all major US-based systems biology research centers (about two dozen sites) within the next three years, to help create a National Systems Biology Institute modeled after the NASA-supported Distributed National Astrobiology Institute. The US National Institute of Health and NSF are jointly sponsoring a workshop in late 2003 on how best to build such an institute; the Howard Hughes Foundation is also looking into this issue.

While the Access Grid has already gained some acceptance in the computational molecular biology community, the challenge is to provide shared access to computational tools and analysis systems for genomics and molecular modeling and visualization. Of particular importance are tools for visualization of cellular networks and gene-expression patterns. These tools require large amounts of screen real estate (tiled displays) for the large-scale networks present in even the smallest of cells, and they can be greatly enhanced by collaborative interfaces (Access Grid-based virtual venue services). Moreover, it is critical that the visualization and collaborative tool interfaces be based on the emerging open grid services models. This will enable them to scale across the TeraGrid and other discipline grids that are being constructed, but it will also let the tools leverage grid tools that are being developed for data and computation.

High-Energy Physics Community

Researchers in the high-energy physics community are working on a variety of theoretical, simulation, and experimental efforts. Typical experiments can include more than a hundred collaborators, representing tens of different organizations around the world. Such collaborations currently occur via email, desktop videoconferencing, and face-to-face meetings, but we envision these researchers collaborating over the Access Grid in a laboratory without walls or location.

Managing the huge amounts of simulation and experimental data this community generates presents an additional challenge.[17] With the Access Grid infrastructure, we can imagine a place where researchers can check out data sets from a virtual venue like books from a library. This would be possible if we used the Access Grid as a portal to the Grid. As part of the Globus Toolkit, researchers already get a secure connection to the Grid infrastructure by simply entering a virtual venue. This infrastructure includes resource discovery, grid-based I/O, and grid-enabled data management systems that make it easy to locate data for analysis, regardless of its actual location. Resources half a world away can process the data, and the results can be stored at some other location, as well. All these events can occur within the virtual venue, which provides a greater sense of collocation than other collaboration technologies currently available. Coupled with the Access Grid's ability to record human-to-computer as well as human-to-human interactions, this lets us fully capture all aspects of an analysis session. The recording's output can be a timeline for the process, a recipe for others to follow to perform similar analysis, or a narration of a major breakthrough.

The numerous challenges to realizing such a scenario include

- coupling the Access Grid to the more general Grid infrastructure;
- developing the appropriate standards, APIs, and examples;
- creating the tools needed to record, catalogue, and play back arbitrary data streams in a various formats;
- developing flexible annotation capabilities that allow users to make both public and secure private annotations; and
- developing an infrastructure to let new investigators know if someone else has already run a given filter, process, or test on the data, and to provide that data's location.

The Futures Lab at Argonne is addressing many of these challenges now.

Conclusions

The growth, acceptance, and increasingly diverse utilization of the Access Grid are indicators that many of the initial ideas expressed in its design have resonated within the user community. We believe that the ideas of persistence, presence, and immersion are key factors in the usability of advanced collaboration environments. As it exists today, the Access Grid is just one part of a more comprehensive computing, interaction, and visualization infrastructure that we think can be deployed in a workspace. Several research groups are exploring the concept of "smart spaces" or "active spaces," which aim to create work environments with embedded computing capabilities that support a broad range of user tasks. The Access Grid provides an important class of collaboration services that can be incorporated into these advanced environments. In addition, some groups are using the Access Grid to explore collaborative visualization modalities that let them share experiences and leverage distributed expertise in analyzing complex phenomena.

All these are exciting paths for Access Grid deployment and research. Our vision reflects the belief that bandwidth, computing, and imaging power will be increasingly available at decreasing cost within the near future, and that high-quality audio and video will be increasingly less expensive. Instead of focusing on how to get more image quality out of limited bandwidth, we expect the new challenge will be how to best organize these robust capabilities to support high-end scientific work, and how to create environments that encourage experimentation and interaction. Toward that end, we will continue with the SWoF program, integrating new devices, deploying more test beds, and building on the new capabilities of the Access Grid 2.0 toolkit.

Acknowledgments

We thank Joe Lambke (animate) for sketches. We also thank the SWoF application team members, particularly Bob Wilhelmson (NCSA), Don Middleton (NCAR), Erik Jacobson (UIUC), Natalia Maltsev (ANL), and Nagiza Samatova (ORNL). The Access Grid effort is the result of work by the entire Futures Laboratory team: Justin Binns, Ti Leggett, Randy Hudson, Ivan Judson, Eric Olson, Robert Olson, Mark Hereld, Ed Frank, Mary Fritsch, Dave Jones, Susanne Lefvert, and Thomas Uram. Funding for this work has been provided by the US National Science Foundation, the US Department of Energy under contract W-31-109-Eng-38, and Microsoft Research.

References

1. R. Stevens, "Access Grid: Enabling Group Oriented Collaboration on the Grid," *The Grid: Blueprint for a New Computing Infrastructure*, C. Kesselman, ed., Morgan Kaufmann, 2003.
2. H. Fuchs, "Beyond the Desktop Metaphor: Toward More Effective Display, Interaction, and Telecollaboration in the Office of the Future via a Multitude of Sensors and Displays," *Proc. Advanced Multimedia Content Processing*, LNCS 1554, Springer-Verlag, 1999, pp. 30–43.
3. T.L. Disz et al., "Designing the Future of Collaborative Science: Argonne's Futures Laboratory," *IEEE Parallel and Distributed Technology Systems and Applications*, vol. 3, no. 2, 1995, pp. 14–21.
4. R. Evard and R. Stevens, *Distributed Collaboratory Experimental Environments Initiative {LabSpace}: A National Electronic Laboratory Infrastructure*, US Dept. of Energy proposal, 1994; www.ccs.neu.edu/research/labspace/.
5. R. Evard, E.F. Churchhill, and S. Bly, "Waterfall Glen: Social Virtual Reality at Work," *Collaborative Virtual Environments*, M.S. Churchill, ed., Springer-Verlag, 2001, pp. 265–281.
6. P. Curtis et al., "The Jupiter Audio/Video Architecture: Secure Multimedia in Network Places," *Proc. ACM Multimedia 95*, ACM Press, 1995.
7. T.L. Disz et al., "Two Implementations of Shared Virtual Space Environments," tech. report ANL/MCS-P652-0297, Argonne Nat'l Laboratory, 1997.
8. J.A. Leigh and T. DeFanti, "CAVERN: A Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments," *Virtual Reality: Research, Development and Applications*, vol. 2, no. 2, 1997, pp. 217–237.
9. S. McCanne and V. Jacobson, "VIC: A Flexible Framework for Packet Video," *Proc. ACM Multimedia*, ACM Press, 1995, pp. 511–522.
10. P. Curtis and D.A. Nichols, "MUDs Grow Up: Social Virtual Reality in the Real World," *Proc. 3rd Int'l Conf. Cyberspace*, Xerox Palo Alto Research Center, 1993.
11. I. Foster et al., "A Security Architecture for Computational Grids," *Proc. 5th ACM Conf. Computer and Communications Security*, ACM Press, 1998, pp. 83–92.
12. T. L. Disz et al., "The Argonne Voyager Multimedia Server," *Proc. High Performance Distributed Computing*, IEEE Press, 1997.
13. H. Gao, R.L. Stevens, and M.E. Papka, "The Design of Network Services for Advanced Collaborative Environments," to appear in *Proc. 3rd Workshop on Advanced Collaborative Environments*, Argonne Nat'l Laboratory, 2003, pp. 53–58.

Rick Stevens is the director of the Mathematics and Computer Science Division at Argonne National Laboratory and professor of computer science at the University of Chicago. His research interests focus on using advanced computing and communications technologies to accelerate scientific research. Stevens received a BS in applied mathematics and philosophy from Michigan State University. Contact him at stevens@mcs.anl.gov.

Michael E. Papka is deputy director of the Futures Laboratory of the Mathematics and Computer Science Division at Argonne National Laboratory. His research focuses on the design of new visualization technologies that combine storage systems, networking, virtual space technology, and virtual environments for scientific research. Papka received an MS in computer science from the University of Chicago. Contact him at papka@anl.gov.

Terry Disz is an experimental systems engineer at Argonne National Laboratory. His research interests include computer graphics, grid programming, and collaboration technologies. He has been a member of the Futures Lab since 1994 and has spoken and written extensively on their work. He received a PhD in electrical engineering and computer science from the University of Illinois, Chicago. Contact him at disz@mcs.anl.gov.