THE ACCESS GRID: PROTOTYPING-WORKSPACES OF THE FUTURE

RICK STEVENS, MICHAEL E. PAPKA, TERRY DISZ,
AND THE MEMBERS OF THE FUTURES LAB OF
ARGONNE NATIONAL LABORATORY AND
THE UNIVERSITY OF CHICAGO

INTRODUCTION

Collaborative immersive virtual reality technology has been in use since the late 1980s. In the mid-1990s these systems were used to investigate multi-user wide-area collaboration scenarios. While these efforts were pioneering in many respects, they proved less suitable as work environments for everyday use. People tire easily when spending extended time in the dark spaces needed for projection virtual reality or when being immersed in a completely synthetic world for hours at a time without access to high-resolution text displays or high-quality interactions devices.

During the past decade, desktop multimedia technology capabilities have also grown considerably, to the point that all modern desktop systems can easily handle video and audio. Indeed, desktop PCs have exceeded 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[1] 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 twenty-first century. This paper shows how we view the Access Grid as a first step toward room-based computing environments [2] that will, we believe, challenge desktop metaphors, desktop user environments, and perhaps even desktop computer deployment in the decades ahead. Instead of simple, single-stream videoconferencing implemented with special-purpose desktop appliances, the environments we envision are entire rooms or laboratories instrumented for rich, full-time, multimodal communications between multiple groups of people connected over inherently multipoint, high-speed networks. We offer three scenarios showing how this vision is becoming a reality.

EVOLUTION OF THE ACCESS GRID

The Access Grid has emerged as one of an important class of Grid application designed to support wide-area, real-time, computer-mediated communications. The Access Grid project grew out of a long-standing set of research directions being pursued at Argonne National Laboratory's Futures Laboratory [3]. Starting in 1994, the Futures Lab began developing new types of collaboration environments with aggressive assumptions about the future of networking and computing. The Access Grid, the result of this work, contains influences of the LabSpace project[4], MOOs[5], Jupiter[6], ManyWorlds[7], and CAVERNsoft projects [6, 8]. Like these earlier works, 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 these earlier works, however, the Access Grid is designed to exploit high-performance peer-to-peer multimedia services (audio [9], video [10], text [11]) in the creation of multi-person, shared virtual workspaces. The goal of the Access Grid is to create computer-augmented workspaces with the capability of supporting natural (full-duplex) audio and video communication for distributed workers.

Figure 1: Photograph of group at Argonne National Laboratory participating with other groups in the National Science Foundation's announcement of the Extended TeraGrid Facility.

While a number of research groups have focused on the concept of connecting users via desktop video and audio services, the Access Grid project took two important steps away from this tradition. The first step was to focus on creating environments for *small groups of users* rather than individual users (i.e., focusing on room-scale systems with large display surfaces and dedicated computing, display, and multimedia infrastructure); see Figure 1~~Figure 1~~. Arguably, large-scale Grid-oriented scientific collaborations often involve dozens of institutions and hundreds of researchers. Indeed the development of the Grid itself is a response in part to the need to support highly distributed scientific projects. Many of these large-scale collaborations have emerged as virtual organizations (including NSF Centers, NSF ITRs, NIH Research Resource Centers, and DOE SciDACs Centers); they have a persistent management structure, significant number of shared tasks undertaken by the group, and significant resources dedicated to achieving shared goals. Within most such large projects, however, one can often identify core teams of researchers and core institutions. The initial design point for the Access Grid, then, was to target groups of approximately six to eight users per site at about eight sites as the ideal grouping for Access Grid. This design goal clearly distinguishes the Access Grid from other collaboration environments. By focusing on small-group–oriented technology, the Access Grid project has been able to specify a high-end resource model with more capabilities than those systems designed for desktop deployment (i.e., more shared bandwidth, higher-end audio equipment, dedicated computing systems).

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 AG 1.0 virtual venue software is web based and manages multicast address allocation and media tool startup. Virtual venues are located by name, linked from a central web site. The media tools used by the Access grid are a version of VIC[10] modified for more efficient rendering, the RAT version of VAT[9] and a text chat tool[11]

An important metadriver for the Access Grid was to imagine a scenario based on two important assumptions: (1) bandwidth would be at least several orders of magnitude more than available today and (2) computing cycles would be inexpensive to the point that we can embed them in the infrastructure for meeting rooms. Both of these assumptions led to the idea that the Access Grid could support many parallel video and audio streams sources from each site; that these streams could be made available to all other sites; and that dedicated computing resources for audio processing, video capture, screen control, and collaborative services could easily be made available at each site. These features combine to provide a rich set of communications channels. Both of these resource assumptions have played a role not only in our initial Access Grid design point but also in the philosophical aspects of the Access Grid development.

Because of the desire to support groups interacting with groups, the Access Grid is built around the idea of a multi-user semi-public workspace that supports audio and video connections to other such spaces. In these workspaces users are free to move around, carry on local conversations and interactions, and enter and exit the room in an ad-hoc manner. The use of large display surface 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.

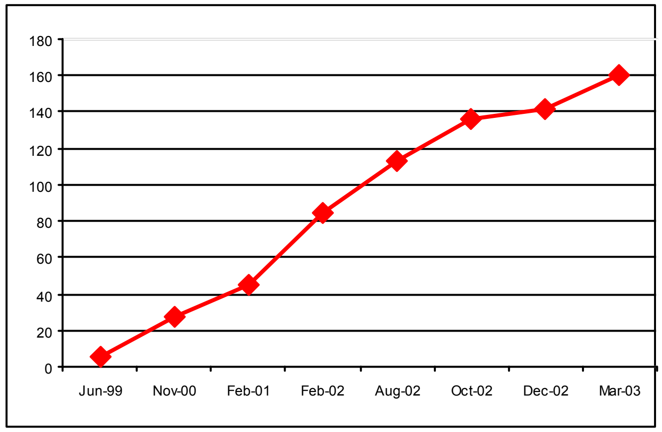

Figure 2: Number of registered Access Grid nodes sites since June 1999 at research institutions.

The success of the Access Grid is attested to by the fact that over 200 Access Grid nodes are deployed worldwide in private companies and research institutions and used almost daily for distributed meetings, seminars, classes, and a host of other collaborative activities, with over 160 of these nodes being registered at the Argonne Access Grid web site. These nodes are typically connected via the high-speed Internet services of Internet2, ESnet, or other regional, national, or international research and academic networks. Figure 2~~Figure 2~~ shows a graph of the growth rate of Access Grid nodes in research institutions around the world since June 1999; Figure 3 shows the number of scheduled meetings per month since September 1999; the spike that appears in October of 2001 reflects a large number of sites preparing for the SCGlobal event that was held at the Supercomputing 2001 conference; the sharp rise in early 2002 could be attributed to world events of the previous fall.

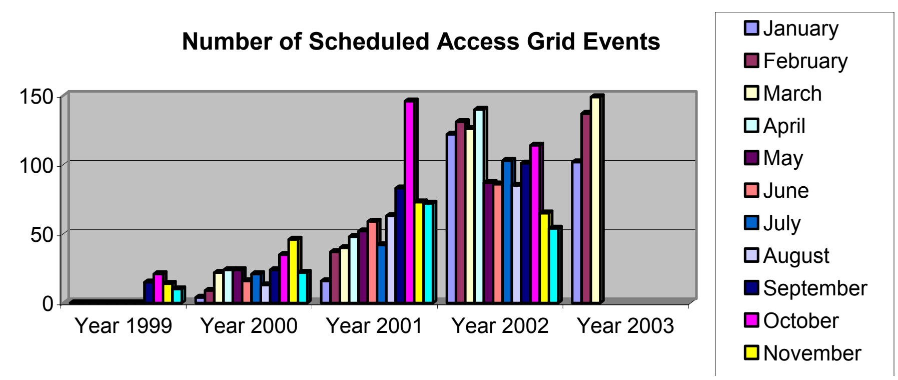

Figure 3: Chart of scheduled events on the Access Grid, does not include the use of institutional rooms.

The latest evolution of the Access Grid is taking place in the form of the AG 2.0 software, to be released in late Spring 2003. The AG 1.0 virtual venue server was driven from a web server and provided minimal security, multicast address mapping onto virtual rooms and launched media clients, but provided no mechanism for adding more functionality. AG 2.0 is based on a highly secure web services model, giving users and developers opportunities to create new services and add them to a virtual venue for the benefit of all. This new release will greatly broaden the audience of potential developers and service providers, giving the user community more ownership and control of the final product.

ACCESS GRID DESIGN REQUIREMENTS

In our design and development of the Access Grid, we have identified several design requirements:

Presence: One goal of an advanced collaborative system is to reproduce as many physical presence cues as possible in an attempt to recreate the feeling of co-presence among the attendees, regardless of their actual location. The Access Grid provides both audio and video mechanisms to deliver a sense of presence in meetings. To provide natural style audio communications, the Access Grid specifies the use of multiple, always on, full-duplex operating microphones, situated in the space to provide maximum coverage. Access Grid video aids in natural style communication by presenting near life scale images, increasing the feel of "being there" with the participants.

High-Quality Audio: The Access Grid implements high-quality audio, including high-quality microphones, high quality codec (16 KHz codec instead of 8 KHz), and commercial echo cancel products to prevent audio feedback and echoing and to suppress unwanted noise.

Context and Continuity: With only a single video stream available to users, a typical video-conference may appear to each user as a shifting landscape of different views with little or no context. The Access Grid has two solutions to this problem. One is to provide multiple camera views – with three cameras on the participants, the Access Grid delivers multiple overlapping perspectives, and the one camera on the display provides important contextual cues to remote sites. Additionally, all these streams are delivered simultaneously, to all sites, providing everyone with enough context to fully understand and participate in the meeting.

Consistent, Designed Meeting Places: Designing a large space with audio and video equipment in a consistent manner is important to providing consistent views on the space. The Access Grid provides a set of best practice recommendations (see the web site www.accessgrid.org) for room layout and architecture, including lighting, camera locations, microphone locations, and seating. At Argonne, "life-scale" means our video images are a little over 4' high, which gives us a little over a 16' display wall using three projectors. These figures dictate everything else about the space; how far away people can sit, how wide the viewing area can be, how many people can be accommodated, how many microphones and their placement and such things as lighting and sound management.

Internet connection: Supporting these decisions means that the Access Grid software infrastructure is required to provide multiple high-quality video and audio streams and manage their distribution. This requires that Access Grid users have a robust internet connection – we recommend at least 20Mbs. Choosing to use multicast as a distribution layer for the media streams creates a second-order requirement for a set of tools to manage multicast.

Multicast Infrastructure Stability: Tools to debug and measure multicast performance have been developed by the National Laboratory for Applied Network Research. These tools center on a multicast beacon that periodically transmits a signal and listens for signals from all other beacons on a given multicast address and include a tool to analyze and visualize the results delivered by the deployed beacons.

A second set of requirements, presented below, is driven by the fact that the AG meeting participants are not in the same place and still want to collaborate and do science.

Meeting Management: All meetings face logistical and management problems as they proceed. In Access Grid meetings, these problems are exacerbated by the fact that not all participants are co-located. The Access Grid offers best practices and technology to address managing remote meetings. A text back channel is supplied that replaces private face to face meetings and whispering. An audio back channel is provided out of band from normal AG communications and is useful where real-time command and control is essential to the meeting. Lastly, there are various meeting schedulers available which can provide automatic scheduling of spaces, document repositories and participant notification.

Navigation, Scope of Interaction: Real-life meetings are held in rooms that control the scope of interaction, and people find these rooms by name or number. The Access Grid provides a similar model with the "virtual venue". Instead of a phone number or a network address, users are directed to a room by name. Entrance to the room is guarded by authorization mechanisms to prevent unwanted attendees. Persistent virtual rooms contain meetings and meeting objects much as real rooms do.

Grid integration: Targeting science communities for fully collaborative meeting places requires that those places support science applications that increasingly require access to Grid resources. The Access Grid supports full integration with Grid tools and resources. AG 2.0 is built using the GSI security toolkit, securing all data, communications and resource usage. AG 2.0 is a Grid services–based environment, allowing users to discover and run Grid applications and to create new, Grid-compliant applications.

Our current work on the development of AG 2.0 [12](the second generation of Access Grid tools) is layered on the Globus Toolkit®, and plans are in place to migrate AG 2.x to the Open Grid Services Interface [13, 14]. The work on the AG is therefore closely coupled to the future evolution of mainstream Grid software infrastructure.

ACCESS GRID: THE NEXT STEP

The Access Grid architecture (see Figure 4) is ultimately driven by the requirement for a workspace designed to provide features and interaction not available in previous systems. We have identified the following elements as high priority for development, all of which will have significant positive impact on end users and those wishing to apply the concepts of the Access Grid to new areas of scientific collaboration.

- **Scalable Virtual Venue Service:** The AG 1.0 currently offers a set of "rooms," or virtual spaces, mapped to multicast addresses. This version is not scalable and does not provide persistence beyond simple presence. The next step, AG 2.0, will implement a peer-to-peer venue service operating much the same way as the Web: Anyone can host a server (a virtual space), and anyone on the network can visit. The new server creates a venues service that scales to thousands of nodes, with no centralized services, where anyone can trivially create new spaces and link them into the peer-to-peer infrastructure.

- **Access Grid Security:** In the past, AG users have generally taken a casual approach to security and privacy during Access Grid sessions. However, in many applications, there is a need for real security, for the availability of robust access control and privacy of shared media and applications. The AG security model leverages the large body of existing work in Internet security in general and Grid security in particular [15]. This model is based on the use of public-key infrastructure standards for identity certificates and the Globus toolkit extensions for proxy certificates to implement a secure authentication framework with support for single sign-on. Transport security is provided by the Globus toolkit as well, which in turn uses the well-understood SSL protocol for point-to-point secure connections. We use these mechanisms to distribute keying information to the endpoints of the multicast-based media tools in order to support privacy of these multipoint media sessions.
- **Application Sharing and Dynamic Workspace Dockin**g: Workspace docking is analogous to docking a laptop into a network to gain access to local services. Our plans for workspace docking allow for the ad-hoc construction of user space Grids where the AG users can "share" a portion of their personal workspace (desktop applications and data) with other AG users, nodes, or sites, both local and remote. The docking infrastructure involves migrating or launching one or more specific application clients (linked with multicast as needed) onto the AG displays and attaching them to the user's server.
- **Node Management and User Interfaces**: 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 the management domain of a node and to enable shared applications to exploit these new capabilities.
- **Asynchronous Collaboration Capabilities:** Work is under way to extend the Voyager[16] system to include streaming data types required to capture the interactions and events that occur in the persistent spaces of the Access Grid. 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.
- **Network Services:** Network services are resources available to an AG session that live somewhere accessible on the network. Examples of network services could be video and audio transcoders or speech-to-text converters. The network services architecture concentrates the required functionality into a small number of key components: 1) The Network Services Engine which will act as the principal point of contact between the users of the AG Network Services, 2) The Virtual Venue which enables the brokering of access to the services, and 3) the Resolution Engines that will provide the detailed support for the determination of the appropriate instances of services to be supplied.

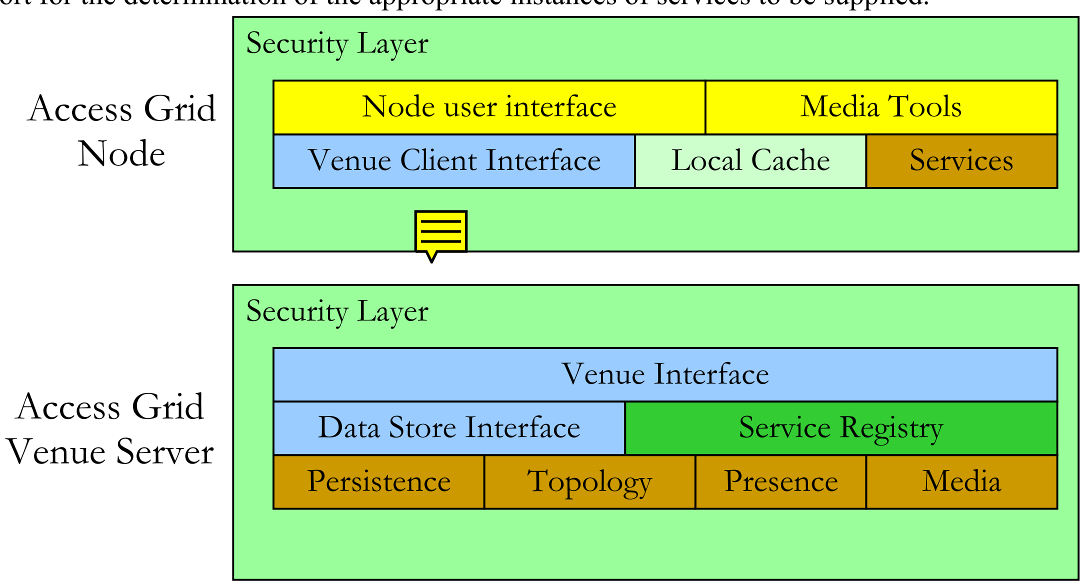

Figure 4: Access Grid services-model architecture showing the division between nodes and venues and the layered architecture for each component. All communications and storage operations are conducted through secure interfaces.

COLLABORATION SCENARIOS AND CHALLENGES

During the past year we have been working with two communities to explore the use of the Access Grid to enable them to do their science with remotely located colleagues , and we are planning to open communications with a third community later this year. The current work is part of the National Computational Science Alliance (NCSA) Expedition projects. An "expedition" is a narrowly focused project bringing together technology and applications partners to create a problem-solving-based model of scientific collaboration. The expedition *Scientific Workspaces of the Future* (SWoF) is led by Argonne and includes participation 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 goal of the SWOF expedition is to develop virtual communities for the fields of atmospheric sciences and molecular biology using Access Grid infrastructure. We describe these projects, as well as our planned work with the high energy physics community, and then discuss the challenges associated with them.

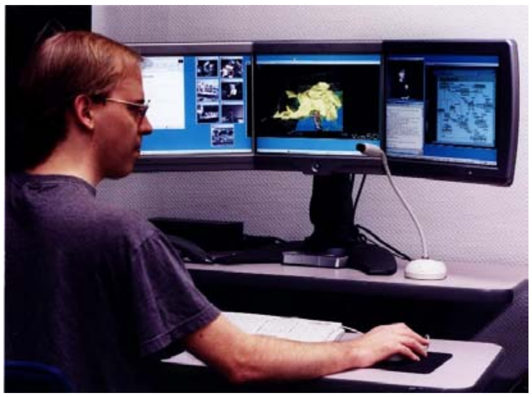

Figure 5: Photograph of user interacting with other researchers and visualization of a climate model via a personal interface to Grid.

ATMOSPHERIC MODELING AND SIMULATION COMMUNITY

A stated major goal for the atmospheric sciences community is the deployment of common collaboration and visualization infrastructure throughout the member sites of the University Corporation for Atmospheric Research (UCAR) within three years. UCAR members include all U.S. schools that grant doctorates in atmospheric, oceanic, and related sciences. There is strong interest in integrating the human and experimental resources of these institutions, but as yet no focused or funded program to accomplish this integration. The proposed Alliance 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 to have multiple institutions actively participate in the formulation and analysis of large-scale computational models of the ocean and atmosphere (see Figure 5). These models produce large amounts of output that need to be visualized greatly exceeding the capabilities at most UCAR institutions; yet much of the national expertise for 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 could be greatly enhanced by having remote courses and seminars from other similar departments.

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, including the first distributed national seminar and tutorials on computational biology and bioinformatics via the Access Grid and a major effort to develop a national computational biology curriculum to be delivered nationally via the Access Grid. The proposed expedition will focus on developing the software infrastructure needed to build collaboratories linking the biomedical research community. The project goal will be to deploy Alliance SWoF expedition technology to all major U.S.-based systems biology research centers (about two dozen sites) within three years, to help create a National Systems Biology Institute, modeled after the NASA supported Distributed National Astrobiology Institute. NIH 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 gained some acceptance in the computational molecular biology community as noted above, the challenge is to provide shared access to computational tools and analysis systems for genomics and for 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) to visualize the large-scale networks present in even the smallest of cells and 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 that will enable them to scale across the TeraGrid and other discipline Grids that are being constructed, but will also enable the tools to leverage Grid tools being developed for data and computation.

HIGH ENERGY PHYSICS COMMUNITY

The high energy physics community comprises researchers worldwide working on theoretical, simulation, and experimental efforts. The work of this community is frequently collaborative; typical experiments can have over a hundred collaborators representing tens of different organizations spread around the entire world. Currently such collaborations are via e-mail, desktop video conferencing tools, and face-to-face meetings. We envision a situation in which these researchers collaborate by using the Access Grid: The Virtual Venue becomes a laboratory without walls or location.

In addition to the challenge of collaboration, however, is the challenge of managing the huge amounts of simulation and experimental data the high energy physics community generates [17]. By using the Access Grid infrastructure to construct a virtual laboratory, we can imagine a place where this data is available in the venue as books are in a library; researchers are able to check out datasets. This can be made possible by allowing the Access Grid to act as portal to the Grid. Simply by entering a Virtual Venue, researchers already have a secure connection into the Grid infrastructure provided as part of the Globus Toolkit tools, an infrastructure that includes access to resource discovery, Grid-based I/O, and Grid-enabled data management systems. The data needed for analysis can easily be found, regardless of its actual location. It can be processed by resources that might be half a world away, and the results can be stored at another location. All these events can occur within the Virtual Venue, providing a sense of collocation greater than in any other technology currently available. With this ability, and the ability of the Access Grid to record not only human-to-human interactions but human-to-computer interaction, one can fully capture all aspects of an analysis session. The output of the recording can be a timeline for the process, a recipe for others to follow to perform similar pieces of analysis, or a narration of a major breakthrough.

The challenges to realizing such a scenario are numerous:

- Coupling of 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 variety of different formats

- Developing annotation capabilities so that users can annotate data in a flexible manner, allowing for secure private as well as public annotation, enabling future users to benefit from (but not be distracted by) other researchers' annotations
- Developing an infrastructure to let a new investigator of the data know that someone else has already run that filter or process or test on the data, and to know the location of that data

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

CONCLUSIONS

The Access Grid is designed to be used by groups in explicit high-end workspaces. This focus on groups and relatively high-end applications has led to the need for high-quality audio and video and the need to move away from ad-hoc deployment environments. Many groups are now designing a studio-like space for Access Grid deployments that give users control over lighting, sound quality, and video, resulting in the capture of high-quality video images. The more professional the video and audio, the more effective the user experience.

Often the Access Grid is only one part of a more comprehensive computing, interaction, and visualization infrastructure 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 capability targeted at supporting a broad range of user tasks. The Access Grid is an important class of collaboration services that can be incorporated into these advanced environments.

Another important use of the Access Grid is as the technology basis for building "persistent virtual project rooms." Persistent project rooms enable groups to maintain project-related materials and applications in a form that is continuously available to team members.

In addition, the Access Grid is being used to explore collaborative visualization modalities that will enable groups to share visualization experiences and leverage distributed expertise in the analysis of complex phenomena.

All these are exciting paths for Access Grid deployment and research. Our vision of the Access Grid reflects the belief that within the near future that bandwidth, computing, and imaging power will become effectively free and that high-quality audio and video capture will be increasingly inexpensive. The primary challenge will not be how to get more image quality out of limited bandwidth. Rather, the challenge will be how to best organize these capabilities to support high-end scientific work—how to create environments that encourage experimentation and interaction.

ACKNOWLEDGMENTS

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

REFERENCES

1. Stevens, R., *Access Grid: Enabling Group Oriented Collaboration on the Grid*, in *The Grid: Blueprint for a New Computing Infrastructure*, C. Kesselman, Editor. 2003, Morgan Kaufmann.
2. Fuchs, H., *Beyond the desktop metaphor: Toward more effective display, interaction, and telecollaboration in the office of the future via a multitude of sensors and displays*, in *Advanced Multimedia Content Processing*. 1999. p. 30-43.
3. Disz, T.L., et al., *Designing the Future of Collaborative Science: Argonne's Futures Laboratory.* IEEE Parallel and Distributed Technology Systems and Applications, 1995. **3**(2): p. 14 -21.
4. Evard, R.S.a.R., *Distributed Collaboratory Experimental*
Environments Initiative
{*LabSpace*}*: A National Electronic Laboratory Infrastructure*. 1994.
5. R. Evard, E.C.a.S.B., *Waterfall Glen: Social Virtual Reality*
at Work, in *Collaborative Virtual Environments*, S. Churchill, Munro, Editor. 2001, Springer-Verlag. p. 265-281.
6. Curtis, P., et al. *The Jupiter Audio/Video Architecture: Secure Multimedia in Network Places*. in *ACM Multimedia 95 - Electronic Proceedings*. 1995.
7. Disz, T.L., et al., *Two Implementations of Shared Virtual Space Environments*. 1997, Argonne National Laboratory: Argonne.
8. Leigh, J., J. A., and T. DeFanti, *CAVERN: A Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments.* Virtual Reality: Research, Development and Applications, 1997. **2**(2): p. 217-237.
9. Jacobsen, V. and S. McCanne, *vat: LBNL Audio Conferencing Tool.*
10. McCanne, S. and V. Jacobsen, *Vic: A flexible Framework for Packet Video.* ACM Multimedia, 1995: p. 511-522.
11. Curtis, P. and D.A. Nichols. *MUDs Grow Up: Social Virtual Reality in the Real World.* in *Third International Conference on Cyberspace*. 1993. Austin, TX.
12. Stevens, R., et al., *Middleware to Support Group to Group Collaboration*. 2001, DOE SciDAC National Collaboratories and High Performance Networks.
13. Foster, I. and C. Kesselman, *Globus: A Metacomputing Infrastructure Toolkit.* International Journal of Supercomputer Applications, 1997. **11**(2): p. 115-128.
14. Foster, I. and C. Kesselman, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations.* International Journal of Supercomputer Applications, 2001. **15**(3).
15. Foster, I., et al. *A Security Architecture for Computational Grids*. in *5th ACM Conference on Computer and Communications Security Conference*. 1998.
16. Disz, T.L., et al., *The Argonne Voyager Multimedia Server*. 1997, Argonne National Laboratory: Argonne.
17. Foster, I., et al. *Chimera: A Virtual Data System for Representing, Querying and Automating Data Derivation*. in *14th Conference on Scientific and Statistical Database Management*. 2002. Edinburgh, Scotland.