

Access Grid: Immersive Group-to-Group Collaborative Visualization

Lisa Childers,¹ Terry Disz,¹ Robert Olson,¹ Michael E. Papka,¹ Rick Stevens^{1,2} and Tushar Udeshi¹
{childers, disz, olson, papka, stevens, udeshi}@mcs.anl.gov
Mathematics and Computer Science Division
Argonne National Laboratory^{1*}

Computer Science Department
University of Chicago²

Abstract

Immersive projection displays have played an important role in enabling large-format virtual reality systems such as the CAVE and CAVE like devices and the various immersive desks and desktop-like displays. However, these devices have played a minor role so far in advancing the sense of immersion for conferencing systems. The Access Grid project led by Argonne is exploring the use of large-scale projection based systems as the basis for building room oriented collaboration and semi-immersive visualization systems. We believe these multi-projector systems will become common infrastructure in the future, largely based on their value for enabling group-to-group collaboration in an environment that can also support large-format projector based visualization. Creating a strong sense of immersion is an important goal for future collaboration technologies. Immersion in conferencing applications implies that the users can rely on natural sight and audio cues to facilitate interactions with participants at remote sites. The Access Grid is a low cost environment aimed primarily at supporting conferencing applications, but it also enables semi-immersive visualization and in particular, remote visualization. In this paper, we will describe the current state of the Access Grid project and how it relates and compares to other environments. We will also discuss augmentations to the Access Grid that will enable it to support more immersive visualizations. These enhancements include stereo, higher performance rendering support, tracking and non-uniform projection surface.

A Platform for Semi-Immersive Collaborative Visualization

Advances in large scale computing have allowed researchers to produce more data than ever, requiring

aggressive research in large scale and distributed visualization in order to try and keep up with this new flood of data. These research visualization projects are frequently combines of different groups, each contributing different specific skills [1, 2] These combines include as many as six different partners, located at various universities and laboratories across the U.S. Typically, each group consists of a small number of researchers and programmers, often less than a half dozen.

Immersive projection technology devices, typically CAVEs [3], immersive desks, and their derivatives are often used as superior tools for data visualization and analysis of these very large amounts of data. Non-located teams can use these devices, along with various collaboration technologies, to visualize and examine data in a collaborative manner with colleagues that have similar capabilities [4]. Some drawbacks to this approach, however, are that these devices are relatively expensive and not commonplace and that they are designed specifically for visualization, not collaboration. We believe that collaborative visualization requires a new model of collaboration environment, one that fosters a sense of presence among the groups, incorporates semi-immersive visualization capabilities, is available all the time, is easy to use and helps encourage the social processes important to team development.

The challenge we are interested in is how to build and use a low cost platform that satisfies these requirements and learn how to use it to bridge into the immersive visualization environments. Existing desktop and telephone options are less than satisfactory [5], not being well suited for multiple group to group interactions, and certainly not conveying a sense of presence among the participants nor providing any sense of immersion. Our approach is to build a semi-immersive collaborative environment by leveraging off-the-shelf commodity hardware and software and to employ an open source

* This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, U.S. Department of Energy, under Contract W-31-109-ENG-38, and the National Science Foundation under NSF Cooperative Agreement #ACI-9619019.

policy whenever possible. In our use, semi-immersive means that we provide a wide field of view, natural audio (hands free, full duplex), multiple perspective video streams and a large pixel count.

In the next section, we describe the Access Grid [6], a semi-immersive collaboration environment developed at Argonne. Following that, we describe some investigation we have started to explore bridging the Access Grid into CAVE collaborative visualization sessions.

The Access Grid: The vision

Our own unsatisfactory experiences with desktop collaborative technology caused us to re-think what was really required to enable wide area collaboration. First, we realized that we most often worked with colleagues as small groups and so we began to think in terms of wide area *group* collaboration. Second, while we all attend structured meetings, workshops, etc., we find we often tend to be more productive in an unstructured manner with lots of brainstorming, problem solving, casual conversation and spontaneous idea generation. From this insight we realized the need to support multiple modes of interaction, from very structured to completely casual. Third, we usually have our personal portable computers with us and often want to share with other individuals or the group, some idea expressed on our computer, be it a visualization, a spreadsheet, a presentation, a web site, a document or a movie. Last, but not least, we realized that one of the problems plaguing existing efforts was the perceived need to accommodate wide ranges of pre-existing equipment, software and capabilities. We could see there would be significant advantages to be gained from having all participants use exactly the same hardware and software.

For our ideal collaborative environment we envision an intentionally designed space, one that would be rewarding to be in, one that provides a sense of co-presence with other groups using similar spaces. We envision a space with ambient video and audio, large-scale displays and with software to enable the relatively transparent sharing of ideas, thoughts, experiments, applications and conversation. We envision a space where we can “hangout” comfortably with colleagues at other places, and then use the same space to attend and participate in structured meetings such as site visits, remote conferences, tutorials, lectures, etc. We imagine the space will support the same capabilities, through remote interaction, that we have now in face-to-face meetings – subconscious floor control through social

conventions, the ability to have private one-on-one, whispered conversations, the ability to gather up small groups in a corner and caucus, and all the other things we take for granted when we are all in the same place. In addition, we envision the space being “smart” enough to recognize that you have brought personal computing resources to it and allowing you to export items from your computing devices to other individuals or groups.

The challenges this vision presents are many and varied, some easily addressed, others requiring groundbreaking research efforts. Other similar efforts to break from the desktop[7] also feature large-scale displays and instrumented spaces.

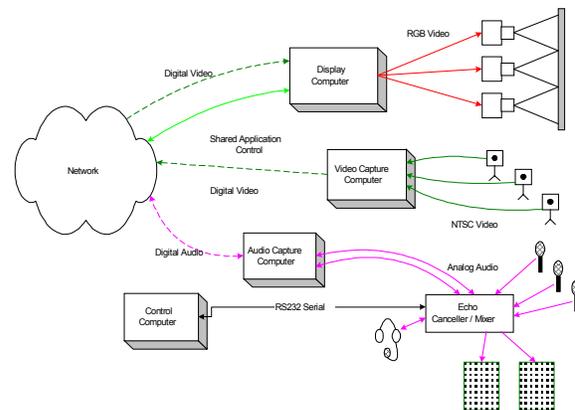


Figure 1. Schematic of Access Grid Node

Access Grid Research Issues

Supporting group oriented collaboration sessions multiple sites requires a system architecture capable of scalable wide area connectivity. A key part of this capability is the efficient transport of audio and video streams over the network. We have chosen to use IP multicasting and the media tools that grew out of the research efforts that brought multicast into general use [8-10]. To control the scope of who is attending meetings and collaborating, we have chosen to implement a spatial metaphor of rooms which are always present, rather than the more usual, but heavyweight operation of creating a session for each time people want to meet. We believe it is important in groups to have a persistent space in both the virtual and physical world. We have created a set of virtual rooms that are always available, and have specified that the physical Access Grid nodes themselves be dedicated spaces, available for drop-in use much as virtual rooms are.

One source of a sense of presence in real life meetings is implicit awareness about where other people are with respect to one's own location in the meeting space. This awareness is gained by being able to see and hear other people and to build a mental map of the space and the people in the space [11]. The challenge is in reproducing this awareness among non-located people. To achieve a sense of presence, we have designed the Access Grid to use hands-free, full duplex audio and to deliver four multiple perspective video streams from each location. This causes a large number of streams to be delivered across the Access Grid, requiring a careful balance between network constraints, encoding and decoding capabilities and the desire for high resolution, high frame rate video and high quality audio.

Due to these large numbers of media streams and the heavy reliance of the Access Grid on wide-area IP multicasting and given the current state of the implementation of multicast on the networks we use, network monitoring is a crucial part of the AG toolkit. Debugging multicast problems based only on subjective user observations is very time consuming and inaccurate. Our approach has been to use a multicast channel beacon that runs at each AG site and constantly transmits a low level signal. It also listens for signals from other beacons and reports to a central server where network connectivity statistics are maintained. [12]

To facilitate network as well as other Access Grid debugging, a secondary communications channel is required. This back channel must be out of the primary communication band and yet be available to all participants. Our solution has been to implement a MOO [13, 14] for the AG community. This has had the effect of providing a solid text based channel for debugging, and at the same time, a persistent community space for use at all hours and has served us very well in creating an Access Grid user community with all the normal social interactions necessary for collaborative work.

When performing scientific collaborative work, being able to record the process is important to good science. With the Access Grid being capable of delivering several dozens of IP based media streams, recording in this environment is a significant challenge. The requirements for a system to record a virtual meeting are that it must be well enough connected and robust enough to be able to sink and save on disk multiple streams of audio and video without loss. The system must be able to play back the multiple streams, synchronized in time so as to

faithfully reproduce the sequence of events. Argonne has built such a capability in the Voyager Multimedia Multistream record and playback engine [15, 16].

Access Grid Capabilities today

In realizing the Access Grid, we focused on basic enabling infrastructure for groups of people to find, talk to, see and share ideas with other groups. Our philosophy is to use open source software wherever possible. This focus generated requirements for displays, computing, audio, video, room architecture, network, and software tools.

Display. An Access Grid Node, as we call a single room outfitted for AG use, requires a tiled display of sufficient physical size to comfortably accommodate a small group of people – up to a dozen or so, sitting around the display, all with good sight lines to the display. Secondly, the display must have sufficient resolution and size to accommodate the projection of multiple video streams from multiple sites, projecting near life size images of people at other sites. Solutions to this vary, but we are most satisfied with a three projector, front projection wall. The projected area is about 18' by 6' with a seating area of about 25' wide by approx 20' deep. The projectors are ceiling mounted and of sufficient brightness that the room can operate in normal light so people can read and interact.

Video. An Access Grid Node must generate multiple video streams from different perspectives in the room in order for people at other sites to get a feel for the room and its occupants. We specify 4 video streams: A wide audience shot, a close-up shot of the presenter or main speaker, a wide area shot of the display screen (it is important for remote sites to be able to see what you are seeing), and last, a roving audience and room camera. They are placed so as to be unobtrusive and to facilitate the feeling of eye contact.

Audio. Being able to freely converse with people at other sites, unencumbered by microphones, wires, floor control protocols or gadgets, is a cornerstone of AG usability. We achieve this ability by placing sufficient numbers and types of microphones and speakers with the space. We make sure there is adequate pickup everywhere in the room there are likely to be people. Secondly, we employ professional quality echo cancellation gear by Gentner Corp. to ensure full-duplex audio. We currently use two speakers placed strategically in the front of the room to project good quality audio into the space.

Computers. An Access Grid Node uses four computers. The Display Computer runs Windows NT and has a multi-headed video card. This machine manages the tiled display and allows us to treat the multiple projectors as a single desktop. It is decoding all of the video streams, which can range into the several dozens, so needs to be as robust as possible. The Video Capture Computer runs Linux and has 4 video capture cards. The Audio Capture computer also runs Linux and performs the audio encoding and broadcasting as well as the audio decoding of the multiple streams being sent from other AG Nodes. Lastly, the Control Computer runs Windows 98 and is used to run control software for the audio gear. This separation of function allows us to optimize each piece of gear for its intended purpose

Software. In addition to the operating system mentioned above, a compliant AG Node requires several pieces of software developed by Access Grid partners, including a Multicast beacon and viewer, distributed PowerPoint tools, a MOO client and the UCL Mbone tools, VIC and RAT. Persistence and scope are provided by using the Virtual Venue software developed at Argonne. The Virtual Venue software contains a set of rooms in which AG node participants can interact. This is a method of allocating, controlling and automatically assigning multicast addresses. This software allows users to leave one group and join another by simple clicks on a web based map interface. The software automatically tears down existing connections and builds new ones as dictated by the addresses related to each room. With several dozens of media streams being delivered on the Access Grid, managing windows on the display space is a challenge. To assist in this, we have provided auto placement software that automatically lays out video windows across the screen real estate, based on pre-selected preference.

Network. The Access Grid tools depend on network multi-cast to work well. Until native multicast connectivity is ubiquitous, we must accommodate sites without multicast capability are with a multicast to unicast bridge. We use the Fermi MultiSession Bridge[17]. Use of the bridge introduces delay, complexity, and significantly increases network load. Sites wishing to become Access grid Nodes should see that multicast capability is supplied to their site. The other practical network consideration is available bandwidth. A full AG session can deliver many dozens of video streams to a Node, typically four from each participant as well as your own. The bandwidth required by each stream is dependent on settings at the origination and can vary from 128Kb/s

to 512Kb/s or more. The effect of inadequate bandwidth on the AG Node is dropped packets resulting in unintelligible audio and jerky-motion video. Other effects can be detrimental performance for, and eventual hostility on the part of, other users on the local network.

Production. The Access Grid as of spring of 2000 has over 20 sites up and running and with nearly 20 more sites planning deployment in the next year. The Access Grid has been used in several major events in 1999 – the ACCESS DC grand opening event in April[18], the three Alliance Chautauquas in the summer[19], and at SC99, where several sites brought Nodes to the show floor while others participated from their home sites. From these events, we have learned a great deal about operating an AG node and conducting a live event using AG technology. An operator's manual is being developed which encapsulates and codifies the practices we have learned.

Visualization and Access Grid

Distance Visualization

Just as the AG enables collaboration at a distance for meetings, workshops, tutorials, etc. it can support group analysis and interpretation of simulation or sensor output through collaborative visualization. Supporting this type of work requires that no assumptions be made on the capabilities of a given site beyond the availability of an AG node. If all sites were required to meet a certain visualization capability, this would severely limit the number of sites able to participate within a session. Therefore, it is important that visualization techniques be developed that are able to be executed at a distance via the standard AG environment. This provides those sites without advanced facilities the ability to still participate by relying on the capabilities of other sites. Research needs to be done in probing the capabilities of node in a particular session in order to optimize the appropriate visualization form. When determining the capabilities of a node, it is not only the computing power of the nodes that is important but also the infrastructure with which it is connected to the grid, requiring close examination of both bandwidth and latency. If the remote visualization overloads the machine to slow the network down, the use of the AG will suffer. Based on the probed results the remote visualization server may choose to either deliver the output in the form of images, if the capabilities are present it may be more appropriate deliver geometry and if even more capabilities are present it might mean the actual distribution of raw

data for processing on the AG node. While the visualization of the data is important, interaction with the data is also important to complete the analysis and interpretation. If the visualization is hindering the performance of the node it is extremely likely that interaction will also be affected and hence the experience. As discussed in the section on presence the overall goal of the system needs to be making the users feel collocated, therefore the visualization needs to feel responsive and local even though it is happening at a distance.

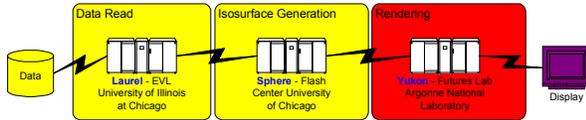


Figure 2. Distributed Visualization Pipeline

Teleimmersion

When talking of scientific visualization and collaborating at the same time one often means multiple people huddled around a desktop workstation pointing at the screen. What if one wanted to enter that world represented on the desktop? To explore the data from the inside out, expand it in size to surround the user, or shrink in down and stand above it, one would employ the use of a VR environment such as the CAVE[3]. Still, collaborating in this space would require bringing colleagues to the space to experience the

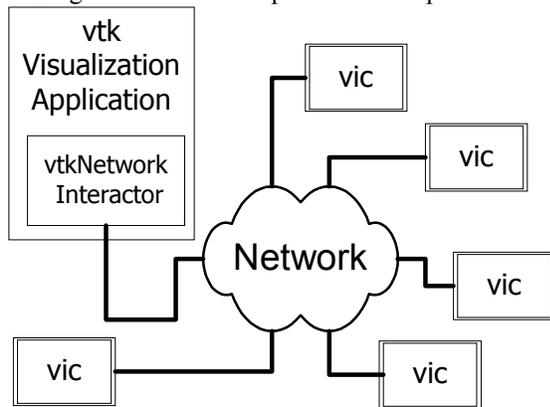


Figure 3. Remote rendering with VTK and Vic

Breaking the Pipeline

The first experiment presented the breaking of the traditional visualization pipeline: different pieces of the visualization pipeline were executed at different locations with the final output scan converted from

environment. The question we are really interested in is how to take the experience to collaborators. We have set out to construct a series of tools to enable the sharing of a VR space at a distance.

Experiments in Semi-Immersive Visualization

Much of the remote visualization work described here has been done as part of the Department of Energy’s (DOE) Next Generation Internet (NGI) Corridor One (C1) [2] project for the development of an integrated distance visualization environment for large-scale computer simulation applications. The C1 project is composed of six sites focusing on five major research activities:

- Grid Middleware and Advanced Networking
- Distributed Visualization and Data Manipulation
- Distributed Collaboration and Display Environments
- System Architecture, Software Frameworks and Tool Integration
- Applications Liaison, Experimental Design and Evaluation

As part of the C1 effort, we have undertaken a number of experiments in remote and distributed visualization in concert with the Access Grid activities, each probing different facets of the distance and remote visualization problems.

an SGI at Argonne for presentation on the AG. Figure 2 shows an overview diagram of the visualization pipeline for the experiment. The reading of the data was performed at the location the data, which in this case was the University of Illinois at Chicago. The data was then transported via an MPI send to the University of Chicago where an isosurface was created from the data. The polygons from the isosurface were then sent via an MPI send to ANL for final rendering. The rendering window at ANL allowed for interaction with the visualization, such as rotating, scaling and translating the final rendered object. Interaction was also allowed in the form of requests for new isosurfaces of different threshold values. In the case of this interaction, a request was made from the rendering stage back to the isosurfacing stage for the creation of a new polygon list, which was then sent back to the rendering stage for display. The output of the rendering stage was scan converted for output onto the AG. This work made use of the visualization toolkit (vtk) for the construction of the visualization

pipeline [20] and MPI for communication between nodes [21].

Remote Rendering with VTK and Vic

The second experiment involved the extension of vtk and vic for use in remote rendering. The vic application was modified to stream mouse and keyboard from a video window to a remote machine on the network. This stream was then received by a standard vtk application that was supplemented with a network-based interactor. This allowed multiple AG sites to interact with the remote vtk application using the standard vtk interactor controls via the vic window. The output of the vtk application was then scan converted from an SGI at ANL and transmitted to the AG as video. This application allowed for a remote user to have direct control of the application without concern for rendering power of the host machine. This allowed the user to interact in real-time with a one million-polygon model on a standard desktop pc. Figure 3 shows a diagram of the system. Figure 4 shows a snapshot of the remote rendering output, and Figure 5 shows the output as seen in the vic window.

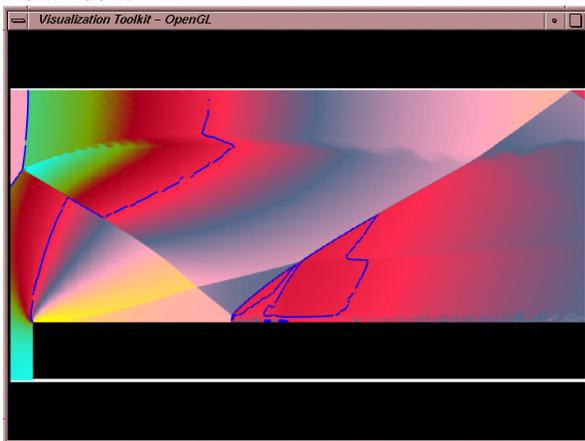


Figure 4 VTK visualization window

Immersive Displays and Synthetic Video

The third experiment used ANL's CAVEav suite of tools for the capture of synthetic video from advanced display devices such as the CAVE, ImmersaDesk, and activeMural tiled display systems. The CAVEav tools allow a developer to place within a standard CAVE application virtual cameras that stream video into the AG. This allows users who are unable to participate in an immersive session the ability to partake via the AG. The virtual cameras can

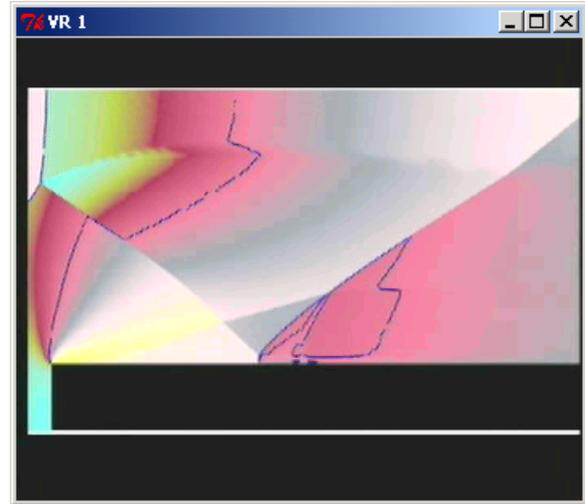


Figure 5 Remote vic display of Figure 4

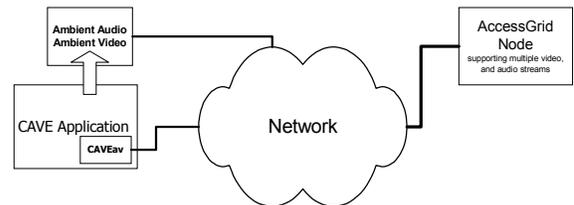


Figure 6 Architecture for synthetic video from CAVE Application

either be stationary within a given scene, attached to a given user, or remotely controlled via shared control-panel that resides outside the CAVE application [22]. Figure 6 depicts the structure and dataflow of this experiment.

This experiment showed the strongest connection to the advanced display environments by directly translating the output from these environments into a video stream for broadcast on the AG. The other experiments demonstrate technologies that can be applied to include more visualization power on the AG nodes than currently resides there. In addition,

the first two examples are currently being incorporated into advanced display environments and in turn connected back to the AG for the extension of these environments into a broader community.

What We've Learned

The connection of an immersive environment such as a CAVE, or even a semi-immersive environment such as a high-resolution tiled display to an AG node poses many interesting problems. In the case of participation in a collaborative CAVE session we have learned multiple issues need to be addressed such as audio, video, and representation. We have learned lessons that are applicable to both the AG spaces and immersive spaces when connecting them to visualization applications. Most CAVEs are designed to support audio feedback but not in a collaborative manner. This means that the CAVE environment needs to support an audio environment that mimics in many ways an AG node. This allows for hands-free full-duplex communication between users in the CAVE and remote users at an AG node. This allows for a very natural interaction between sites.

The user in the CAVE is surrounded by visual images that are for the most part computer generated. With the exception of the input device and fellow users within the CAVE, the physical space is very much part of a virtual space. In contrast, the users in an AG node are surrounded by the real world and interact with the space via standard laptops or terminals. The AG users view the virtual space through video streams. These streams are both real and synthetic. It has been found that the use of both types of streams makes the interaction with the users of the virtual space easier. This is because the real streams will often provide the AG users with a global context to the synthetic views they are receiving from the virtual cameras. A trickier problem is how to represent the remote AG users back to the users of the CAVE space.

Based on the experiments that we have performed to date on using a variety of different visualization applications within an AG setting including; immersive, semi-immersive, and standard desktop environments; it is important that a sense of presence is achieved. Early tests of just synthetic video from the CAVE or ActiveMural provided the AG user with information on what was happening in the space but did not provide context. The addition of standard ambient video streams of the CAVE or ActiveMural user began to provide the needed context for the remote users. Users are left with the "I'm at the

movies feeling" without the ability to interact with the remote CAVE users, therefore hands-free full-duplex audio is essential in the CAVE environment. Even given the inclusion of audio, the more you can plug the user into the remote environment the more strongly the sense of presence and contribution can be felt. By building tools like the vic based network interactor we have allowed the user to directly interact with the remote data. By extending that notion the user can remotely pan, tilt, and zoom the over the shoulder cameras, allowing the remote AG user the ability to control the information presented to them. Within the CAVE environments the virtual cameras can also be controlled in a similar manner. We have also learned the commodity hardware that makes the AG nodes affordable is incapable of fully supporting an entire visualization pipeline; some of the immersive environments are also incapable of supporting a complete pipeline; therefore it is extremely important to be able to break up the visualization pipeline and distribute only the information that can be accommodated on the host machine to it. While this is not directly related to bringing immersive environments to an AG node, it is an important point that for the environments to interact in a collaborative fashion. This reason holds equally true in the case of the immersive environments like the CAVE: if the host machine is overloaded the experience becomes frustrating, therefore this work directly carries over into the immersive arena.

Discussion and Future Work

Access Grid. The Access Grid has grown from two nodes in the same building to over twenty nodes across the U.S. in the last year, with plans to build more approaching one hundred this year. In that time, well over one hundred events have been held using Access Grid technology, giving good evidence of the soundness of the concept and the execution. Much technical progress has been accomplished so far, but this progress falls far short of the vision and much is still to be done in order to bring us closer to the vision. Hands-free full duplex audio support, once the most difficult aspect of the Access Grid, has now reached an acceptable state and we can address more advanced audio features such as spatialized audio, private audio channels and automatic audio setup and level balancing. There is considerable room for improvement in the way we use the current H.261 video codec and for the development of new codecs allowing more flexibility with respect to frame size, compression techniques and frame rate. It has been proposed that an open source MPEG codec be

developed for integration into the current VIC tools. As with audio levels, there are no tools available to monitor and control color balance and video quality across the sites.

Today, all streams are sent on the same network channel, all contending equally for access to the network resources. There is no way to designate some streams as being more important than others and to dynamically switch and re-route traffic. Several things must happen for this part of the vision to become a reality: QoS service must be deployed throughout the networks being used, and some kind of network flow engine must be developed to act on the AG application network requirements and dynamically reconfigure, reshape and switch routes based on AG and network policy requirements. In the vision, we imagine entering an AG space with personal devices, from cell phone and pager to laptop computers, and having the devices automatically integrated into the AG. We envision a method of exporting or importing content from these devices, either to or from other individuals, groups of individuals to the AG session participants as a whole. This implies much with respect to automatic device identification, network negotiation and formation and the development of protocols and methodologies for data and application transfer.

For structured events, a script is required. For an AG event, that script should be a tool to automate and synchronize each activity for each event across the network. The tool should be able to verify that required resources are in place at each site and marshal them if not. The tool should launch applications, start media players, remind the players, request and schedule resources and generally automate the performance of a network based structured meeting.

Immersive Visualization. In the computational community, visualization is a key component of collaboration. Integrating the various and sundry visualization environments into the Access Grid is a matter of primary urgency and requires imagination, skill and cooperation. In particular, we are interested in immersive visualization environments as discussed earlier.

We see a continuum of capabilities with respect to immersive visualization integration into the Access Grid. At one end is the semi-immersive environment described in our experiments, where both ambient and application video from multiple perspectives is presented to the AG and viewers have a sense of immersion, albeit once removed from the CAVE

dweller. At the other end we see an Access Grid node being run in stereo with tracking and with enough local rendering capability to faithfully render a scene as in a fully immersive environment. From the one to the other requires a corresponding increase in Access Grid capabilities. We are interested in experimenting at both ends, seeing what user environment improvements can be made with no additional hardware expense, and seeing what infrastructure requirements are necessary to achieve full peer status in immersive collaborative visualization. In between these two extremes are a number of technologies and experiments we are interested in:

- Controlling the immersive environment from passive video streams
- Exporting sufficient CAVE or immersive environment video to re-create an accurate, high fidelity version of the immersive experience on the Access Grid.
- Exporting control interfaces to the Access Grid
- Exporting a portion of the rendering pipeline to the Access Grid to allow active applications to run natively.
- Increasing Access Grid rendering power.
- Integrating Access Grid and collaborative immersive visualization with very large display wall technology such as the ActiveMural.
- Investigating non-coplanar displays for immersion using the Access Grid

Summary

The Access Grid has demonstrated successful distance collaborations and holds great promise. Many significant events are planned for 2000, including tutorials, lectures, workshops and even large scale distributed conferences – the 2000 Chautauquas to be held this summer. Many more sites than the current 20 are planning to build Access Grids this year and many events are in the planning. People are finding the Access Grid experience to be an effective antidote to the expensive and time-consuming process of building up their frequent flyer miles to visit with collaborators. Effectively integrating immersive visualization experiences into the Access Grid will further enhance its usability and attractiveness to those collaborators who don't have immersive environments and can't always travel to sites that do. We have already demonstrated a minimal capability with satisfactory results. We believe that further work and experiments in exporting key elements of immersive visualization experiences will produce a new immersive collaboration paradigm, widening the audience for

both collaborative and immersive visualization technologies.

References

- [1] Advanced Visualization Technology Center Website. <http://www.avtc.org/>.
- [2] "Data and Visualization Corridors," presented at 1998 DVC Workshop Series, DOE and NSF Sponsored, 1998.
- [3] C. Cruz-Neira, J. Leigh, M. E. Papka, C. Barnes, S. M. Cohen, R. Engelmann, R. Hudson, T. Roy, L. Siegel, C. Vasilakis, T. A. DeFanti, and D. J. Sandin, "Scientists in Wonderland: A Report on Visualization Applications in the CAVE Virtual Reality Environment," presented at IEEE 1993 Symposium on Research Frontiers in Virtual Reality, 1993.
- [4] J. Leigh, A. E. Johnson, M. Brown, D. J. Sandin, and T. A. DeFanti, "Visualization in teleimmersive environments," *Computer*, vol. 32, pp. 66-+, 1999.
- [5] M. Bogen, C. Bonkowski, R. Rodriguez-Val, and C. Wermelskirchen, "Desktop Video and Its Obstacles for Collaborative Work," presented at INET 97, Kuala Lumpur, 1997.
- [6] Access Grid Web Pages. <http://www-fp.mcs.anl.gov/fl/Access Grid/>.
- [7] 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," in *Advanced Multimedia Content Processing*, vol. 1554, *Lecture Notes in Computer Science*, 1999, pp. 30-43.
- [8] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications," IETF Audio-Video Transport Working Group RFC1889, January, 1996.
- [9] S. Mccanne and V. Jacobsen, "vic: A Flexible Framework for Packet Video," presented at ACM Multimedia '95, San Francisco, CA, 1995.
- [10] S. Varakliotis, O. Hodson, and V. Hardman, "A software platform for multi-way audio distribution over the Internet," presented at Audio and music technology: the challenge of creative DSP, IEEE Colloquium, London, 1998.
- [11] G. Mantovani and G. Riva, ""Real" presence: How different ontologies generate different criteria for presence, telepresence, and virtual presence," *Presence-Teleoperators and Virtual Environments*, vol. 8, pp. 540-550, 1999.
- [12] K. Chen, B. Nickless, and R. Olson, "The Multicast Beacon: User-level Verification of IP Multicast Group Reachability and Quality," *Submitted to Special Issue of Computer*

Communications on Integrating Multicast into the Internet, 2000.

- [13] D. J. Schiano, "Lessons from LambdaMOO: A social, text-based virtual environment," *Presence-Teleoperators and Virtual Environments*, vol. 8, pp. 127-139, 1999.
- [14] R. Evard, "Collaborative Networked Communication: MUDs as Systems Tools," presented at Seventh Systems Administration Conference (LISA VII), Monterey, CA, 1993.
- [15] T. L. Disz, R. Olson, and R. Stevens, "Performance model of the Argonne voyager multimedia server.," presented at IEEE International Conference on Application-specific Systems, Architectures, and Processors, Zurich, Switzerland, 1997.
- [16] T. L. Disz, I. Judson, R. Olson, and R. Stevens, "The Argonne voyager multimedia server," presented at IEEE Computer Society 6th High Performance Distributed Computing, Portland, OR, 1997.
- [17] G. A. Roediger and W. P. Lidinsky, "The multi-session bridge," *Computer Physics Communications*, vol. 110, pp. 149-154, 1998.
- [18] NCSA Alliance ACCESS Center. <http://www.ncsa.uiuc.edu/alliance/access-dc/>.
- [19] "Chautauquas Build Electronic Communities," in *NCSA Access*, September 21, 1999.
- [20] W. Schroeder, K. Martin, and B. Lorenson, *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 2nd ed: Prentice Hall PTR, 1998.
- [21] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI Portable Parallel Programming with the Message-Passing Interface*, Second ed. Cambridge, Massachusetts: MIT Press, 1999.
- [22] J. Bresnahan, J. Insley, and M. E. Papka, "Interacting with Scientific Visualizations: User-Interface Tools within Spatially Immersive Displays," Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL ANL/MCS-P789-0100, 2000.