

ActiveSpaces on the Grid: The Construction of Advanced Visualization and Interaction Environments

Lisa Childers¹, Terry Disz¹, Mark Hereld^{1,3}, Randy Hudson^{1,2}, Ivan Judson¹, Robert Olson¹, Michael E. Papka^{1,3,4}, Joe Paris¹, and Rick Stevens^{1,2,3,4}

¹ Futures Laboratory, Mathematics and Computer Science Division; Argonne National Laboratory; Argonne, IL 60439; USA

² ASCI FLASH Center; University of Chicago; Chicago, IL 60637; USA

³ Computation Institute; University of Chicago; Chicago, IL 60637; USA

⁴ Department of Computer Science; University of Chicago; Chicago, IL 60637; USA

Abstract. The Futures Lab group at Argonne National Laboratory and the University of Chicago are designing, building, and evaluating a new type of interactive computing environment that couples in a deep way the concepts of direct manipulation found in virtual reality with the richness and variety of interactive devices found in ubiquitous computing. This environment provides the interactivity and collaboration support of teleimmersive environments with the flexibility and availability of desktop collaboration tools. We call these environments ActiveSpaces. An ActiveSpace is a physical domain that has been augmented with multiscale multi-screen displays, environment-specific and device-specific sensors, body and object trackers, human-input and instrument-input interfaces, streaming audio and video capture devices, and force feedback devices—and has then been connected to other such spaces via the Grid.

1 Toward the Evolution of ActiveSpaces

The Futures Lab group at Argonne National Laboratory and the University of Chicago is developing prototype collaboration and visualization environments that we call *ActiveSpaces*. ActiveSpaces are the workspaces of the future, places that combine existing workspace infrastructure with high-tech information technology. The goal is to construct a workspace that enhances the work experience, enables the user to be more productive, and does not intimidate the user. The construction of ActiveSpaces is a cross-cutting research project that combines research in display technology, collaboration environments, networking, and many other areas into a seamless environment. Other groups are working on similar cross-cutting problems [6,10,14,15] or similar focused problems addressing a specific area [8,11].

The key point of ActiveSpaces is not the construction or integration of any one technology into the workspace, but the integration of many diverse

components. These components are individually developed to address a certain research goal or problem, and they are then incorporated into the larger picture of the workspace to become an ActiveSpace. Our premise is simple: in order for advanced visualization and collaboration technology to succeed, it must be widely used and used in a variety of application and interaction domains; and this wide use will happen only when the technology is empowering and the environments are compelling. The technology needs to be integrated into the places and modalities that characterize modern intellectual work.

Our experience has indicated that a considerable number of research activities require people to work in loosely coupled collections of small groups. This multiple-level group structure is the target of our technology work. We are interested in developing working environments that enable groups to visually and interactively investigate large scientific datasets using large-format and immersive visualization technologies in the context of shared collaborative spaces. These shared collaborative spaces are characterized by a number of attributes that distinguish them from both current desktop-oriented IP-based teleconferencing systems and from traditional low-bandwidth videoconferencing. They create the illusion of being in a shared workspace that is permanently connected to other workspaces. The illusion is supported by using multiple cameras, large-format displays, and full-duplex ambient audio, all of which enable natural conversations between participants as if they were in the same work room.

Our goal in this work is to understand how to engineer into future workspaces the technology that can support high-performance collaboration and scientific visualization. We believe that the following areas of work are needed to pursue this goal.

- Group workspaces need to be designed to be comfortable, flexible, attractive, and compelling environments that encourage users to congregate and naturally express themselves.
- Computing and communications technologies that are needed to support collaboration (e.g., cameras, displays, microphones, interfaces, and controls) should be highly integrated into the physical design of the space.
- The user interface should support a natural set of interaction modalities (for collaboration, this could mean hands-free full-duplex high-quality audio, multiperspective video streams; for visualization this could mean that 3D datasets should be experienced in 3D, with direct interaction interfaces of virtual reality devices).
- User-owned and user-managed resources should be easily and temporarily integratable into the active scope of the workspace. For example, a user's laptop, phone, and personal digital assistant [PDA] should include a simple mechanism that enables it to become part of the collective resource set for the time the user is participating in the activity of the workspace.
- Interconnecting collaboration and visualization provides a uniquely challenging set of requirements for the systems design, the physical and soft-

- ware integration, and the use of high-performance networking and communications for enabling shared interactions.
- Comprehensive sets of middleware services are needed to provide a high-level set of abstractions for multimodal communication, security, scheduling, and resource management on which the ActiveSpace environments can be layered.

Below we expand on each of these areas. We will then discuss the progress in our collaboration and advanced display technology research.

1.1 Creating Compelling Group Workspaces

A major challenge for designers of modern workspaces—rich in information technology—is to design something that people want to use. We believe that while fully immersive VR environments can be technologically exciting and highly attractive for occasional use, they do not (yet?) replace a well-designed physical space that has comfortable seating, excellent lighting, variable lines of sight, multiple work surfaces, and a flexible layout. One can create images of similar environments in a VR system, but there are many problems involved. Supporting multiple users in the same proximity is difficult. Many physical constraints rapidly tire the user, including lack of mobility, occluded views, the persistent use of low-resolution displays, and constant-focus display planes. In addition, the amount of computing and graphics power needed to provide just the ambient environment is often many times more than what is currently available and focused on the task (e.g., design model or scientific visualization dataset).

Thus we believe that we need to design spaces that work first as physical spaces; then, with sufficient introduction of information technology, they will become truly compelling spaces. We show in Fig. 1 and Fig. 2 how we have modified a small meeting room/library and a workshop room with communications technology that in our opinion creates a compelling space to work, particularly with the displays enabled. The implication of our analysis and experiments is that we need to return to the more basic thinking of how IT should be built into the workspace. It would not be a stretch to say that these ideas have been influenced by the Arts and Crafts style of interior architecture and the realization that display and communications technology is poised to leap off the desktop into the very fabric of our work and living spaces.

1.2 Integrating Technology into the Structures

Soon we will have a vast assortment of computing and communications device, form, and function factors to use in developing new environments. As a starting point for our work, we have made the following assumptions.

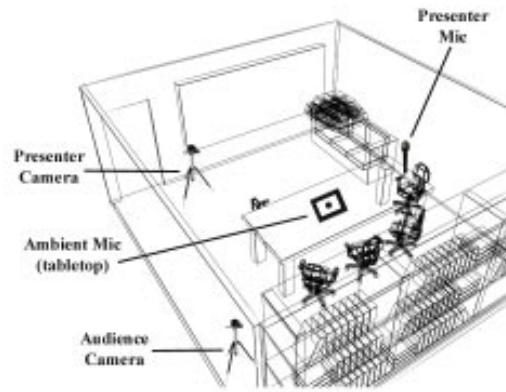


Fig. 1. AccessGrid node in a small meeting room setting.

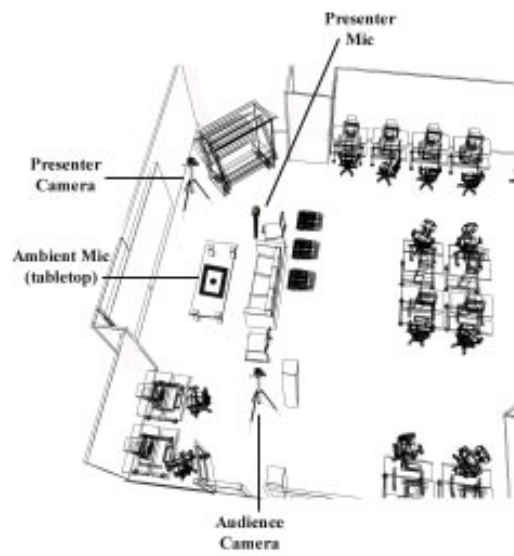


Fig. 2. AccessGrid node in a larger workspace setting.

- Within a few years, the functionality of the desktop PC will be available in a variety of physical form factors ranging from handheld and wearable devices to embedded systems.
- PC-based computer systems will become increasingly componentized, enabling one to assemble new types of ensemble and clustered devices that may share subsystems with desktop computers, but that will not be recognizable as desktop systems per se.
- New classes of devices (e.g., sensors, lights, controllers, switches, transducers) will become available as computer peripherals.
- Existing analog and digital peripherals will continue to decrease in price and increase in quality as chip technologies and economies are more fully used.
- Building, room, and personal-area networking will explode with new types of low-power and low-cost wireless capabilities.
- Moving off the desktop will create a huge new number of application opportunities, as well as the need for new user interaction mechanisms and organizing metaphors.

These six factors will drive computers into the environment in a way that is only partially understood at present. We certainly believe that rooms of all types will become one of the next application environments to be explored (along with automobiles, aircraft interiors, luggage, and the personal environments people carry with them). An essential point here is that we recognize that it will be possible to integrate substantial technology into the physical environment; the important questions are what type of technology it will be and how the integration will work. In our current and planned experiments, we are focusing on integrating cameras, projectors, microphones, speakers, screens, and tracking systems. However, we expect in the future to increase the number of environmental sensors and controls; we also expect to include lighting, seating, interactive work surfaces, and boundary (door and window) sensors. These devices will be built into the space and become a permanent part of the environment.

1.3 Natural Interaction Modalities

After the room and related structures become the application environment, it becomes immediately obvious that we need new user interfaces. Interaction modalities that these interfaces might use include voice, gesture via camera images, intentful motion via trackers, and handwriting or touch via embedded sensor mats or tracked pens.

In the collaboration application, natural interaction can mean a number of things. It might mean using full-duplex always-open voice channels to have a natural conversation with a remote participant. It might mean using perspective video to support a casual nonverbal visual cue. When these environments are used for scientific visualization, it can mean having a variety of

3D direct manipulation interfaces, each permitting the user to interact, manipulate, and modify the data without breaking the visualization metaphor in use. Workspaces of the future will include haptics, audio feedback, and audio commands.

1.4 User Resource Docking

It is very common today for users to have one or more personal computing and communication devices that they routinely take with them throughout the workplace throughout the workday. These devices include laptop computers, cellular telephones, two-way pagers, and PDAs. It is becoming clear that making all these tools available for use in a group activity would be an improvement over the current situation. A common example of this is a user who brings a presentation to a meeting on a laptop and needs to make it available to the group.

In more complex examples, the user's personal resources can augment the group environment by providing private channels of communication, or by providing local user interfaces to group- or room-oriented tools. Exploiting this capability requires software interfaces in the ActiveSpace that can "dock" with a personal resource, as well as common-room networking or personal-area networking (like Bluetooth) to provide the seamless physical interface for the duration of the session [4]. We are investigating the types of group applications that this capability would enable.

1.5 Collaboration and Visualization

Of particular interest to our group is the merging of collaboration technology with high-end scientific visualization systems. Part of this vision is the notion of people working together over large distances while engaged in close cooperative visualization of a large dataset that might be part of a computational science exploration or a complex design project. One approach to this problem is to understand how people use advanced visualization tools when they are physically co-located and to approximate that experience via networking by combining teleconference and tele-immersion capabilities with the visualization tools.

Another approach is to create a virtual shared space in which multiple types of visualization might be pursued simultaneously (e.g., immersive visualization and large-format visualization), and use collaboration technology to provide multiple foreground and background channels to link the groups together. A third approach is to focus on techniques that scale the human interaction beyond what is possible in the physical world (i.e., multicasting a visualization session to hundreds of people, giving them each a large degree of freedom to investigate the data, and synchronized to one or more lead sites). The goal here is to create an environment that is intrinsically more powerful than any single system. An example of this would be to link multiple displays

and groups of people that can explore a single dataset in multiple modalities simultaneously or at multiple space- and time-scales simultaneously.

These latter cases are our primary motivation; however, we recognize the need to support a wide variety of collaborative approaches to visualization. We are also interested in how collaboration technologies can be scaled. Currently our focus is on small group collaboration, with on the order of five connected sites, with up to six people per site. This appears to be a reasonable goal for the near future because in our experience, this is the most common scale of scientific and engineering collaborations.

1.6 ActiveSpace Middleware Needs

As we develop these systems, we have identified a need for a common set of system infrastructure capabilities. This set includes the following.

- Widespread and robust support of IP-based multicast, and tools for converting multicast to unicast and unicast to multicast.
- A common set of authentication tools that provide a one-time-per-session logon to the collection of ActiveSpace resources.
- High-performance communications libraries that support both best-effort service as well as emerging differentiated service.
- Resource discovery and management tools, which enable sites to advertise and discover resources that can be allocated as part of a distributed ActiveSpace session.
- Comprehensive scheduling tools that can integrate both human and non-human resource availability into coherently managed schedulable resource sets.
- Database-driven software distribution, configuration, updating, and monitoring capability that will enable a small group to manage a large number of distributed resources nearly automatically.

The importance of middleware cannot be overstated. It is critical to enable the rapid experimentation and development of ActiveSpace-like systems. By building on this layered structure, we not only have the ability to avoid overly specific implementations but also can leverage the development resources of a number of groups that are developing middleware-related tools and libraries.

1.7 Our Testbeds

The prototype environments we have built include a number of “AccessGrid Nodes” aimed at exploring group-oriented room-scale collaboration mechanisms. For visualization research we have developed software systems that use the CAVE and the ImmersaDesk, and we have been actively involved in the development of CAVE-related technology since 1992. In the past year, we have developed a large-format tiled display system we call the *ActiveMural* and a smaller version known at the *μMural*. We have initiated several

long-term projects to further develop the technology and software systems for these devices, and we have a number of collaborative projects that are aimed at networking together multiple sites that have similar environments [1,3,2].

2 The Grid

A new concept is shaping the way the research community is thinking about the evolution of the Internet. We call this concept the Computational Grid or simply the Grid. The name Grid comes from the analogy of the electric power grid, which provides a pervasive service via a complex network of providers, and yet has functional interoperability of devices and services. Perhaps we stretch the analogy a bit far, but the term Grid is short and works.

The concept of a Computational Grid has emerged in the university and national laboratory research communities in the past few years. Its origin can be traced to early (i.e., early in Internet time: 1991) metacomputing experiments, where multiple supercomputers were linked together to create a single large virtual machine that could be used to attack problems larger than any single machine could handle. In these early efforts to build distributed virtual machines, the supercomputer sites were interconnected by high-speed networks—the modern Internet precursor NSFnet being one of them. However, there were no additional capabilities available beyond the high-speed transport of data.

The early metacomputing experiments were sometimes successful, and they led to the thinking that interconnecting additional high-performance devices (e.g., virtual reality systems, telescopes, electron microscopes, terabyte-scale data archives) via common high-performance Internet-based networks would be even more interesting. After prototyping dozens of distributed high-performance applications on these ad hoc testbeds, it became apparent that many of the same software services and functions were needed by multiple applications. In each case, the application was having to solve the same problems; among these were authentication in a distributed multidomain environment, remote access to data without the benefit of a common name space, lack of network performance interfaces, and lack of high-performance wide area data transport interfaces. It soon became clear that these and some additional services formed a common collection of services that would benefit many applications. It is this layer of services that has formed the basis for the middleware layer of the Grid.

By deeply incorporating into the architecture this notion of common middleware services, the Grid represents a much more comprehensive view of one possible future direction for the Internet. Central to this view is the idea that the network can be the conduit of advanced applications delivery beyond the near-term evolution of the Web and desktop-to-server-based applications. In one type of Grid application, the user is presented with the capability of

harnessing a diverse set of computational, informational, collaborative, and possibly remotely controlled systems to build an application. These applications could range from simulation-based design of products, to distributed data mining, to inclusion of networked virtual reality demonstrations in a sales seminar.

The Computational Grid extends the current notion of the Internet to include a variety of integrated services, including computational servers, remote data services, collaboration infrastructure, remote instrumentation control, distributed computation, and tele-immersive visualization (distributed virtual reality). These high-level services are implemented on a common software substrate of middleware that provides application-independent services, such as data security, user authentication using public key infrastructure, distributed resource management, directory services, resource brokering services, and distributed resource scheduling.

The key advantage of the Grid is that through these advanced middleware services, applications that are much more powerful than today's Internet applications can be developed and broadly deployed. Applications developers will be able to develop new types of applications (e.g., distributed data mining with integrated collaborative information visualization) that can include high-performance access to data and devices in generalized ways. To build the Grid, we need to add new middleware capabilities to the existing suite of IETF protocols that support the Internet, and we need to begin to build large-scale testbeds that can validate the Grid concept on real applications with real users. Fortunately, this is happening. The federal Next Generation Internet program is funding a significant amount of work in this area, the university-based Internet2 initiative is building testbeds that support further exploration of the Grid concept, and the commercial sector is starting to view the Grid as a new modality of the Internet that may open up significant business applications and markets for advanced networking services.

3 Advanced Display Environments

An advanced display environment takes one beyond the traditional desktop display system, beyond a single monitor-based system. As part of an active space, the advanced display environments work to integrate the physical space, the user, and the computing environment. We have been investigating advanced display environments from two directions: immersive stereo-based environments, such as the CAVE and ImmersaDesk, and high-resolution tiled environments.

3.1 Immersive Stereo Environments

Immersive stereo displays, such as the CAVE and ImmersaDesk, provide a rich environment for the development of interactive 3D applications. The

immersive environments couple real-time tracking to the software-generated images for correct viewer-centered rendering of a given scene. The CAVE environment uses a single projector to display surface configuration, with a single desktop resolution projected onto a 10×10-foot surface. Our work in immersive environments has focused on developing immersive scientific visualization applications and on building these applications to work within the real-time requirements of VR [13]. We have also worked on understanding what value is added by the use of immersive environments and how people work within them.

3.2 Tiled Display Environments

Tiled display environments allow for the construction of arbitrary resolution surfaces. The display surfaces are constructed by tiling display devices together to form a single display. Several groups are working on the design and development of tiled displays that are being driven by a variety of different computing platforms [8,10,11]. We have constructed two such devices: the ActiveMural, a fifteen-projector system; and the μ Mural, a portable six-projector version. Both systems can be driven by either a high-end SGI machine or a Linux cluster.


ActiveMural. The ActiveMural (AM) is a rear-projected large-tiled display with a 16×8-feet visible projection surface (see Color Plate 1 on page 21). The current configuration is driven by either a Linux cluster, with each computer having a graphics accelerator card, or an SGI Onyx2 with eight Infinite Reality2 graphics cards, with two-channel output. Table 1 shows the various resolutions that the ActiveMural is capable of based on the current configuration. The screen material for the AM is JenMar Visual Systems BlackScreen, with a resolution greater than 200 lines per inch. The screen material is extremely good at the rejection of ambient light, allowing the AM to be used in normal room-lighting conditions, unlike the immersive environments described above. Interaction with the AM is currently limited to keyboard/mouse input; we have experimented with a 3D tracked joystick, which operates much like the CAVE wand, with mixed results. A promising input research path we have taken is the use of the Fujitsu Stylistic 2300, running WindowsNT and using a wireless networking card as a mobile interface to the AM.

μ Mural. The μ Mural is a smaller portable version of the AM described above. The μ Mural currently uses six projectors and can also be driven by either a Linux cluster or the SGI Onyx2. When being driven by the Onyx2, the μ Mural can use either one IR (Infinite Reality2) board with the eight-channel option or three IRs using the dual-channel option. Interaction with the μ Mural is the same as interaction with the AM. The major difference between the two, other than size, is the inclusion of hardware shadow masks

Table 1. Different Resolutions of ActiveMural and μ Mural.

Device	Compute Engine	Graphics Cards	Total Resolution
ActiveMural	Linux	15	5120×2304
ActiveMural	SGI	8	5120×2304
μ Mural	Linux	6	3072×1536
μ Mural	SGI	3	3072×1536
μ Mural	SGI	1	1948×1140

to help with edge blending. Figure 3 shows a head-on picture of the μ Mural using hardware-based edge blending.

**Fig. 3.** Image of a computer model rendered at full μ Mural resolution.

3.3 Visualization Systems

The advanced display environments discussed above are output devices with no knowledge of the Grid. It is the software infrastructure that makes these devices endpoints on the Grid. Using the advanced display systems, we are

constructing visualization systems to integrate the advanced display software environments onto the Grid and into ActiveSpaces. For immersive display systems, we built on the CAVE library [5], layering on top of it additional tools such as the Visualization Toolkit [12]. In addition, we have worked to enable the CAVE library to support more than eight display surfaces for use on tiled displays. Color Plate 2 on page 22 shows the same application running in both a CAVE and on the ActiveMural. The CAVE library has been ported to a variety of different compute platforms and provides a mechanism for rapidly deploying applications on a variety of different display environments.

In addition to developing tools for the CAVE library, we are constructing a set of open source tools and libraries for the tiled display environments. We have begun by building a set of tools for maintaining these environments, which includes alignment, color correction, and blending for the tiles.

4 Advanced Collaboration Spaces: AccessGrid

Today, the most common vision of computer-based collaboration tools is one of people sitting at their computer terminals, trying to look at their little cameras while also looking at a small, grainy video image and all the while saying, “Can you hear me? How do I run this software? Can you hear me?” This low image is the result of a lofty goal: to use the Internet to provide access to and collaborate with people in other places without having to travel.

Our own unsatisfactory experiences with desktop collaborative technology caused us to rethink 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, although we attend structured meetings, workshops, and conferences, we often tend to be more productive in an unstructured manner with a lot of brainstorming, problem solving, casual conversation, and spontaneous idea generation. From this, we realized the need to support multiple modes of interaction, from very structured to completely casual. Third, we usually have our computers with us and often want to share with others some information stored in our computer, be it a visualization, a spreadsheet, a presentation, a Web site, a document, or a movie. Finally, 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 gear and software.

We envision our ideal collaborative environment as an intentionally designed space: one that is rewarding to be in, and 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 software to enable the relatively transparent sharing of ideas, thoughts, experiments, applications, and conversation. We envision a space where we can “hang out” comfortably

with colleagues at other places, and also use that same space to attend and participate in site visits, remote conferences, tutorials, lectures, and other structured meetings. 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 a small group in a corner and caucus; and all the other things we take for granted when we are a group in the same physical place. In addition, we envision that the space will be “smart” enough to recognize that someone has brought personal computing resources to it and will allow the export of items from one computing device to other individuals or groups.

The challenges this vision presents are many and varied; some are easily addressed, while others will require groundbreaking research efforts.

In realizing the first AccessGrid, 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. First, this avoids forcing participants to purchase from and be slave to a particular vendor. Second, this allows every AG organization an equal chance to participate fully in research and development in AccessGrid technology.

4.1 Display

An AccessGrid 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. Second, 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 feet by 6 feet, with a seating area about 25 feet wide and about 20 feet deep. The projectors are mounted in the ceiling and are of sufficient brightness that the room can operate in normal light, allowing people to read and interact. We are experimenting with a corner display that uses four projectors, two per wall, to see if this enhances the sense of co-presence.

4.2 Video

An AccessGrid 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 four 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 the originating site sees), and a roving audience and room camera. We use remote control pan-tilt-zoom cameras for maximum flexibility. They are placed to be

unobtrusive and to facilitate the feeling of eye contact. We place cameras just below the projected screen area and place video images of people with whom we are conversing just above one of the cameras.

4.3 Audio

Being able to converse freely 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 within the space. We make sure there is adequate pickup everywhere in the room that there are likely to be people. We also use professional-quality echo cancellation gear by Gentner Corp. to ensure full-duplex audio. We currently place two speakers strategically in the front of the room to project good-quality audio into the space.

4.4 Computing

An AccessGrid Node uses four computers. The Display Computer runs windows NT and has a multiheaded video card. This is the machine that manages the tiled display and allows us to treat the multiple projectors as a single desktop. This machine is decoding all of the video streams, which can be several dozen, so it needs to be as robust as possible. The Video Capture Computer runs Linux and has four video capture cards. This machine encodes all the video streams captured at a node and then broadcasts those streams to the AccessGrid. It too must be a robust configuration to keep up with encoding demands placed on it. 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. 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.

4.5 Software

Aside from the operating system mentioned above, a compliant AG Node requires several pieces of software developed by AccessGrid partners.

A distributed PowerPoint master and server Software that allows one node to control the flow of PowerPoint presentations at all participating nodes.

A status beacon Software that runs on AG Nodes and regularly reports the status of the node to an AG database.

Reflector A piece of software that reflects network packets back to the node used for audio debugging.

- Vic** Network video capture and display software, originally written at LBL, modified and distributed by University College London, and further modified by AG developers.
- Rat** Network audio capture and playback software, written and distributed by University College London.
- Virtual Venue software** A room-based metaphor to control the scope of interaction on the AccessGrid. 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 with 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.
- Auto-layout** Software that automatically lays out windows across the screen real estate, based on preselected preferences. It eases managing the placement of windows on a large display.
- tkMOO** A text-based virtual space client used as a reliable communications back channel during live events and as a virtual community meeting place at all other times.

4.6 Network

The AccessGrid tools depend on network multicast to work well. Sites without multicast capability are forced to use some kind of traffic reflector. We use the Fermi National Accelerator Laboratory's MultiSession Bridge. Use of the bridge introduces delay, complexity, and significantly increases network load. Sites wishing to become AccessGrid 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 the originating one. The bandwidth required by each stream depends on settings at the origin and can vary from 128KB/s to 512KB/s or more. The effect of inadequate bandwidth on the AG Node is dropped packets, which results in unintelligible audio and jerky-motion video. Other effects can be detrimental performance for, and possible hostility from, other users on the local network.

4.7 Operations

The AccessGrid has been used in several major events in 1999: the Access DC grand opening event in April; the three Alliance Chautauquas in the summer; and SC99, where several sites brought Nodes to the 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. Figure 4 shows a diagram of a basic AccessGrid node.

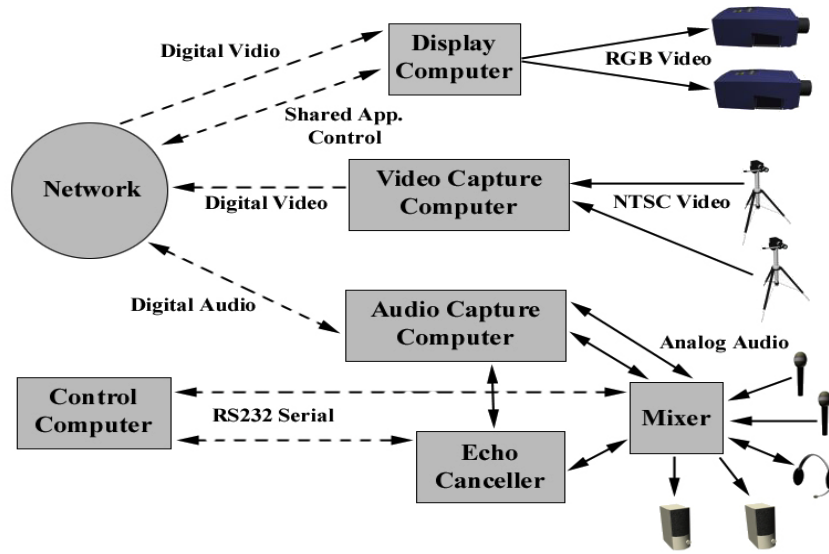


Fig. 4. Conceptual diagram of an AccessGrid node.

5 Conclusion

We have initiated several long-term projects [1,3,2] to develop the technology and software systems for ActiveSpaces. ActiveSpaces represent a convergence of collaboration, advanced visualization, and smart spaces concepts and technology. We have identified six trends in the development of information technology that will provide the driving forces for ActiveSpace environments. We are actively prototyping testbed systems to support exploration of the software, user interfaces, and applications that ActiveSpaces will enable. We firmly believe that our software must be designed and developed in a way that can leverage the emerging Grid Middleware infrastructure [7], and that these concepts need to be evaluated with real applications. Our applications partners include dozens of researchers from academia and national laboratories. Our goal is to release our prototype software to the community via open source software [9].

Acknowledgments

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.

References

1. Argonne National Laboratory. *AccessGrid*. <<http://www.mcs.anl.gov/fl/accessgrid/>>.
2. Argonne National Laboratory, Lawrence Berkeley National Laboratory, Electronic Visualization Laboratory, Los Alamos National Laboratory, Princeton University, University of Utah. *CorridorOne*. <<http://www.mcs.anl.gov/fl/corridorone/>>.
3. Argonne National Laboratory, Los Alamos National Laboratory, University of Utah. *Advanced Visualization Technology Center*. <<http://www.avtc.org/>>.
4. C. Bisdikian, S. Bouet, J. Inouye, R. Mettala, B. Miller, K. Morley, T. Muller, M. Roter, and E. Solboom. Bluetooth protocol architecture: Version 1.0. White Paper I.C.120/1.0, Bluetooth, August 1999. <<http://www.bluetooth.com/>>.
5. C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *SIGGRAPH '93 Annual Conference Proceedings*, pages 135–142, 1993.
6. T. L. Disz, M. E. Papka, and R. Stevens. Ubiworld: An environment integrating virtual reality, supercomputing, and design. In *Heterogeneous Computing Workshop Proceedings*. International Parallel Processing Symposium, April 1997.
7. I. Foster and C. Kesselman, editors. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1998.
8. G. Humphreys and P. Hanrahan. A distributed graphics system for large tiled displays. In *IEEE Visualization '99 Proceedings*, October 1999.
9. Mathematics and Computer Science Division, Argonne National Laboratory. *Futures Laboratory Website*. <<http://www.mcs.anl.gov/fl>>.
10. R. Raskar, M. S. Brown, R. Yang, W.-C. Chen, G. Welch, H. Towles, B. Seales, and H. Fuchs. Multi-projector displays using camera-based registration. In *Visualization '99 Proceedings*, October 1999.
11. R. Samanta, J. Zheng, T. Funkhouser, K. Li, and J. P. Singh. Load balancing for multi-projector rendering systems. In *Eurographics '99 Proceedings*, 1999.
12. W. Schroeder, K. Martin, and B. Lorensen. *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Prentice Hall PTR, 1998.
13. H. Tufo, P. Fischer, M. E. Papka, and K. Blom. Numerical simulation and immersive visualization of hairpin vortices. In *International Conference of High Performance Computing and Communications Proceedings*, Portland, Oregon, November 1999.
14. M. Weiser. The computer for the 21st century. *Scientific American*, 265(3), September 1991.
15. T. Winograd. Towards a human-centered interaction architecture. Unpublished draft available at <<http://graphics.stanford.EDU/projects/iwork/papers/humcent/>>, April 1999.

Presenters

Michael E. Papka


Virtual Environments Engineer

Mathematics and Computer Science (MCS) Division, Argonne National Laboratory, USA

Michael Papka received a B.S. degree in physics from Northern Illinois University in 1990, then continued master's studies in physics for another year while working as a research assistant at the Fermi National Accelerator Laboratory. He received an M.S. in electrical engineering and computer science in 1994 from the University of Illinois at Chicago. Since 1990, he has held various research assistantships within the Mathematics and Computer Science Division of Argonne National Laboratory <<http://www.mcs.anl.gov/>>. He is now a staff member at Argonne doing research in experimental virtual reality, collaborative science, and multimedia as applied to high-performance computing. He is also a Ph.D. student in the Department of Computer Science at the University of Chicago.

You can find out more about Michael E. Papka on the Web at <<http://www.mcs.anl.gov/~papka/>>.

Color Plates



Color Plate 1: Playback of a large-format movie on the ActiveMural. (See Sect. 3.2 on page 10.)



Color Plate 2: A CAVELib-based application running in the CAVE and on the ActiveMural. (See Sect. 3.3 on page 11.)