Middleware to Support Group to Group Collaboration

Rick Stevens, Justin Binns, Terry Disz, Mark Hereld, Ivan Judson,
Susanne Lefvert, Ti Leggett, Robert Olson, and Michael E. Papka
Argonne National Laboratory

Summary

The goal of the *Middleware to Support Group to Group Collaboration* (G2G) project is to address key middleware issues facing large-scale deployment of high-end collaborative environments. The research focuses on development of the scalable infrastructure and tools needed for group collaboration regardless of geographic location. The Access Grid is the testbed on which the developed infrastructure is evaluated.

The Access Grid (AG) is an ensemble of resources that supports group-to-group (3–20 people per site) interaction across the Grid. It consists of large-format multimedia displays, presentation and interactive software environments, and interfaces to Grid middleware and remote visualization environments. The Access Grid enables distributed meetings, collaborative teamwork sessions, seminars, lectures, tutorials and training. Large-format displays integrated with active meeting rooms are a central feature of AG nodes. The AG technology was developed by the Futures Laboratory at Argonne National Laboratory with support from the DOE 2000 program and is now used at over 150 institutions, including DOE national laboratories and NSF PACI universities.

The research effort is divided across five different areas.

- *Scalable Virtual Venues Service*: The Access Grid 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. We envision the AG peer-to-peer venue services operating much the same way as the Web: anyone can host a server (a virtual space), and anyone on the network can visit. Our goal is to create 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 order to gain widespread deployment, the AG implementation must address the security concerns of its users. At the least, the system should provide as much privacy that one expects from the telephone system. The AG security model will be based on a rigorous analysis of each AG component in terms of its vulnerability to attack.
- *Application Sharing and Dynamic Workspace Docking*: AG users often desire to "share" a portion of their personal workspace (desktop applications and data) with other AG users, nodes, or sites, both local and remote. The docking step amounts to 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 are developing a software layer that will improve node operations through simplified user interfaces, automated node configuration, and node

management functions. We will produce an AG node architecture document including specifications for AG hardware, software, and services.

- *Asynchronous Collaboration Capabilities*: We plan to extend the Voyager 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.

The Access Grid has demonstrated over the past three years the power of connecting geographically distributed teams. Groups that in the past met infrequently now use the Access Grid to meet on a regular basis. We expect to see an even greater increase in use as new software developed as part of the G2G effort is made available. This is in part due to new capabilities being added, including a much wider choice of machine configurations and seamless data sharing.

We have been working closely with the National Fusion Collaboratory to meet the needs of groups collaborating from within fusion control rooms with members of the experimental team that are not physically at the experiment. We have also been working with the Earth Systems Grid project to test applications that are layered on top of the AG for scientific visualization.

With SciDAC support supplemented by outside funding, we released version 2 of the base AG infrastructure. This new version enables AG participants to launch Grid-based jobs and to analyze and study results in a more inclusive manner. We will work with other SciDAC collaboratories and projects to encourage widespread use of the new infrastructure.

We plan to continue developing the base infrastructure as well as domain specific applications for the AG over the next two years. Improvements will include the integration of legacy applications into the shared environment, development of APIs for the construction of new shared applications and tools, and improvement in the reliability of the overall AG system. Reliability is especially challenging because the network infrastructure extends beyond a given location's control. Over the next two years we will also work to put automated recovering methods into the base AG software to enable meetings to continue even in the multicast failure, and without operator intervention.

We will also investigate higher-resolution video. As network capacity increases, the quality of the content between the distributed participants can increase, as can the amount of data that can be shared in real time between sites.

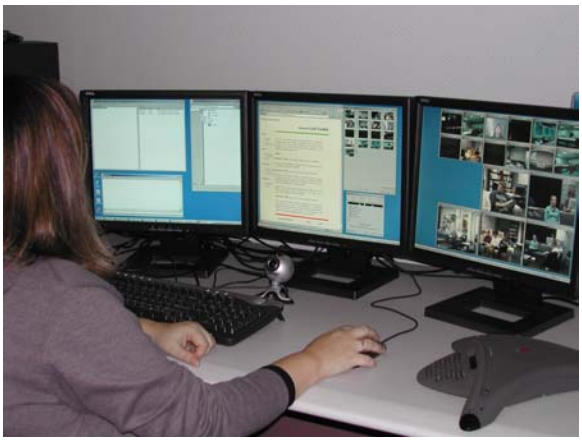

Figure 1: User working at a personal interface to the Grid running the new version 2.0 of the Access Grid software.

For further information on this subject contact:

Mary Anne Scott
Program Manager for Collaboratory Research
Phone: (301) 903-6368
Email: scott@er.doe.gov