

# USING THE ACCESS GRID AS A TESTBED FOR NETWORK MANAGEMENT RESEARCH

C.S. Hood, S. Devarapalli, N. Gadhia, S. Hegde, V. Mallikarjuna, S. Shankar,  
S. Yoginath

*Department of Computer Science  
Illinois Institute of Technology*

**Abstract:** There are many open issues in the field of network management. System data and/or operational experience can provide insight into these issues. Researchers often do not have access to such data or practical experience. To overcome this, we propose using the Access Grid as a testbed for network management research. The Access Grid is a multimedia collaboration environment that utilizes many different network technologies. We are using the Access Grid to collect data, gain operational experience and evaluate new management techniques.

**Key words:** Network management, testbed, monitoring, multimedia, dataset

## 1. INTRODUCTION

Network management is a very practical, yet complex and theoretically interesting problem. Understanding the problem requires exposure to the issues faced in managing a complex system. We can get insight into the issues through data analysis and operational experience. Toward this end, we are proposing the Access Grid (AG) as a testbed for network management research. The AG is a large-scale distributed collaboration environment where audio, video, and data is exchanged over high-performance networks using multicasting [2]. It is a research environment that is actively used for collaboration across geographic distances. Over the past year, we have installed our own AG node at IIT.

Our AG node primarily serves as an environment for fault management research. Data for this type of research is not readily available, so to begin with we focus on generating a comprehensive dataset. The performance requirements of the audio and video streams provide service expectations that can be used to label the data. In

addition to the data collection, we are creating simulation models to broaden the research capabilities. The goal is to put both the data and the simulation models in the public domain. This paper provides background about the access grid and describes our ongoing efforts in using the AG as a network management testbed.

## **2. ACCESS GRID**

The Access Grid is a collection of AG nodes interconnected by high-speed networks. There are currently over 100 AG nodes around the world [8]. An AG node is comprised of the computational resources and equipment necessary to send and receive audio, video, and data streams. Each node typically sends multiple video streams and a single audio stream. The AG facilitates both formal and informal group communication. There are eight virtual meeting rooms that can be utilized by participants. These virtual meeting rooms accommodate events ranging from large-scale distributed meetings to smaller informal collaborations.

Each node has one or more operators. The operator installs and maintains the node hardware and software. During an event, the operator troubleshoots in real-time. One of the tools most frequently used by the community is the AG mud, a text-based chat system. During an event, as problems occur the operators communicate to resolve the problems via the mud. All mud sessions are archived providing significant insight into fault, performance and configuration management issues. The established community of AG node operators and technical experts is willing to discuss and assist on topics ranging from operational issues to research problems. The resources of this community can be drawn upon to answer questions, arrange multi-node experiments and test solutions.

The AG community welcomes new members interested in deploying nodes. Documentation covering necessary equipment, software and installation instructions can be found at [3]. The cost of equipment for an AG node is approximately \$45,000.

## **3. MEASUREMENTS**

The AG spans a variety of interesting networking technologies involving LANs, WANs, multicast and multimedia. It is a large-scale networked system that allows us to study problems in a realistic manner. The ability to collect significant amounts of meaningful data is key. The data collected may include statistical measurements, events, configuration information, problem descriptions, and any other information that will provide insight into the system or network state. Additional information on the network infrastructure can be found through the Quilt [10] and Abilene Network and Operations Center website [1].

We focus on audio and video applications. Application performance is used to identify periods of degraded service. There is a measurement infrastructure in place to collect network, system, and application level measurements across the grid. This infrastructure takes advantage of existing measurements and tools to create a unified measurement log. Links to many of the measurement efforts can be found at [7]. CAIDA has comprised a list of monitoring tools at [4]. Using these tools, different types of information can be collected in different ways.

Measurements can be collected using active or passive methods. Active methods add synthetic load to the system or network and observe the resulting performance. Passive methods collect measurements resulting from existing load. Our measurement infrastructure utilizes both of these methods. We consider the infrastructure to be a starting point for data analysis. The goal is to determine the utility of different types of information under different circumstances. We are collecting

- SNMP Management Information Base (MIB) information from hosts and routers. MIB data contains both static and dynamic information
- End-to-end measurements including Round Trip Time (RTT), packet loss, one-way delay, out of order packets and duplicate packets
- Routing histories and changes
- System information regarding processor, memory, disks, network and other devices
- Application (vic and rat) performance measures. Vic and rat are the video and audio streaming applications on the AG
- Multicast beacon measurements [6].

We have been logging many of these measurements for several months. We are in the final stages of enabling measurement collection from Abilene and regional gigapops. Presently, we are getting readings on most of the measurements every fifteen seconds. As the utility of different measurements is better understood, the monitoring frequency can be adjusted. Additionally, since the AG software is in the public domain, we can add instrumentation if necessary. The goal of this measurement infrastructure is to provide a data set for studying the management problem, not for real-time monitoring.

## **4. SIMULATION**

To complement the “real” data, we also have a simulation of the AG. This allows us to simulate performance problems and test solutions. A simulation model of the AG can be an important tool for understanding how problems occur and propagate, as well as for testing the trigger models that are developed as part of this research. The simulation is particularly useful in terms of the network, because the opportunity for inserting faults and testing models is very limited there.

We simulate the AG [5] using the OPNET [9] network simulation tool. OPNET provides models for different types of equipment and protocols across many different technologies. The simulation project is broken into two pieces, (1) the AG node simulation, and (2) the network simulation.

The research described in [5] focuses on simulating the AG node. A typical AG node contains four machines; video capture machine, video display machine, audio machine and control machine. Our simulation models the first three machines. The primary function of the control computer is audio control. Since we have a good environment for studying the audio control problem directly, we have decided not to model the control machine. The three-machine model is sufficient to simulate the video and audio applications (vic and rat). Vic and rat use the Real-time Transport Protocol (RTP) to facilitate QoS. OPNET does not provide an RTP model, so this

effort primarily involved implementing RTP. In our AG node model, we embedded RTP within the audio and video application models.

The network modeling piece is straightforward, largely involving gathering the information about specific network nodes, topology, and configuration. For Abilene and several regional networks, this information is readily available. We are in the final stages of modeling Abilene. The simulation will be validated using some of the data collected as part of the measurement infrastructure.

## **5. SUMMARY**

The AG provides an opportunity to learn about management issues in a large-scale networked system. Implementing an AG node at IIT has been beneficial in many different ways. One key benefit is the collection of comprehensive, end-to-end labeled data. The goal of this project is to put the collected data into a public domain repository.

A second benefit is the operational experience. We have gained expertise installing the AG node and continue to face the challenges of operating the node. This first hand experience has helped us to identify several candidates for automation. Given the significant manual effort required to operate the nodes, automating pieces of configuration or troubleshooting can have big impact. The public domain node software enables us to deploy the mechanisms we develop. This is critical for evaluating the effectiveness of the mechanisms under realistic situations.

## **ACKNOWLEDGEMENTS**

The authors gratefully acknowledge support from NSF 013067 and NSF 9984811. This work was also supported in part by the U.S. Department of Energy, under Contract W-31-109-Eng-38.

## **REFERENCES**

- [1] <http://www.abilene.iu.edu/>
- [2] <http://www.accessgrid.org>
- [3] <http://www.accessgrid.org/agdp>
- [4] <http://www.caida.org/tools>
- [5] Chowdhury, Sadia, "Simulating the Access Grid," IIT MS Project, May 2002.
- [6] <http://dast.nlanr.net/Projects/Beacon/>
- [7] [http://e2epi.internet2.edu/tools\\_list.shtml](http://e2epi.internet2.edu/tools_list.shtml)
- [8] <http://www-fp.mcs.anl.gov/fl/accessgrid/ag-nodes.htm>
- [9] <http://www.opnet.com>
- [10] <http://www.thequilt.net>