Tools for Building Virtual Laboratories[a]

Deborah Agarwal, William E. Johnston, Stewart Loken, Brian Tierney

Ernest Orlando Lawrence Berkeley National Laboratory,
University of California, Berkeley, CA 94720, USA
http://www.lbl.gov/

There is increasing interest in making unique research facilities accessible on the Internet. Computer systems, scientific databases and experimental apparatus can be used by international collaborations of scientists using high-speed networks and advanced software tools to support collaboration. We arc building toots including video conferencing and electronic whiteboards that are being used to create examples of virtual laboratories. This paper describes two pilot projects which provide testbeds for the toots. The first is a virtual laboratory project providing remote access to LBNL's Advanced Light Source. The second is the Multidimensional Applications and Gigabit Internetwork Consortium (MAGIC) testbed which has been established to develop a very high-speed, wide-area network to deliver real. time data at gigabit-per-second rates.

1 Motivation

A collection of emerging technologies, ranging from distributed data handling and distributed computing to the underlying multicast network infrastructure, have the potential to make it possible to create distributed scientific and industrial laboratory environments that provide complete location independent access to instruments, data handling and analysis resources, and enable remote human collaboration. These technologies will produce not so much an incremental change of today's routine use of PCs and LANs in laboratory environments, but rather the introduction of a new paradigm more akin, perhaps, to virtual reality. The vision is for a scientist thousands of miles away to get the same (or even better) sense of presence and control as if he or she were at the experiment site.

The expected value of these geographically distributed environments includes substantially increased effectiveness in doing science by more closely integrating university and federal researchers, and an enabling capability for analytical and high-value production use of unique facilities by industry.

2 Introduction

The scientific community will benefit from advanced information infrastructure for several reasons. Breaking down the barriers of distance will permit routine and frequent par-

---

[a] This work is supported by the Director, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences (MICS) Division, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098 with the University of California. This document is LBNL report number LB L-37724. The authors may be reached at: wejohnston@lbl.gov, http://www-itg.lbl.gov, DAAgarwal@lbl.gov, BLTierney@lbl.gov, SCLoken@lbl.gov.

ticipation by senior academic and research scientists in planning and conducting experiments that would otherwise require scheduling long in advance, flying, housing arrangements, site safety and operations training, etc. These circumstances are completely typical for today's investigators who routinely have to visit large-scale facilities to do experiments and to meet with graduate students and post dots to plan new experiments.

Establishment of remote collaborator environments at unique or scarce facilities can also enable new capabilities for industry. This is because, in general, industry will not pack up their key personnel and go through the elaborate and time consuming procedures noted above. This is especially true of analytical and production use of remote facilities, as opposed to the few industry scientists who work at academic-like research labs. We therefore anticipate that establishing the technology for routine remote collaboration environments will add a new capability for use by industry.

The need for a distributed laboratory / scientific environment can occur by virtue of research and engineering involving the use of large and / or scarce facilities (e.g. large electron microscopes, synchrotrons light sources, various types of particle accelerators, etc.); by virtue of a single experiment being unique due to its scale (e.g. a fusion reactor); or (probably most commonly) because the collaborators and / or industrial partners are at several different institutions. We anticipate that a number of technologies will be necessary to solve the problem of providing location transparency in these environments. These technologies include high-speed distributed systems; use of network-of-workstations (NOW) -like computational system to do the theoretical calculations[b] use of shared networks for real-time, remote and automated control of instrumentation and sensing systems; routine and easily integrated telepresence (e.g. multimedia communication among all participants); and effective electronic publishing mechanisms to support intellectual communication during the scientific process and dissemination of results as they mature.

Research, development, and experimentation in testbed environments is necessary to extend, enhance, refine, and integrate the relevant technologies in ways that contribute to dramatic increases in the effectiveness of small and large scale scientific research collaborations where the scientific expertise, facilities, and instrumentation are geographically distributed. These testbeds environments need to identify, evaluate, and enhance the use of information, computing, and telecommunications technologies that will support routine "connection" among people and laboratory facilities that are scattered all over the country.

3 Establishing the Requirements for Virtual Laboratories

The Distributed Collaborator Experimental Environments (DCEE) Program of the E. O. Lawrence Berkeley National Laboratory and DOE, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences Division is intended to establish requirements and testbeds for Virtual Laboratories.

---

b. The U. C. Berkeley Network of Workstations project. http://now.cs.berkeley.edu/index.html

Five projects, beginning in early 1995 and running for approximately two years, will build and operate testbeds for "collaboratories"[c]. The testbeds will provide remote access to the kinds of expensive, hard-to-duplicate facilities, ranging from electron microscopes to a tokamak fusion reactor. By building these testbeds and using them for real-world experiments, we will be able to study both the technical and the social aspects of controlling apparatus, taking data, and interacting with colleagues by wire. The testbeds are:

- Electron microscopy, physics, and object-oriented virtual reality: Argonne National Laboratory. A shared work space with persistence and history will add a virtual-reality/multi-user dimensions touch to this project.
- Remote control for fusion experiments: Livermore, Princeton, Oak Ridge, and General Atomics. This testbed will attempt to replicate the function and feeling of real-time presence in a control room at a tokamak.
- Remote collaboration for environmental molecular sciences: Pacific Northwest Laboratories. The two parts of this testbed will serve, respectively, large, highly shared NMR facilities and smaller-scale molecular-beam instruments with small user groups.
- Remote SpectroMicroscopy at the Advanced Light Source (ALS)[d]: University of Wisconsin-Milwaukee and Lawrence Berkeley National Laboratory. Distinctive features of this project include a large and geographically distributed collaboration, which is part of a much broader (and rapidly expanding) user community in the field of synchrotrons radiation.

3.1 Requirements

Preliminary work at the ALS Beamline 7 (BL7) collaborator has established a set of requirements that are similar to those of the other virtual laboratory DCEE projects: remote experiment monitoring and control; on-line shared laboratory notebook; tele-presence; security for access control, safety, and data confidentiality; collaboration between multiple researchers at multiple sites; cross-platform compatibility (PC, Sun and Macintosh); maintainability; interface coherence and usability; control coordination; access from industrial sites.

Many of these points relate directly or indirectly to network capability and infrastructure, so we will characterize each briefly.

3.1.1 *Experiment Control*

A surprisingly small number of control system strategies are in use. National Instruments' LabView[e] is commonly used in small to medium scale systems (up to a few hundred monitoring and control points). Large-scale experiments (like the DIII-D Tokamak fusion reactor) have thousands to tens of thousands of control and monitoring points. These

---

[c] http: //www-itg.lbl.gov/DCEE

[d] http://beanie.lbl.gov:8001/als/als_homepage.html

[e] http: //www.natinst.com/labview/lvindex.html and http://k-whiner.pica.army.mil/info-labview/info-labview.html

experiments either use custom built control systems or EPICS[f] (Experimental Physics and Industrial Control System).

The ALS BL7 uses LabView to provide experiment control. We have modified this program to allow it to operate in either a client or a server mode. When in the server mode, the program provides control of the equipment at the beamline. In the client mode the user is simply monitoring the experiment parameters and data during the experiment. The user interface indicates which mode the user is in but otherwise is virtually the same. This has allowed us to develop a single version of the software for use at the remote sites as well as at the ALS BL7 itself. We are currently using both the UDP and TCP network protocols to provide communication between collaborating sites, but in the future we will be using reliable and unreliable multicasts for communication between the sites.

3.1.2 On-Line Experiment Notebook Requirements

An on-line notebook, especially for organizing and recording experiment parameters, data locations, etc. is a common need. This notebook must be shared in real-time by the collaborators and provide at least the capabilities available using a paper notebook.

In studying the ALS BL7 experiment notetaking methods we have been able to determine many of the characteristics needed in the on-line notebook. The on-line notebook must allow for both automated and manual entry of data along with an ability to reference other files. The support of cut and paste for images and pictures is also required since the analysis packages are not always shared by all of the collaborators. Other essential properties include: custom forms for formatted entry, sketching/drawing capability, support for scientific characters and symbols, secttrity/access restriction. The user interface must also allow search functions and easy interfacing to other software packages (Labview, statistics, graphics). The on-line notebook does not currently exist.

3.1.3 Multimedia Conferencing and Tele-presence

It is early in the human factors evaluation process, but it seems clear that rich multimedia inter-human communication and experiment tele-presence capabilities will be needed for successful remote operation and collaboration. Again, for the ALS, a specific set of capabilities has been identified, however it appears that variations of these will be common to all such environments.

Videoconferencing tools which use 1P multicast have been developed at LBNL. The “vat” tool sends and receives compressed audio and the “vie” tool sends and receives compressed video. These tools allow the user to set the desired transmission and compression rates and the specific compression algorithm to be used. A shared whiteboard (wb) has also been developed to allow sharing of screen images and a shared drawing space. In addition to this work, we have placed significant effort on improving the Internet transport and routing protocols.

Videoconferencing alone is not sufficient for the collaborator environment. Beamline 7 covers a large area consisting of the beamline itself and several computers and read-

---

[f] http://csg.lbl.gov/OverView.html

[g] http: //www-nrg.ee.lbl.gov/nrg/

outs. To provide the remote researcher with the ability to walk around the floor and view the various instrument settings, equipment and other researchers at the site we are developing and installing robotic cameras (remotely controlled) for area monitoring, fixed cameras at workstations, and data cameras monitoring the sample chamber. The various video sources can be switched using a remotely controlled video switch. The remote researcher will also have a video camera mounted by the remote workstation to transmit video of the researcher to the ALS and the other collaborators.

In order to provide natural movements during an experiment without losing audio contact with the remote collaborators, researchers at the ALS are equipped with wireless headsets. These headsets provide two-way communication between the researchers at the ALS and the remote collaborators. The background noise level at the ALS is relatively high but, with headsets, the researchers are able to carry on natural conversations. An audio mixer provides the means of joining the various researchers' voices together into one single conversation.

The software tools in use for the video and audio feeds are vie, vat, wb, CU-SeeMe, nevot, ivs. Remote robotic camera control is implemented in LabView. Both high and low bandwidth video channels are required, and all video and audio will be multicast. The robotic cameras and wireless communication systems are still under development.

3.1.4 Security/Safety

The advent of virtual laboratories brings a whole new class of user to the experiment. These users are likely to be much more occasional and less experienced with the equipment then has been the case in the past. The system will provide network based access to very expensive equipment and must be designed to avoid several potential security and safety problems. The system must be designed to have automated equipment failure modes with sanity checks on all incoming data and be resistant to network-based tampering.

The safety mechanisms are experiment specific and must be built directly into the server control software. Fail-safe recovery mechanisms are also required by the experiment server. With respect to remote users, one significant issue is to ensure that use-conditions of the remote resource have been met. Authentication and authorization certificates can be used in this environment to verify that some action (e.g. training) has taken place with respect to an individual. Facility managers may, for example, require certain kinds of training prior to allowing one to participate in remote operation. This manager will also probably never physically meet many of the people who want remote access. Assuming that there are training programs at the remote institutions that qualify an operator, one model is that this operator will "present" a security certificate as proof of identity in order to be allowed access by the local security mechanisms, and as proof of a level of training, or other authorization criteria, to engage in operations.

This scenario. if built. can provide one important component of the technology needed to create distributed scientific and industrial laboratory environments which provide complete location-independent access to instruments, data handling and analysis resources, and collaborators [1].

3.1.5 Communication Among Collaborators

The virtual laboratory environment consists of several components and most of these components run distributed at each of the collaborating researchers' sites and communicate with each other by passing messages. The virtual laboratory uses several different types of communication, including video conferencing, dissemination of experiment parameters and results, collaborator control information, security information, and notebook updates and transactions.

The different communication usages give rise to a multicast transport model with four types of message delivery service: unreliable unordered, unreliable ordered, reliable unordered and reliable ordered. Each of these message delivery types provides service characteristics that are useful to the virtual laboratory environment. For example, a status display can make use of unreliable unordered or reliable ordered delivery of messages. Parameter settings can be disseminated periodically to the collaborators using unreliable ordered delivery and points on a graph can be transmit using reliable unordered delivery. Experiment parameters and values that are sent only once will use reliable ordered message delivery.

The existing 1P multicast mechanisms provide unreliable unordered delivery of multicast messages. The other protocols must be developed from current research on reliable ordered multicast message delivery protocols.

3.1.6 Resource Arbitration

When many people at remote sites all have the possibility for controlling various aspects of an experiment, there has to be a robust mechanism for implementing whatever policy is established to permit multiple party access to the experiment. We are in the process of building a resource arbiter. The resource arbitrator/broker will be implemented as a completely distributed process and monitors resource status, arbitrates between sites requesting control of a resource, and monitors requestor status. The resource arbitrator must be scalable to an environment containing many resources and collaborators, and must provide flexible access restriction and arbitration mechanisms.

3.1.7 *Throughput Requirements*

The bandwidth required for communications to support virtual laboratories is fairly laboratory-specific in terms of the data that has to be transferred, but the multimedia streams for human collaboration and monitoring has a fair bit of commonality across different types of labs.

Considering specifically the spectro-microscopy experiment at BL7 of the ALS, the following parameters can be established. (Note that all numbers are in terms of application-level data.) The data transmission rate requirements at BL7 are between 8 -280 Kbits/second. The video and image transmission rate requirements are 2- 10Mbits/second. An additional video signal provided to monitor the sample chamber requires 2 Mbits/second. The telepresence videoconferencing transmission rate requirements are 128 Kbits/second for each compressed video stream.

4 Wide Area Network Experiments

The virtual laboratory scenario is a nice application for a high speed network **testbed:** The data flows are higher than current production networks could sustain (5 -15 **Mbits/s** of application data, 7-21 Mbits/see of network-level bandwidth for the ALS BL7 example, which is an unusually low data rate device); there are real-time (though as yet **uncharacterized**) requirements in the control system; multiple coordinated data and multimedia streams must be multicast (using several different models) to several users and sites that are widely spread across the country; and a versatile, strong, and widely distributed security infrastructure must be in place. These are **all** issues that have to be addressed in research, development, prototype deployment, and experiments in the environment of a high speed network where multicast strategies, QOS, bandwidth reservation, etc. can be experimented with in the wide-area network infrastructure.

4.1 High Speed Distributed Storage Architectures

A critical aspect of many laboratory experiment environments is the ability to collect and analyze high speed data streams. There are several aspects to this problem: The basic ability of systems to handle the bandwidths and volume of data involved; the mechanisms to bring computational capability to bear on the high speed data streams, and; the practical issues of system architectures that can be easily assembled and afforded by the scientific community.

The distributed-parallel storage architecture (**DPSS**) developed in the MAGIC gigabit network testbed[h] is intended to address these issues. The DPSS provides an economical, high performance, widely distributed, and highly scalable architecture for caching large amounts of data that can potentially be used by many different users and processes in laboratory and scientific collaborative environments. The current implementation of the DPSS technology is called the Image Server System ("1SS'')[4], and is optimized for provialing access to large, image-like, read-mostly data sets. In the MAGIC **testbed** the 1SS is distributed across several sites separated by more than 1000 Km of high speed network that uses 1P over ATM as the transport protocol, and stores very high resolution images of several geographic areas. In the Magic project, the **"TerraVision"** terrain visualization application uses the 1SS to allow a user explore / navigate a "real" landscape represented in **3D** by ortho-corrected, one meter images and digital elevation models [2]. **TerraVision** requests from the 1SS, in real-time, the sub-images ("tiles") needed to produce a view of the landscape. Typical use requires aggregated data streams of from 100 Mbits/see to 400 Mbits/see that are supplied from several servers on the network. Even in the current prototype system the 1SS is easily able to supply these data rates.

---

[h] MAGIC (Multidimensional Applications and Gigabit **Internetwork** Consortium) is a gigabit network **testbed that** was established in June t 992 by the U. S. Government's Advanced Research Projects Agency (ARPA)[3], and as **part** of DOE's high speed distributed computing program. The **testbed** is a collaboration between Mitre, LBL, Minnesota Supercomputer Center, SRI, Univ. of Kansas, Lawrence, KS, USGS - EROS Data Center, Sprint, Northern Telecom, U. S. West, Southwest Bell, and **Splitrock Telecom.** More information about MAGIC may be found on the WWW home page at: **http://www.magic.net/.**

The 1SS architecture is that of multiple network disk servers that are based on Unix workstations. The system coordinates multiple servers to aggregate high-bandwidth data streams to a network-based client application (e.g. TerraVision), as shown in Figure 1. Alternatively, many lower data rate streams can be supplied to many applications simultaneously (in a “video server” style of operation). The DPSS implementation uses an open systems, platform-independent, software approach. High performance is achieved in two ways: First, the functionalist y of the disk servers has been kept very simple - they are essentially “block” servers (a block being a fixed size unit of data, such as an image tile). Second, image data sets are easily partitioned over network distributed servers in a way that ensures parallel operation of many independent servers in order to supply a high bandwidth data stream to an application.

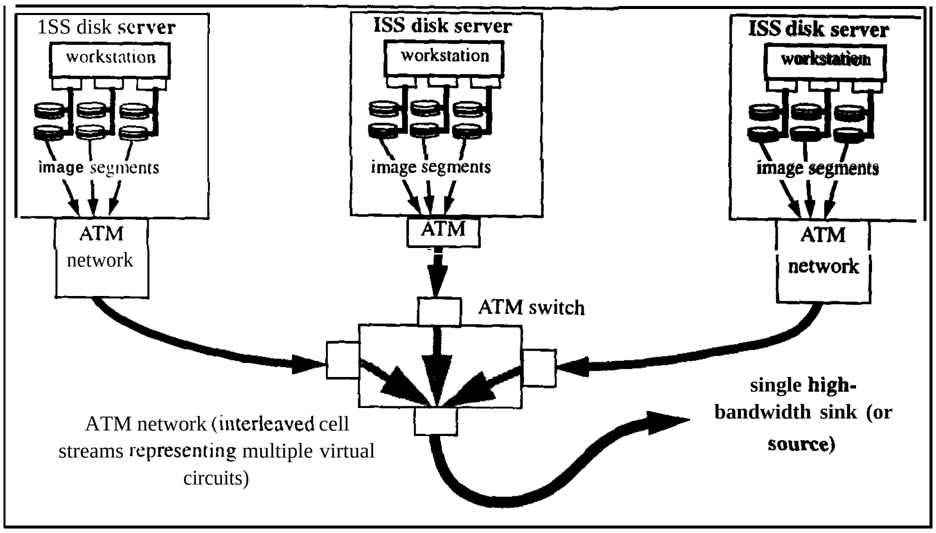

Figure I: Architecture for Distributed-Parallel Storage System (MAGIC testbed, Image Server System)

The DPSS technology potentially fits into the experimental science environment in various ways. For example, the DPSS might be used for buffering data coming from data collection systems prior to archiving, and it might be used as a large scale query results cache to support various sorts of data analysis. Second, DPSS technology also has potential use as an active element of a data analysis system for rapid reorganization of large volumes of data. It might be used as a cache for high speed in-line processing operations in conjunction with highly distributed processing.

5 Conclusion

The many recent advances in videoconferencing, collaborative software, and networking provide the backbone on which to create virtual laboratories. Large scale prototype virtual laboratories and related projects are currently under construction. These prototypes are

---

i. The most current description of the DPSS / 1SS technology is at http://www-itg.lbl.gov/ISS/papers.html identifying and developing the tools needed to successfully implement laboratories that allow remote control and monitoring of experiments by scientists. Emphasis is being placed on creating tools that are natural to use, secure, and safe. These tools will also allow the collaboration of many scientists working on a single experiment at disparate locations.

References

1. Johnston, W., Agarwal, D.,"Issues in Architectures and Technologies for On-line, Global-Scale Scientific Environments", presented at "America in the Age of Information: A Forum on Federal Information and Communications R&D," Bethesda, MD, July 6-7, 1995 (http://www.hpcc.gov/cic/forum/CIC_Cover.html). LBL report 37352.
2. Lau, S., Leclerc, Y. "TerraVision: a Terrain Visualization System,", Technical Note 540, SRI International, Menlo Park, CA, Mar. 1994. Also see: http://www.ai.sri.com/~magic/terravision.html
3. Richer, I. and Fuller, B. B., "An Overview of the MAGIC Project," M93BOOO0 173, The MITRE Corp., Bedford, MA, 1 Dec. 1993. (Available from http://www.magic.net/MAGIC_Summary.ps.)
4. Tierney, B., Johnston, W., Chen, L. T., Herzog, H., Hoo, G., Jin, G., Lee, J., "Using High Speed Networks to Enable Distributed Parallel Image Server Systems", Proceedings of Supercomputing '94, Nov. 1994, LBL-35437. Available from http://george.lbl.gov/ISS/papers.html.)