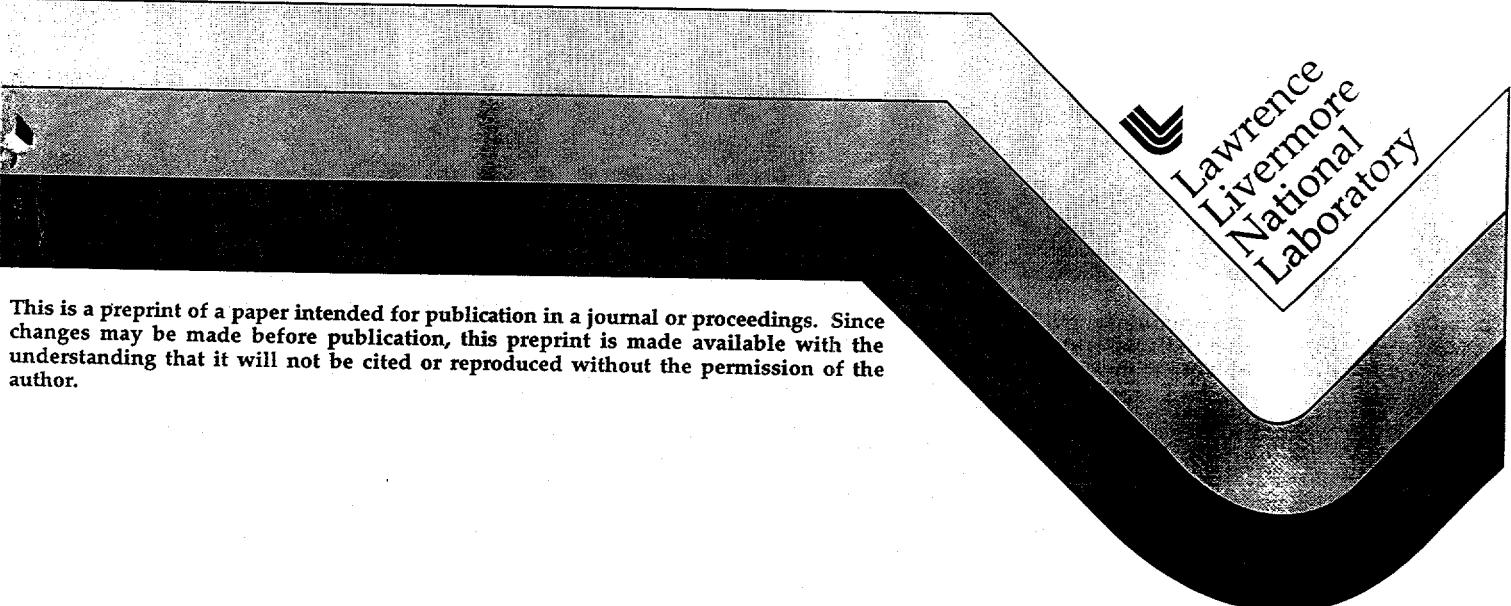


## Distributed Computing Testbed for a Remote Experimental Environment

D.N. Butner, T.A. Casper, B.C. Howard, P.A. Henline, S.L. Davis, D. Barnes,  
D.E. Greenwood

This paper was prepared for submittal to the  
16th Symposium on Fusion Engineering  
Champaign, Illinois  
September 30-October 5, 1995

September 18, 1995



This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

**MASTER**  
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED  
DLC

#### DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

#### **DISCLAIMER**

**Portions of this document may be illegible  
in electronic image products. Images are  
produced from the best available original  
document.**

# Distributed Computing Testbed for a Remote Experimental Environment\*

D. N. Butner<sup>1</sup>, T. A. Casper<sup>1</sup>, B. C. Howard<sup>1</sup>, P. A. Henline<sup>2</sup>, S. L. Davis<sup>3</sup>, D. Barnes<sup>3</sup>,  
D. E. Greenwood<sup>4</sup>

<sup>1</sup>Lawrence Livermore National Laboratory, P. O. Box 808, Livermore, CA 94550

<sup>2</sup>General Atomics, P. O. Box 85608, San Diego, CA 92186

<sup>3</sup>Princeton Plasma Physics Laboratory, P. O. Box 451, Princeton, NJ 08543

<sup>4</sup>Lockheed Martin Energy Systems, P. O. Box 2003, Oak Ridge TN 37831

## ABSTRACT

Collaboration is increasing in importance as physics research continues to become concentrated on a few large and expensive facilities. This is particularly true in magnetic fusion energy research, where the next generation experiments will be operated as facilities with national and international participation. In addition, these facilities are designed for steady state operation, for which interactive, real-time experimentation becomes an important issue. We are developing tools to provide for the establishment of geographically distant centers for interactive operations; such centers would allow scientists to participate in experiments from their home institutions.

We are developing a testbed for a Remote Experimental Environment (REE), a "Collaboratory". The testbed will be used to evaluate the ability of a remotely located group of scientists to conduct research on the DIII-D Tokamak at General Atomics (GA). The REE will serve as a testing environment for advanced control and collaboration concepts applicable to future experiments. Process-to-process communications over high speed wide area networks provide real-time synchronization and exchange of data among multiple computer networks, while the ability to conduct research is enhanced by adding audio/video communication capabilities. The Open Software Foundation's (OSF) Distributed Computing Environment (DCE) is being used to test concepts in distributed control, security, naming, remote procedure calls and distributed file access using the Distributed File Services (DFS). We are exploring the technology and sociology of remotely participating in the operation of a large scale experimental facility.

## INTRODUCTION

Modern research is often conducted on large and expensive experiments that utilize collaborative efforts on a national or international level to explore physics and engineering issues. This is particularly true for the current US magnetic fusion energy program, where collaboration on existing facilities has increased in importance and forms the basis for future efforts. As magnetic and inertial fusion energy research approach reactor conditions, the trend is towards fewer large and expensive experimental facilities, leaving many major

institutions without local experiments. The problem of integrating many experimental teams into the ongoing fusion program is an important one, since their expertise is quite valuable. To sustain a continued involvement in experiments, scientists are now often required to travel frequently, or to move their families to the new large facilities.

The currently proposed next-generation tokamaks, such as the Tokamak Physics Experiment (TPX) and the International Thermonuclear Experimental Reactor (ITER), will operate in steady-state or long pulse mode and produce fluxes of fusion reaction products sufficient to activate the surrounding structures. As a direct consequence, remote operation requiring robotics and video monitoring will become necessary, with only brief and limited access to the vessel area allowed. Even the on-site control room, data acquisition facilities, and work areas will be remotely located from the experiment, isolated by large biological barriers, and connected with fiber-optics[1]. Current planning for the ITER experiment includes a network of control room facilities to be located in countries from each of the four major international partners, the US, the Russian Federation, Japan, and the European Community.

Considerable work has been done over the last few years to promote effective participation in fusion research experiments by collaborators at distant locations[2][3]. Software tools have been developed and enhanced, and network capabilities have been extended in support of this activity. It is currently possible for researchers to control diagnostics, operate data acquisition systems, integrate high-capacity computing and storage facilities at the remote site, view and provide results, and generally have access to all the computer-based information about the progress of the experiment.

As an example, using internal funding at LLNL, a Remote Experimental Site (RES) was implemented that provides for remote interactive involvement in experiments[4]. This concept relies on many of the rapidly developing modern computational and communications technologies and has been made practical by the appearance of high performance workstations in a distributed computing environment. In particular, process-to-process communications over high speed wide area networks (WAN) provide real-time synchronization and exchange of data among multiple

computer networks. An internally developed asynchronous interprocess communication system (IPCS) provides a framework for application of WAN-distributed computations fully synchronized with experimental operations[5]. Information used for data analysis, display, and control system applications is routinely exchanged between geographically distant sites and the experiment, limited only by network latency times. Fast response is required, for example, for control of a subsystem or diagnostic, and for analysis and data display that may affect experimental operations. Although this concept has been developed and successfully used, as the number of collaborations grow, standardization, including security, becomes an issue that can be addressed with the newer distributed computing technologies.

The Distributed Computing Group at the National Energy Research Supercomputing Center (NERSC) works with customers to explore the possible benefits of new technologies to existing and planned scientific collaborations. The Distributed Computing Pilot Project at NERSC demonstrated the value of distributing applications between supercomputers and workstations, and of using distributed file systems. In particular, visualization tools running on high-performance workstations were shown to be an effective means of assimilating large amounts of data. Also, the distributed Andrew File System (AFS) is now being used more widely by collaborators in the fusion community to share data.

The Open Software Foundation's Distributed Computing Environment (OSF/DCE) software addresses the urgent need for application interoperability across major vendor platforms. It facilitates the creation, use and support of distributed applications and provides a transparent and efficient method for sharing resources. The OSF/DCE Distributed File Service (DFS) builds upon and enhances the features of AFS. The OSF/DCE and DFS software is now offered by a sufficient number of vendors to justify the effort of building a testbed for porting DIII-D applications.

While selected instruments are now routinely controlled by a remotely located diagnostician, the next step for active participation in experimental operations is to improve remote participation in operations in a major facility like the DIII-D tokamak. Considerable additional information associated with the control room environment is required to allow more complete participation in experimental operations. Sufficient audio and visual information will provide a remote presence in the control room by presenting crucial operations information while nurturing closer personal interaction between researchers at multiple sites. This provides the "look and feel" of local operations to remote groups by allowing views of the control room environment and integrating remote physics staff into discussions and meetings relevant to operation of the experiment.

To create a testbed, standardized DCE cells are being implemented at GA, LLNL, ORNL and PPPL. Rapid access to data at remote sites is being provided through optimization of WAN bandwidth usage, including data caching techniques within a distributed file structure. Wide area network access to major facilities requires a greater emphasis on connection security to protect against errors by researchers with only a limited knowledge of operations and against attacks by malicious intruders. Security is being addressed with the tools available through DCE. We are enhancing the collaboration environment by using advanced visualization techniques to implement a standardized graphical user interface into the data and operations. By integrating data communications and audio-video tools with distributed resources, we will provide a REE in which remote operations can be conducted.

## OBJECTIVES

We are developing a DCE based collaboratory for evaluating the ability of remotely located groups of scientists to conduct research on a world-class experimental facility such as the DIII-D tokamak. The configuration will be sufficiently advanced to support real-time operations and serve as a prototype from which future facilities, dependent on remote operations, could be designed, for example, TPX and ITER. Use of standards based technologies will result in a testbed that could be used to interact with other facilities connected to Internet, either within the US or internationally. This requires utilization of remote LAN's at all participating sites, fully synchronized across the WAN (ESnet) with the real-time experimental operations, and sufficient audio/video communications to make the remote researchers feel they are functionally present in the control room.

This collaboratory offers an excellent opportunity for testing an OSF/DCE system. The OSF/DCE environment provides:

- a secure, open, and distributed operating environment
- a standard naming service
- a standard remote procedure call
- an enhanced distributed file service with access from all sites
- a hierarchy of shared data files with optional automated replication

In present applications where remote computational resources are applied to an experimental operations environment, it is necessary to use an asynchronous message/data passing system to provide transparent and vendor-independent synchronization of tasks running on the WAN. Message passing is being implemented by developing an Asynchronous Message System (AMS) in the DCE environment.

Distributed data management, data caching, security, and visualization within the DCE environment are also being evaluated and addressed. Additional objectives of the OSF/DCE test bed are:

- Test conversion tools that will be used for the eventual AFS to DFS transition.
- Test AFS client interoperability with DFS servers.
- Join with other DCE cell sites, such as OSF, to test inter-cell operation.
- Evaluate DCE administration tools (utilities and scripts) as they become available in public domain and commercially.
- Demonstrate the extensibility of DCE to multiple, widely distributed sites (GA, LLNL/MFE, ORNL, PPPL) integrated into the DIII-D operations.

Another objective of this collaboratory is to provide a demonstration of distributed tools in the DCE testbed. To accomplish this, a distributed scheduling system, a graphics distribution and subscription system, and distributed visualization tools will be demonstrated in the environment. The systems are being built upon AMS and rely on the DCE distributed file system (AFS/DFS). Using DCE tools, methods for remote sites to synchronize with DIII-D real-time events and other status and operational information are being developed. A distributed data environment for accessing DIII-D experimental data is being developed for the DCE testbed.

The distributed scheduling system will provide the ability for various tasks to be aware of the status of other tasks, and provide the ability for logical task dependencies to be specified. In our current fusion efforts, both in relation to local and remote experimentation, enabling such task and data dependencies is an important requirement[6][7]. The graphics distribution system will provide a mechanism for distributing a variety of standardized preformatted graphics throughout the experimental environment. Experimenters will be able to provide their graphics output to the system, allowing other interested collaborators to subscribe to the graphic.

The process used for analyzing experimental data is currently tedious and labor intensive, requiring the scientist to move the output of one code to another platform as input to the next code. We are exploring techniques for extracting relevant data from large data sets, such as those created during DIII-D operations, and delivering it to one of several possible visualization tools in the required format. An objective is to provide a data interface as a pipeline to automate the analysis process while enabling the analyst to display data.

These systems are viewed not only as a direct demonstration of DCE features, but also as tools on which other demonstration applications will rely. These tools are expected to have applicability beyond just the fusion community.

A final objective of this collaboratory is to contribute to the ongoing experimental operations by processing data on multiple remote workstations, and controlling selected instruments and subsystems in collaboration with the DIII-D

group. Performance of a security standard, such as Kerberos, in the real-time operations and data processing environment will be tested. The use of network-based audio/video tools will be demonstrated and their utility in meeting our operational objectives will be evaluated. A review of both the technical and sociological aspects of conducting operations from a remote location will be provided.

## TECHNICAL APPROACH AND PLAN

An interactive data and control link has been implemented between LLNL and the DIII-D tokamak at GA which is in use during experimental operations associated with an existing GA/LLNL collaboration. This link represents a prototype for the proposed collaborative environment and has shown the importance of having the support and commitment of the local group for providing systems and operational support. Utilizing this link during current experiments, issues related to performance associated with data acquisition, data analysis, data display, and control have been explored, while addressing some of the sociological issues inherent with such remote interactions.

The fundamental design concept for the REE is multiple computer LAN's connected over a WAN and synchronized by a process-to-process asynchronous message transfer technique implemented with DCE standards. In the current implementation at LLNL, the LAN includes the existing RES, the office automation network, compute servers in the User Service Center, and links to NERSC. We are enhancing and generalizing this concept to include locations at LLNL, ORNL, and PPPL by adding the video/audio communications needed to support remote operations and by using the DCE standards. We are improving and generalizing our data caching based on a distributed file system and adding security features. We will provide a graphical interface based on industry standards, a distributed scheduling system and a graphics distribution system. These enhancements will allow full participation in the operation of the DIII-D tokamak by scientific staff located at remote sites.

## CURRENT STATUS

We have begun the initial development of the DIII-D data handling model using AFS clients at all sites and selected data stored on the AFS server at NERSC. With the help of OSF, we have installed and configured a DCE cell at NERSC and are bringing up DFS. NERSC is administering this cell. Access is provided via the ESnet ATM service to this DCE cell from clients installed at GA, ORNL and PPPL. DFS will be provided at a later date.

The first version of AMS has been developed under DCE and is undergoing testing. It provides for both authenticated and unauthenticated asynchronous message passing among distributed processes.

We are conducting an investigation into the application of audio/video (A/V) tools to support our remote operations environment. We are evaluating the suitability of existing Internet-based A/V tools to provide effective telepresence:

- Broadcasting of status and operations meetings with extensibility to support interactive participation
- Desktop conferencing between participating scientists including A/V communication and shared "whiteboard" applications.
- Remote monitoring of the experiment and the control room environment
- Access to and control of the remote monitoring from multiple sites

Rather than developing our own methods, we are concentrating on the availability and suitability of existing A/V methods using both public domain and commercially available software as needed to demonstrate and evaluate functionality and performance.

### FUTURE DEVELOPMENT

While investigating the use of visualization tools as an effective means of assimilating large amounts of data, we found that incompatible file formats often prevented the scientists from using the graphics tool of choice. Design of a data interface can alleviate this difficulty while also expanding the choice of tools for viewing the data, such as AVS, Collage, XImage, IRIS Explorer, PV-Wave, and IDL. We will explore the following capabilities:

- Browse the contents of files written in a variety of formats
- Randomly select data variables
- Rearrange variables into the desired form
- Modify the value of certain variable attributes
- Save variables into new files
- Feed variables directly to a visualization environment.

The distributed scheduling and graphics distribution systems will demonstrate operations in the DCE environment, enabling collaborators to automatically run analysis programs and to communicate graphics results among colleagues.

With the addition of audio/video communication, distributed file data caching, security, and visualization, we will be properly positioned for implementing and enhancing the remote operation of selected subsystems of a major facility such as the DIII-D tokamak, where a significant established collaboration consisting of on-site and off-site researchers has already been established. This standards-based testbed will also provide an ideal facility for testing concepts being planned for remote control of the future fusion experiments.

We will have the opportunity to investigate the technological and sociological aspects of scientific research using a widely distributed work group.

### ACKNOWLEDGMENTS

\*Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-Eng-48, by General Atomics under contract number DE-AC03-89ER51114, by Princeton Plasma Physics Laboratory under contract DE-AC02-76-CHO3073, and by Lockheed Martin Energy Systems under contract DE-AC05-84OR21400. This work is funded by the Distributed Collaboratory Experimental Environments Initiative at the Lawrence Berkeley National Laboratory.

### REFERENCES

- [1] S. Davis, P. Funk, G. Oliaro, W. Rauch, W. Stark, D. Butner, and T. Casper, "Conceptual Design for the TPX Central Instrumentation and Control System," *Proceedings of the 15th IEEE/NPSS Symposium on Fusion Engineering*, Oct 11-15, 1993, pp. 565-568.
- [2] D. Barnes, S. Davis, P. Roney, P. Henline, T. Casper, M. Greenwald, "Collaborations in Fusion Research, A Fusion Computing Council White Paper," Sept. 1994, unpublished.
- [3] R. Fonck et al., "Remote operation of the TFTR BES experiment from an off-site location," *Rev. Sci. Instrum.* 63 (10) Oct. 1992.
- [4] W.H. Meyer, D.N. Butner, and T.A. Casper, "Remote Experimental Site (RES) Release 1.0," *Bull. Am. Phys. Soc.* 38,2112, 1993
- [5] G.G. Preckshot and D.N. Butner, "A Simple and Efficient Interprocess Communication System for Actually Using a Laboratory Computer Network," *5th Conference on Real-Time Computer Applications in Nuclear Particle and Plasma Physics*, San Francisco, CA., 1987.
- [6] M. Thompson, T. Carroll, T. Gibney, J. Snyder, "Real-Time Remote Access to TFTR Experimental Data," *Proceedings of the 15th IEEE/NPSS Symposium on Fusion Engineering*, Oct 11-15, 1993, pp. 142-145.
- [7] B. B. McHarg, Jr., "Access to DIII-D data located in multiple files and multiple locations", *Proceedings of the 15th IEEE/NPSS Symposium on Fusion Engineering*, Oct 11-15, 1993, pp. 123-126.