Remote Experimental Environment: Building a Collaboratory for Fusion Research

T. A. Casper, B. Howard,
R. Hunt, W. H. Meyer,
J. M. Moller, B. McHarg,
P. Henline, K. Keith,
K. Greene, S. Davis, D.
Barnes, T. Gibney,
P. Roney, and D. Greenwood

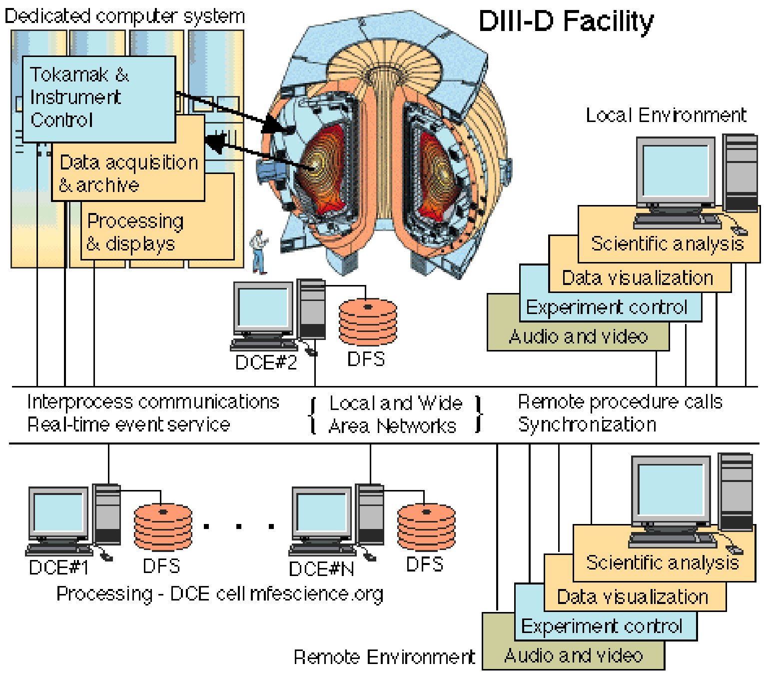

Figure 1. DIII-D distributed-computation environment. Interprocess communications, remote procedure calls, and event services provide synchronization and scheduling for real-time operations, scientific analysis, and data visualization. Integrated audio and video provide interactive communications.

Collaboration is increasing in importance as physics research continues to focus on fewer, larger, and more expensive projects. Experiments in the U.S. research program on magnetic-fusion energy are currently operated as national projects using teams of scientists from many institutions, most with international representation. A next-generation experiment is being designed as an international effort and will be operated with worldwide involvement. Future facilities will support continuous operation for which interactive, real-time experimentation becomes an important issue. To minimize the need for relocation and travel by researchers and their families and to sustain scientists' continued active involvement from their home institutions, we have been exploring techniques for interactive remote participation in experiments.[1,2]

High-performance wide-area networks and powerful workstations are helping us to create a distributed computing and information environment. In our approach, process-to-process communications over high-speed wide-area networks provide real-time synchronization and exchange of data among multiple computer networks. Considerable additional information associated with a control-room environment is also made available to the off-site collaborators, so that they can be integrated into experimental operations. Shared audio and visual environments help to nurture close personal interaction among researchers at multiple sites. This sort of organization of a research project is often referred to as a "collaboratory."[3]

Background

Magnetic-fusion-energy research involves the scientific exploration of methods for confining and heating an ionized gas, or plasma, to temperatures in the range of tens of millions of degrees for energy-recycling times of a few seconds. In a plasma, isotopes of hydrogen, deuterium and tritium, undergo fusion reactions in which these ions combine to form helium with the release of energy in the form of neutrons. These reactions are the same that fuel our sun and the stars and can in principle be used to generate a nearly limitless supply of energy for future generations. In addition to fusion, the rele-

---

T. A. Casper, B. Howard, R. Hunt, W. H. Meyer, and J. M. Moller are at Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, CA 94550. E-mail: casper1@llnl.gov
B. McHarg, P. Henline, K. Keith, and K. Greene are at General Atomics, 3550 General Atomics Court, San Diego, CA 92121. E-mail: mcharg@gav.gat.com
S. Davis, D. Barnes, T. Gibney, and P. Roney are at Princeton Plasma Physics Laboratory, James Forrestal Campus, Princeton, NJ 08543. E-mail: davis@pppl.gov
D. Greenwood is at Oak Ridge National Laboratory, Bethel Valley Road, Oak Ridge, TN 37831. E-mail: greenwoodde@ornl.gov

S0894-1866(98)00303-4

vance of plasma science also extends to interstellar matter and astrophysical research, to plasma-processing technology used in the semiconductor industry, and even to the potential optimization of engines required for interstellar travel.

Experimental research on plasmas having conditions comparable to those present in the sun is conducted in large, toroidally shaped devices called tokamaks. Strong magnetic fields several thousand times as intense as the Earth's field are generated and shaped with electromagnets to create a confining force effectively replacing the huge gravitational attraction in the sun. Millions of amperes of electrical current heat the plasma to nearly the temperature required for fusion and create an additional magnetic field needed for plasma confinement. To reach the final reacting temperature of nearly 100 million degrees, several million watts of power for additional heating is supplied by beams of accelerated neutral particles or by electromagnetic radiation at radio and microwave frequencies. Currently operating experiments exhibit density and temperature conditions relevant to that needed for fusion reactors.[4]

Like many other types of physics-research equipment, such as telescopes used in astronomical observations, accelerators for probing the nuclear structure of matter, and vehicles for space exploration, tokamaks capable of reproducing solar reactions are expensive to build and maintain. Such tokamaks are few in number, and their use must be optimized if we are to make progress in the science of fusion plasmas. Research on these tokamak experiments is conducted by teams of scientific investigators from laboratories and universities worldwide. Experiments on tokamak-confined plasmas are interactive in nature, with multiple groups simultaneously investigating different aspects of the associated physical processes. The national and international programmatic involvement, combined with the interactive experimental environment, has led to the exploration of techniques for supporting enhanced participation in experiments from geographically distributed locations.

Our project, known as the Remote Experimental Environment, is a response to such needs. The Remote Experimental Environment relies on three major capabilities:

- Distributed utilization of data, resources, and people.
- Robust, interactive communications and information exchange.
- Control of instrumentation and experimental systems from remote locations.

These requirements have all been addressed to some level in the Remote Experimental Environment. Our objective was to provide a testbed for exploring remote-collaboration research, essentially building an initial "collaboratory" for national fusion research. The testbed is being used to evaluate the ability of remotely located groups of scientists to conduct research on a world-class experimental facility, the DIII-D tokamak[5] at the General Atomics Corp. The Remote Experimental Environment serves as a testing environment for advanced computing, control, and collaboration concepts applicable to future experiments while enhancing the ability to conduct research on the existing facility. Participants in developing the Remote Experimental Environment likewise functioned as a distributed workgroup and used the collaborative tools during the development process.

Wide-area DCE testbed

A wide-area testbed for distributed computing provides researchers with seamless access to computing services irrespective of their physical location. It uses the Open Software Foundation's Distributed Computing Environment (DCE) to develop and test concepts in distributed analysis, visualization, and control. DCE security and naming services are used with global, distributed file access from the distributed file system (DFS). An initial DCE cell was installed and operated with cell-management services provided by the Lawrence Livermore National Laboratory (LLNL) Livermore Computing Center. This

A large magnetic-fusion research program uses the Internet for team communications, control of experiments, and the analysis of experimental data.

prototype fusion DCE cell, a portion of the larger LLNL computing environment, formed a nationwide distributed-computing "cluster" with workstations located at LLNL (northern California), General Atomics (southern California), Oak Ridge National Laboratory (ORNL, Tennessee), and the Princeton Plasma Physics Laboratory (PPPL, New Jersey).

After gaining expertise in cell management and operation, we have now formed a separate magnetic-fusion cell (mfescience.org) managed by LLNL. During development, the cell was configured with three Hewlett-Packard (HP) and four Silicon Graphics Inc. (SGI) workstations. Recently we added four more HP workstations and an additional SGI, and we will soon add ten Sun Ultra Sparcstations. Figure 1 shows a diagram of the computing environment available to researchers in the experimental program. Workstations constituting this DCE cell are connected over a wide area via the Energy Sciences network. The energy-research community is particularly fortunate to have this robust, high-performance T3 (45-Mbits/s) network connection among most sites.

Interprocess communications software[6] (IPCS) developed at LLNL provides asynchronous communication among processes in a heterogeneous computing environment[7] distributed among participating sites. IPCS provides an interface to the hardware in the experimental timing system and forms the basis for wide-area network synchronization to the real-time experimental operations. This software provides for coordination of tasks running both inside and outside the DCE cell.

An event-management service,[8] originally developed for local-area network applications in a VMS cluster at PPPL, was modified to utilize IPCS messages for communication over the wide-area network and coordination of information flow in a distributed task environment. This modified service minimizes the latency time between data availability and initiation of processing to obtain results required for both the operations group in the DIII-D control room and for remotely participating researchers. Similarly, using data-availability and data-processing events, the necessary visualization applications are launched to generate the desired displays and make data available to Web-browser applications.

Start setup for shot 92574, shot in 28 secs

```
24/6/1998 8:09
 Welcome to the Magnetic Fusion Energy Program
Lawrence Livermore National Laboratory
Meyer(#DIII-D): steve, if you want some slices without mse let me know
casper(#DIII-D): I don't see any cer results in /scratch/cerneur ?
Shotsync: Setup for DIII-D Shot 92570 Fri Jun 20 09:48:33 1997
Meyer(#DIII-D): efits for 92569 complete
----Available Logs for DIII-D shot 92570 Session Leader, Diagnostics, DIII-D
Shotsync: Acquiring DIII-D Shot 92570 Fri Jun 20 09:50:44 1997
Meyer(#DIII-D): efits for 92570: 1550-4550ms, 60slices
casper(#DIII-D): Steve, Is Dan going to run cer or should we pursue it?
casper(#DIII-D): I don't see any files showing up yet
Allen(#DIII-D): Dan is going to run them--he will start--he is going to do 69
Shotsync: Setup for DIII-D Shot 92571 Fri Jun 20 10:05:26 1997
----Available Logs for DIII-D shot 92571 Session Leader, Diagnostics, DIII-D
Shotsync: Acquiring DIII-D Shot 92571 Fri Jun 20 10:07:26 1997
Allen(#DIII-D): Dan is putting them in now -- 69 and 70 are in there
Meyer(#DIII-D): efits for 92571: 1550-4550ms, 60slices
Meyer(#DIII-D): efits coming in now
Shotsync: Setup for DIII-D Shot 92572 Fri Jun 20 10:24:22 1997
----Available Logs for DIII-D shot 92572 Session Leader, Diagnostics, DIII-D
Shotsync: Acquiring DIII-D Shot 92572 Fri Jun 20 10:26:18 1997
Meyer(#DIII-D): efits for 92572: 1550-4550ms, 60slices
Shotsync: Setup for DIII-D Shot 92573 Fri Jun 20 10:44:58 1997
----Available Logs for DIII-D shot 92573 Session Leader, Diagnostics,DIII-D
Shotsync: Acquiring DIII-D Shot 92573 Fri Jun 20 10:46:02 1997
```

ja_775
meyer
casper
moller
lasnier
allen
chuckg
rice
Shotsync

What is going to change on the next shot?

Figure 2. Java user interface to multiple-channel IRC-based communications showing countdown clock (top), dialog box (bottom), participants (right), and running history of dialog and operations (main box).

Using the results of a review by General Atomics and ORNL of the data flow within the DIII-D computer system, General Atomics has distributed the experiment's data-file system. A library of data calls[9] provides remote access to data elements within the experiment's file system. Using event identification and distributed acquisition of data, we can better schedule information flow and processing during experiments. With these enhancements, we are now evaluating the performance and suitability of DCE and the DFS to support experimental operations.

Our first full DCE application accomplished the distribution of the operations-critical processing required to analyze the magnetic equilibrium produced during an experimental "shot." Current fusion research uses pulsed experiments in which typically a high-temperature plasma discharge is formed for several seconds with about 15 min between discharges. Information from several instruments is sampled by analog-to-digital converters; this information is stored in local memory and sent to the computer environment for intershot analysis and display, which help to direct experimental operations. Data are archived for off-line, post-shot scientific analysis. Using many of these experimental measurements as constraints, the EFIT code[10] solves a mathematical inverse problem to give a reconstruction of the magnetic equilibrium. This code was modified to run in DCE. It allows us to increase greatly the number of reconstructed equilibria available to researchers in the control room by running parallel processing within our DCE cell.

With the cell as currently configured, we have already realized a considerable gain in processing throughput for support of experimental operations. We typically produce over 200 such reconstructed equilibria within 5 min after the shot. This provides for interpretation of the time-dependent evolution of a plasma discharge and meets the demands for intershot processing for both scientific analysis and control of the experiment. This processing gain will increase roughly linearly as workstations of comparable performance are added to the cell.

It is important to be able to use legacy applications in the new DCE. Many codes that process and display data depend on the Interactive Display Language (IDL, not to be confused with the IDL compilers associated with DCE) from Research Systems Inc.[11] This language is a commercial, non-DCE software package that allows easy development of analysis and display capabilities. In order

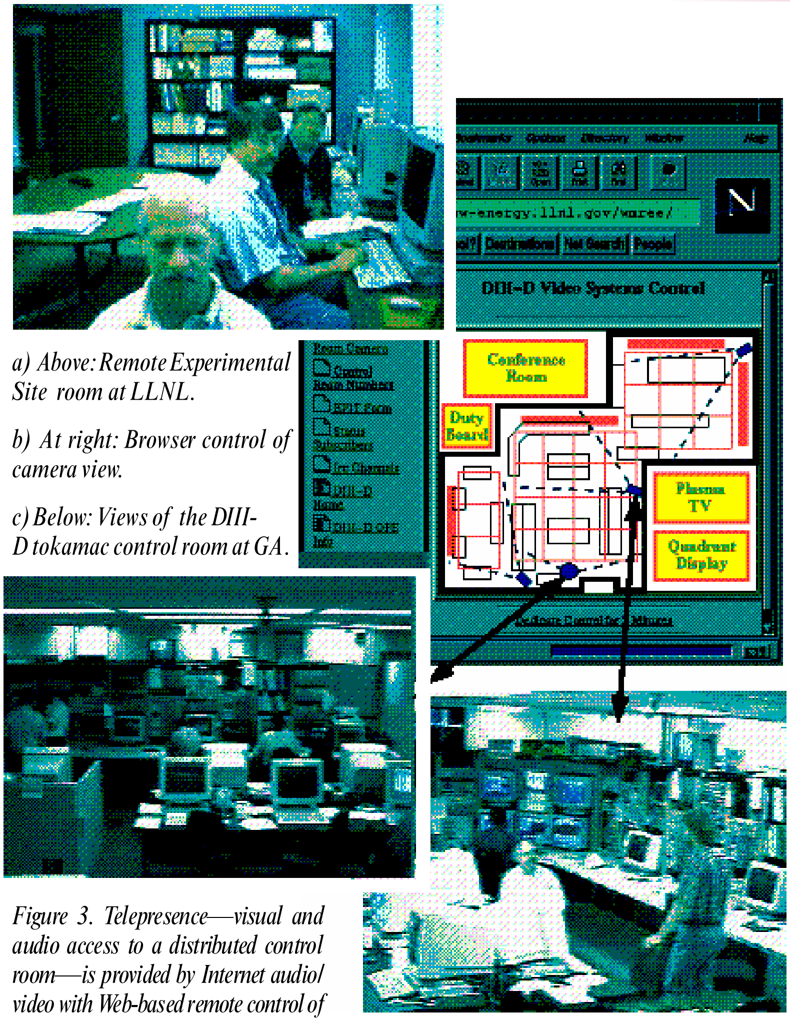

a) Above: Remote Experimental Site room at LLNL.

b) At right: Browser control of camera view.

c) Below: Views of the DIII-D tokamac control room at GA.

Figure 3. Telepresence—visual and audio access to a distributed control room—is provided by Internet audio/video with Web-based remote control of four cameras. Two views of the DIII-D control room (bottom) are shown along with a view of the LLNL remote site (top left).

to utilize the many user-developed IDL-based processing procedures, we have implemented a subroutine library that allows these processes to exchange messages with conventional processes written in C. IDL applications using IPCS include programs for processing and displaying data in synchronization with the shot cycle and an automated system for digitizing and processing video images. With a combination of asynchronous messaging and event declarations, we are able to use IDL code in parallel with DCE processing to generate the desired data displays. Both static displays and animations showing the temporal evolution of the experimental equilibria are generated and available between shots to researchers using either X-Windows or World Wide Web browsers.

In summary, a fully automated, multisite distributed-computing environment asynchronously slaved to the real-time experimental operations is operational. All users on- and off-site have access to data, synchronization messages, and event declarations. A secure, distributed environment provides high performance for compute-intensive applications such as the equilibrium reconstruction and soft x-ray and radiated-power-tomography applications.[12] Results from the DCE-based calculations are stored in the global file system. Data are accessible from the IDL visualization code to produce the required displays needed for tracking the experimental outcome. Since the DFS is employed both for storing results and for the IDL procedures (scripts), analysis can be run on workstations at off-site locations. Collaborators can therefore have access to both the data and analysis code.

Personal communications: IRC and audio/video

Remote participation in an interactive experiment such as the DIII-D tokamak requires personal communications among collaborating scientists. This is more convenient if provided as an integral part of the computations environment. The needed levels of interaction range from robust and interactive sharing of limited amounts of information to highly interactive exchanges within an audio/video environment. Internet Relay Chat (IRC)[13] has been adopted as a highly reliable and interactive means of communicating among sites. Reliability of communications is extremely important during operations, and IRC often succeeds when other "more elegant" computer-based solutions fail.

Instantaneous communication is maintained among collaborators as an aid in diagnosing fault situations arising during operations while allowing for a continuing dialog on the progress of the experiment. Automatic recording of the experimental state is also implemented by use of the messaging system to send information concerning the operations-state transitions—for example, next shot in progress, local time, or status of data-acquisition sequence—to the IRC session. This provides a meaningful time stamp for all communications and gives a permanent record of the operations. Summary dialog describing each experimental discharge is also displayed and provides an online log of the operations from which remote participants can follow the research in progress. A Java applet provides users with a Web-browser interface to the IRC session. This interface contains multiple channels dedicated to different aspects of the research such as, for example, physics operations or predominantly diagnostic instrumentation discussions. The Java interface, along with a short segment of an online discussion, is shown in Fig. 2. This tool has become quite popular and is actively used by several researchers during experiments.

Only a limited amount of information can be transmitted by an application as simple as IRC, and it does not provide the feeling of "presence" at the experiment. To better provide this feeling, Internet-based audio and video communications have been added to the DIII-D control room. SGI Indy workstations provide audio/video access to the control room using IP-multicast.[14] From their desktops, remotely located researchers can attend the daily morning pre-operations meeting at which the experimental plan for the day is presented and hardware status is discussed. Following this meeting, the camera view is switched to the control room so that remote participants can view operations and hear announcements. Two typical views of the control room are shown in Fig. 3; several other views are also available. Using the Web-browser

http://www-energy.llnl.gov/cgi-bin/

DIII–D Status – Shot is 93586

Date	Time Started	Approximate Secs to T0
10/1	16:59:20 PDT	Waiting for next

Last Shot State	Last Shot Event
[illegible]	priority_done

DIII–D Subscribers

●	BOLOMS	theodore	Busy	Diagnostic
●	DSPRED	theodore	Busy	Diagnostic
●	RESBOLOMS	ga_d3d	Busy	Diagnostic
●	D3D_MASTERTIMER	ga_d3d	Ready	Diagnostic
●	CERDATA	d3dvaxd	Ready	Diagnostic
●	DAPS_SERVER	ga_d3d	Ready	Display
●	[illegible]	[illegible]	Ready	Display

Figure 4. Status of instrumentation and experimental subsystems and state of the data-acquisition sequence, as displayed on a Web browser.

camera-control interface shown in Fig. 3, researchers often change their view of the control room depending upon their interest. The ability to "look around the room" as if present on site provides a greater feeling of involvement in the operation. Experienced participants can often infer the success of a run by merely viewing the operations. Also shown in Fig. 3 is the remote-site room at LLNL where researchers regularly participate in experiments and control instrumentation and data analysis.

The control-room audio environment is much more difficult to export, as is the capability for providing feedback to operations in a noninterfering manner. We are exploring methods for improving our use of audio communications. In the present arrangement, off-site researchers can hear announcements concerning the operation in progress such as, for example, the startup and purpose of a new shot. Many of the experiment's session leaders are now routinely announcing details of the operation for the benefit of local and remote participants alike. Even though two-way, interactive communications are available, they are only selectively used at present. Researchers can gather at the workstation audio/video session to discuss issues concerning the direction of the experiment. We envision the day when system availability, software, and network bandwidth better support the use of audio/visual connections for collaborative research.

In spite of their limited fidelity, the audio/visual connections are gaining in popularity for another application, that of broadcasting selected programmatic meetings. Important planning meetings are routinely broadcast over the Internet. Video conference sessions held over the Internet are supporting interactive meetings to discuss the ongoing physics research. For example, a biweekly meeting is currently being held to discuss continued development and application of UEDGE,[15] a large simulation code for modeling the edge region of the confined plasma.

Remote control of systems

Data for DIII-D are produced by several complex, special-purpose instruments. Among these instruments are optical and ultraviolet spectrometers, laser scattering systems, interferometers, soft x-ray imaging systems, neutron detectors, and electromagnetic probes. Many of these instruments that we use have been developed and are operated by collaborators from participating laboratories or universities. The interpretation of each experimental shot requires information associated with several instruments and with combinations of measurements. Because much of the instrument control now uses X-Windows–based interfaces, remote operation of the instruments is possible with a minimum of modification of those interfaces. Several such instruments have been operated remotely over the Internet. Their data have been analyzed at remote sites, and results have been returned to the control room during operations. Remote procedure calls provide another mechanism for connecting remote user interfaces to local (to the experiment) data-acquisition, control, and processing daemons; in this approach X-Windows interfaces do not need to be exported over the network.

Control of the magnetic shaping, heating, and fueling systems has been implemented using X-Windows. Due to the nature of these systems, some of which may supply millions of watts of power, machine- and personnel-safety considerations limit the amount of remote control. Security features with authentication and resource authorization are needed to address this issue. However, the feasibility of remote operation of a fusion experiment has already been demonstrated. In a collaboration[16] between LLNL and the Plasma Fusion Science Center at the Massachusetts Institute of Technology, the Alcator Cmod experiment in Cambridge, MA, was operated from Livermore, CA. Remote operation of the DIII-D tokamak has now also been demonstrated.[17] Using the Remote Experimental Environment, a team of scientists at LLNL conducted a set of experiments on the DIII-D tokamak at General Atomics. The demonstration included remote operations of the plasma-shape control system, a 20-MW neutral-particle-beam heating system, several instruments, and the intershot analysis system. Operation was supported by four audio/visual channels and proceeded at a rate similar to experimental operations conducted locally.[18]

Web-browser applications

Early in the Remote Experimental Environment project, we recognized that, unlike the Unix environment, X-Windows does not work as efficiently on PCs and Macintoshes, which require the use of X emulators. To achieve more heterogeneity, selected applications were migrated onto Web browsers. The previously mentioned camera control is solely based on a browser interface. We developed shot-synchronized browser displays that are automatically updated via the message-passing system. One such display, shown in Fig. 4, provides experimental status information such as identifying the instrumentation that is online, the state of data acquisition, the time to the next discharge, and which systems are busy. This display is routinely used off-site at LLNL to monitor the status of our instrumentation systems during experiments.

Web-browser access to the DCE is provided for scientists so that they can launch the equilibrium-reconstruction analysis previously discussed. In Fig. 5, we show the interface that provides users easy access to DCE compute cycles for interactive, personalized analysis of experimental data. The interface allows users to select processing options and parameters used in the processing. It is routinely used for post-run data analysis. Results of the online analysis can be distributed to Web browsers elsewhere on the network.

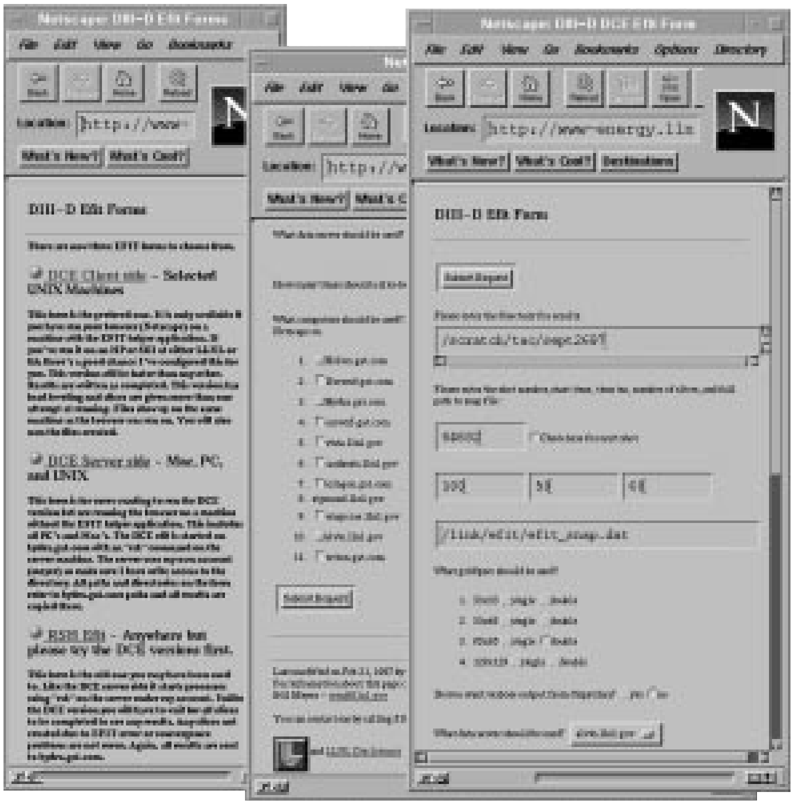

DIII–D Efit Forms

DCE Client side – Selected UNIX Machines

DCE Server side – Mac, PC, and UNIX

RSH Efit – Anywhere but please try the DCE versions first.

DIII–D Efit Form

/scratch/tac/expt.2697

84682

100 5 6

/link/efit/efit_snap.dat

Figure 5. Web-browser interface to the DCE EFIT application: forms provide user-selected initiation of processing (left), control of parameters used in analysis (right), and selection of computers used for DCE processing (center).

Between shots, as part of the data analysis, image files and animations that show the temporal evolution of the equilibria are generated. These files are available for interactive access from Web browsers. In addition, "live" Web displays deliver reconstructions of the equilibria to users with server-push techniques coupled to event declarations and IPCS messaging. These live displays are synchronized with the real-time operations. Automatically distributed waveform displays of discharge parameters (data from instrumentation controlled during the operations) as shown in Fig. 6 are available along with results of the DCE processing of the equilibria.

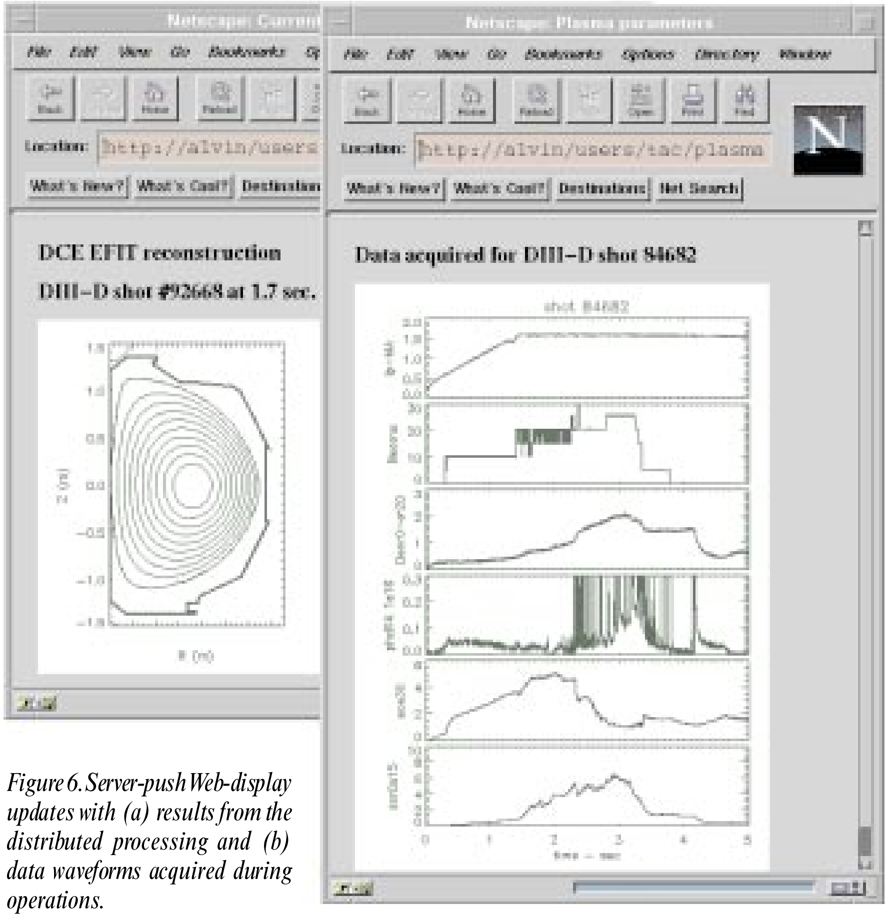

Figure 6. Server-push Web-display updates with (a) results from the distributed processing and (b) data waveforms acquired during operations.

Remote Experimental Environment in use

Now that hardware modifications and software have been installed to support the Remote Experimental Environment collaboratory, we are in a position to explore the technology and sociology[19] of remotely participating in the operation of a large-scale, highly interactive, and complex experimental facility. This is an evolutionary effort; we are still discovering many of the issues associated with research at a distance. Based on the experience gained and a sociological evaluation of our efforts, we shall optimize our techniques and add more distributed applications. Additional security enhancements, which are currently being pursued, are expected to provide more flexibility in the use of control applications from off-site locations.

DIII-D experimental campaigns are scheduled for several weeks of operation throughout the year. During these periods, the collaboration environment actively supports off-site participation in experiments. The audio/video broadcast of operations is available on a daily basis during experiments and is routinely monitored by several participants. Many researchers at the LLNL site have it available on their office workstations and regularly stay tuned to the experiments. The Java-IRC session has become a popular method for discussing operations; it, too, is available on a daily basis.

As with any program, some experiments are more popular than others; on a good day there may be 10 to 15 researchers using the communication tools. Even on-site staff who may not always be able to be in the control room are often logged in. The Web-based status-monitoring display finds routine use by people who are responsible for operation of instrumentation systems. The automatic updating of shot commentaries in the chat session allows researchers to monitor progress, review experiments completed earlier in the day, or, via Web searches of the archived sessions, go back and review summaries of previous operation.

With the addition of more processors to the DCE cell, it has taken on a production-processing role during experiments. Currently, all equilibrium reconstruction needed for operations, both for those running the shape-control system and for the online scientific data analysis, are provided by the distributed application during most experiments. The daily monitoring of this processing during experiments is now done from the LLNL site, where responsibility for this activity resides. Active participation during experiments from off-site locations is now available on a regular basis.

At the present time, this collaboratory is growing as physicists at more remote sites learn of its capabilities or begin collaborations on the DIII-D tokamak experiment. Magnetic-fusion experiments currently under construction are likely to incorporate some level of remote participation into their operations. We expect other experiments in the United States, as well as international magnetic-fusion research programs, to utilize this work and extend it to fit site-specific needs. The U.S. DOE2000 program[20] is currently funding the development of security architectures, applications of object-oriented technology, and pilot projects to enhance collaboratories further. With the advent of the Next Generation Internet (CIP 11:5, 1997, p. 405), bandwidth and network services that overcome present performance limitations should be available. We look forward to using these advancements in our experiments.

Acknowledgments

Funding for the Remote Experimental Environment project was provided under the Distributed Collaborative Experimental Environments Initiative at the Lawrence Berkeley National Laboratory. Work was also performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract number W-7405-Eng-48, by General Atomics under contract number DE-AC03-89ER51114, by the Princeton Plasma Physics Laboratory under contract DE-AC02-76-CHO3073, and by Lockheed Martin Energy Systems under contract DE-AC05-84OR21400.

This work has been a collaborative effort among the four geographically separated laboratories and many individuals within those laboratories. Special acknowledgment is extended to the National DIII-D Tokamak Program and its operations group for their patience, understanding, and willingness to allow us to implement and test these concepts during experimental operations. We also wish to acknowledge the useful interactions with S. Bly of Sara Bly Consulting for her very insightful evaluation of the way we work in our experimental scientific environment.

References

1. T. A. Casper, *Remote Experimental Site Concept for Diagnostic Collaborations in Fusion*, International School of Plasma Physics, Workshop on Diagnostics for Contemporary Fusion Experiments, Varenna, Italy, September, 1991.
2. T. A. Casper and W. J. Lennon, "Remote Experimental Site: A Command and Analysis Center for Big Physics Experimentation," in *Proceedings of the 14th IEEE/NPSS Symposium on Fusion Engineering* (IEEE Press, Piscataway, NJ, 1991), p. 704.
3. Committee on a National Collaboratory, National Research Council, *National Collaboratories—Applying Information Technology for Scientific Research* (National Academy Press, Washington, DC, 1993).
4. B. W. Rice *et al.*, "Observations of Enhanced Core Confinement in Negative Magnetic Shear Discharges with an L Mode Edge on DIII-D," Nuclear Fusion **36**, 1271 (1996).
5. DIII-D Tokamak Long Range Plan, General Atomics report GA-C19596, Rev. 3 (January 1993).
6. G. G. Preckshot and D. N. Butner, "A Simple and Efficient Interprocess Communications System for Actually Using a Laboratory Network," IEEE Trans. Nuclear Science **NS-34**, 858 (1987).
7. D. N. Butner *et al.*, "An Integrated, Multi-vendor Distributed Data Acquisition System," Rev. Sci. Inst. **59**, 1786 (1988).
8. T. Gibney, "QCS: Driving automatic data analysis programs for TFTR," in *Proceedings of the 8th Topical Conference on High Temperature Plasma Diagnostics*, Rev. Sci. Inst. **61**, 3271 (1990).
9. B. McHarg, "Access to DIII-D Data Located in Multiple Files and Multiple Locations," in *Proceedings of the 15th IEEE/NPSS Symposium on Fusion Engineering*, (IEEE Press, Piscataway, NJ, 1993), p. 123.
10. See http://lithos.gat.com/efit.
11. IDL is a registered trademark of Research Systems Inc., 2995 Wilderness Place, Boulder, CO 80301.
12. W. H. Meyer and T. A. Casper, "Tomography in a Distributed Computing Environment," Bull. Am. Phys. Soc. **37**, 1504 (1992); demonstrated at Supercomputing '92 in Minneapolis, MN (LLNL press release dated November 16, 1992).
13. See http://www.irchelp.org/irchelp/ircprimer.html.
14. See http://www.mbone.com/.
15. T. D. Rognlien *et al.*, J. Nucl. Mater. **196-198**, 347 (1992).
16. S. Horne *et al.*, "Remote Control of Alcator C-Mod from Lawrence Livermore National Laboratory," Fusion Technology **32**, 152 (1997).
17. T. A. Casper *et al.*, "Remote Experimentation on DIII-D: Operations and Concepts," 39th Annual Meeting of the Division of Plasma Physics, Pittsburgh, PA, November 17–21, 1997; see http://aps.org/BAPSDPP97/abs/S5000.html#S1WedP1.017.
18. T. A. Casper *et al.*, "Collaboratory Operations in Magnetic Fusion Scientific Research," Interactions **5**, 56 (1998).
19. S. Bly, K. M. Keith and P. A. Henline, "The Work of Scientists and the Building of Collaboratories," General Atomics Report #GA-A22619, June 1997.
20. See http://www.er.doe.gov/production/octr/doewpex.html.