PFC/JA-95-27

# First Demonstration of Remote Tokamak Control:
# Operation of Alcator C-Mod from LLNL

S. Horne, V. Bertolino, J. Daigle, T. Fredian,
M. Greenwald, S.N. Golovato, I. Hutchinson,
B. LaBombard, J. Stillerman, Y. Takase, S. Wolfe

August, 1995

# First Demonstration of
# Remote Tokamak Control:
# Operation of Alcator C-Mod from LLNL

S. Horne, V. Bertolino, J. Daigle, T. Fredian, M. Greenwald,
S.N. Golovato, I. Hutchinson, B. LaBombard, J. Stillerman,
Y. Takase and S. Wolfe,
*MIT Plasma Fusion Center;*
T. Casper, D.Butner, W. Meyer, J. Moller,
*Lawrence Livermore National Laboratory*

# 1   Abstract

Operation of a tokamak from a remote site has been demonstrated for the first time. The Alcator C-Mod tokamak, located at MIT in Massachusetts, was operated over the Internet from a remote control room set up at the Lawrence Livermore National Laboratory (LLNL) in California. Control of the physics parameters, such as plasma current, density, shape, heating power and active diagnostics was accomplished entirely from the remote site. Because the C-Mod control, data aquisition and display systems were designed to operate in a locally networked environment, extension of these systems to a wide area network was transparent, allowing use of the same user interfaces as when operating locally. Engineering control of subsystems such as vacuum, cooling, and power supply limits, remained under local control, providing appropriate equipment and personnel security. The operation was highly successful; important new physics data were obtained and valuable insight was gained into the potential and limitations of remote operation.

# 2 Introduction

As tokamaks have increased in size and cost they have relied increasingly on collaboration between the local staff and teams of scientists and engineers from other laboratories. In order to save travel time and money there has been a trend towards greater dependence on remote collaborators controlling tokamak diagnostics from their home institutions. Examples of this are Wisconsin – TFTR[1] U. Maryland – MIT[2], and LLNL – DIII-D[3]. Future machines such as ITER [4] are being planned with the capability for remote teams to carry out experimental programs without travelling to the tokamak site. This would require operation and monitoring of not just diagnostics but of plasma current, shape, density, heating, etc. Until now, however, there has been no demonstration of this mode of operation.

The advantages of remote experimentation are of course not unique to the controlled fusion program. The fields of radio and optical astronomy have perhaps taken the largest steps in this direction (at least for terrestial based experiments). Remote observing is now a routine part of the operation of most observatories[5] . The technical implementations are quite different from our own, but several common themes are clear. First, new installations, which incorporate modern computer control and standard communication protocols into their designs, have the least trouble in implementing remote control. Second, there is great concern in protecting the integrity of expensive and unique facilities. Third, social issues, which concern the ways in which teams of researchers work together, cannot be overlooked.

Because C-Mod physics operation is carried out on workstations with ordinary network links to various hardware interface systems, and because of the data system software design, it was recognized, from an early date, that the machine could be run from any location with a suitable network connection. The possibility of controlling C-Mod from LLNL was first discussed informally at the 1994 APS meeting. LLNL staff are actively involved in the research program of the DIII-D tokamak at General Atomics. As part of this collaboration the distributed computing environment for DIII-D includes computers at LLNL and supports remote access to DIII-D data and experiments. A plan for developing full remote control capabilities (the Remote Experimental Site, or RES[3]) is in place. While significant work remains to accomplish this, the control room itself was sufficiently functional to support remote operation of the C-Mod experiment.

## 2.1 Objectives

The objectives for the remote run were: 1. To demonstrate the technical feasibility of remote operation of a magnetic confinement experiment; 2. To test the capabilities of current technology including desktop video and other interpersonal communications tool; 3. To uncover problems or needs for new or improved tools; 4. To estimate infrastructure requirements such as network bandwidth and compute power for full scale operation of remote or distributed control rooms.

Early in our planning, it became clear that to meet these objectives it would be necessary to attempt a challenging experimental run with a serious physics objective. We selected an experiment designed to investigate the ability of C-Mod's dissipative divertor regime[6] to accommodate a step change in the total input power and remain in the detached state. This experiment required precise control of the position of the plasma and careful coordination among those controlling the machine, RF heating, diagnostics, and engineering systems.

# 3    Background

## 3.1    Machine Description

Alcator C-Mod is a compact high-field tokamak incorporating many design features relevant to a tokamak reactor. These include a closed poloidal divertor, all metal plasma facing components, a low resistance vacuum vessel not conformal to the plasma shape, and a sparse set of control coils. ICRF is the principal heating scheme. Like the other machines in the Alcator line, C-Mod features high particle, current, and power densities. The purpose of the Alcator research program is to study the confinement, heating, control, and power and particle exhaust of plasmas in divertor tokamaks. Some C-Mod parameters achieved to date are given in Table 1.

| Major Radius | R | 0.67 m. |
|---|---|---|
| Minor Radius | a | 0.22 m |
| Toroidal Field | $B_T$ | 8.0 T |
| Plasma Current | $I_P$ | 1.2 MA |
| Pulse length | | 1.7 sec. |
| Elongation | $\kappa$ | 1.8 |

Table 1: C-Mod Parameters.

## 3.2 Control, Data Acquisition, and the Computing Environment

### 3.2.1 Data Acquisition and the C-Mod Computing Environment.

The C-Mod computing environment consists of a cluster of fast workstations, running DEC's OpenVMS operating system, and X-terminals, connected via an FDDI and switched-ethernet LAN. Data aquisition during the pulse is accomplished mainly with CAMAC digitizers and counters distributed in the test cell and power conversion areas. All raw and processed data are read, manipulated and stored into hierarchical data structures (trees) by MDSplus, the Model Data System[7]. There are two features of that system that are important to mention at this point. First, the MDSplus data structures are as self-descriptive and complete as we can make them. That is they contain all the information for the data acquisition setup, diagnostic configuration, geometry, calibrations, along with the inputs and outputs for all analysis codes. This allows us to interact with the data system with a relatively small set of powerful software tools. Second, data system applications run under the client-server model in a highly distributed environment. This is a critical feature which allowed C-Mod to be easily operated remotely. A small set of dedicated servers with special interfaces to control and data acquisition hardware are linked via fiber optics to the tokamak systems; computers running client processes can be located anywhere on the network.

### 3.2.2 Real-Time Control

Real-time control of Alcator C-Mod is accomplished by two separate systems. Engineering control - vacuum, gas, and cryogenic cooling systems along with power supply current and voltage limits - is accomplished via

distributed Programmable Logic Controllers (PLC). Each PLC is linked to operator displays and controls running on PC's. The PLC system also implements a State Machine[8], which controls the C-Mod shot cycle during an experimental run. The PLC's operate at a relatively slow speed; thus engineering control is principally aimed at providing tokamak availability while protecting the hardware system's integrity. (Personnel safety is provided by redundant hardwired controls.)

PLC compute speeds are inadequate for feedback control of the plasma during an experimental pulse. During this period, the power supplies are under control of a hybrid[9, 10] digital-analog computer. This device implements a linear feedback control system. Signal paths in the hybrid are analog, resulting in high control bandwidth; gains however are digital and switchable during the pulse at time intervals $\geq 1$ msec[10]. Programming of the hybrid computer is carried out on workstations and the results downloaded through a specialized interface.

### 3.2.3 Display of Data

Time series are the most common form of raw data acquired during a plasma discharge. The most convenient method of examining this data is Scope[11], an X application which simulates a multichannel digital oscilloscope. Each panel displays one X-Y trace; various zoom and pan features are available. Each panel can be associated with an "update event" which is typically set when the data to be displayed in that panel becomes available. IDL[12] , a scientific visualization package, is also used extensively for analysis and display of the data.

### 3.2.4 Plasma Control System Software (PCS)

The data used by the hybrid computer to define a shot are stored in one of the branches of the C-Mod data tree. Data are read from and written to this tree by the PCS software, a large X-windows application written in IDL, which serves as the operator interface for the hybrid control computer. Through this interface, an operator can load old shots and modify them as required. Typical operations involve defining control quantities to be used for feedback, drawing waveforms, and computing and setting gains.

5

### 3.2.5 Remote Procedure Calls

A set of Remote Procedure Calls (RPCs) has been implemented to provide access to the data stored by MDSplus to remote collaborators and to users on non-VMS platforms. The standard MDSplus interface is provided via servers running on local nodes. The server receives an expression on the network from its RPC interface, evaluates it, and returns the result on the network. This made the porting of the Scope and PCS to non OpenVMS platforms possible.

### 3.2.6 Remote Experimental Site

The RES available to us at LLNL consisted of four SGI computers - two with desktop video (Video/Audio Tool - VAT[13, 14]) capability - and one HP workstation. While it would be possible to use the remote computers only as X-displays, the RPC interface to MDSplus described above allowed an alternative mode of operation for display intensive applications. These machines communicated via the LLNL LAN to the local ESnet hub. WAN connectivity is available from this hub at T3 speeds (45 Mbps) to Fermilab (FNAL) where a connection to MIT at T1 (1.5 Mbps) is made. This installation provided a sufficient number of display screens and sufficient bandwidth to carry out the experiment.

## 3.3 Interpersonal Communications

Responsibility for running the machine is divided between engineering and physics teams. The Engineering team is led by the Engineering Operator (EO); the physics team is led by the Session Leader (SL), who is responsible for the overall physics objectives of the run, and the Physics Operator (PO), responsible for implementing the shot-to-shot changes in plasma conditions required by the SL. Communication within the engineering and physics groups is quite informal, but communication between the groups is, by design, passed only between the EO and the PO.

During a run, information from the physics team is passed to the session leader, who directs the physics operator to modify the next shot; the PO, possibly in consultation with the EO, makes the programming changes to implement the wishes of the session leader. When the PO is ready, he so

6

informs the engineering operator. When all engineering constraints have been met, and the PO reports "ready", the EO begins the next shot by forcing the state machine to the "INIT" state at which time the download of programming information to the hardware begins. From this point on, no programming changes can be made for the upcoming shot.

Besides this level of routine shot-to-shot communication, a certain amount of less routine but still formal communication occurs between the EO and PO. Examples include diagnostic gate valve permissives, access to power room or cell, questions about or changes in current limits or gas plenum pressure. The EO and PO often engage in a significant amount of less formal interaction as well, usually to resolve cause/effect questions involving unexpected behavior of the plasma or power systems.

Operating the RF heating system also requires the coordinated operation of engineering and physics control systems. The RF group must accomplish two tasks between shots. They must establish suitable antenna matching for the next plasma shot, an "engineering" task which must be done locally, since it is a between-shot activity controlled by a PLC subsystem. They must also program the time-dependent power waveform for each transmitter; this is a "physics" task which can be done remotely. It is the responsibility of the PO to verify that the RF group is ready before starting the next shot.

To ensure that the scientific objectives of the run are being fulfilled, the session leader needs to stay in close contact with members of the physics team (PT). The PT varies in number from run to run but typically is comprised of 10-20 scientists and students, who operate diagnostic systems, monitor the raw data as it is collected, and perform preliminary analyses between shots. This level of communication is very informal and by its nature, (one-to many) is difficult to arrange when the SL or some members of the PT are at a remote location.

# 4  Planning and Preparations for Remote Control Experiment.

Before committing a day's run to the remote control experiment, a number of preliminary steps were required. First, we needed to decide which experiment would be carried out during the run. This determined which subsystems,

in addition to the basic tokamak control, would need to be controlled and monitored remotely. For the selected experiment, we required remote control of the Fast Scanning Probe (FSP), and of the RF power and timing. These choices also determined the composition of the remote team. The session leader (who also controlled the FSP), RF operator, and physics operator (responsible for plasma control) would compose the remote team; in addition one of our computer staff was included. Secondly, from MIT, we installed and tested all the software that would be required and verify that it would work properly from the remote site. Finally, we conducted a preliminary visit to the LLNL site for a final checkout of systems and procedures.

## 4.1  Remote Operation of MDSplus Software

The software used to control and monitor C-Mod was designed to operate transparently in a client/server network environment, either in "local" mode (remote computers used only as X-displays) or in "remote" mode (transferring data via the RPC protocol). With only a minor effort, we were able to verify the proper operation of all the necessary tools. Of greater concern was the performance of these tools given the reduced network bandwidth available to us. In the control room, data was distributed over a 100 Mbps backbone; cross country, we would be sharing a T1 link with total capacity of 1.5 Mbps. Preliminary experiments convinced us that we could count on $\approx$ 200 Kbaud of this for our own traffic and that it would be adequate for our purposes. Various techniques were used to improve the interactive feel of the display and control tools over this (relatively) slow connection.

Next, we tested the two modes of operation for applications which were to be used in the remote experiment. In the first, an application is run on computers at MIT with its display set to the remote site. Because the data transfer from the data acquisition archive is all within the C-Mod control room with its very fast communications infrastructure, this mode works best for data intensive applications. Interactive operations are somewhat slower, however, since key clicks or mouse motion must be transmitted both ways over the link via the X protocol and the display updated over the slower WAN. Alternatively, an application can run at the remote site using the RPC protocol to pull data over the WAN. This approach works best for display intensive applications. While the data transfer is slower, once the data are cached at the remote site, changes in the display (zooming, panning, etc.)

8

are done locally and require no X transport across the slower link.

## 4.2 Communications

One of the objectives of the remote run was to test various communications options. Since these tools were new to us, significant redundancy was included in our planning. Thus three techniques were used in various combinations: video and audio on our VAT capable workstations (packetized and sent over the Internet using the MBone, typing via the Internet Relay Chat (IRC), and ordinary telephones. Since the EO/PO communications were considered to be the highest priority, the VAT system was arranged with this in mind, though some effort was made to facilitate face-to-face communication between the SL and physics team as well. In addition to preliminary tests, a planning meeting with MIT staff in Massachusetts and LLNL staff in California was carried out using the VAT.

To maximize our use of the single VAT computer in the C-Mod control room (an SGI Indy), four video sources were combined with split-frame hardware and the composite image then transmitted. Two of the four video inputs came from cameras in the engineering and the physics portions of the control room. A third input showed the experimental cell. The last input came from one of several cameras which image the plasma; these views are available on monitors in the C-Mod control room. Audio from microphones in the engineering and physics areas was mixed with the control room PA system and transmitted.

Two VAT-capable computers were available at the RES in Livermore. One was used to communicate directly with the C-Mod SGI. The other was used to transmit video only using nv to a 486 PC running Linux[ref?]. The nv process on the PC had its display set to a VAX workstation in the physics area normally used by the PO. (We could not display directly on the C-mod machines from the RES, due to Motif or X version incompatibilities). This arrangement provided two "stations" in the C-Mod control room, each with a microphone, camera, and display of the RES, and allowed "face-to-face" communication among the SL, PO, EO, and a representative of the physics team. The PO/EO channel was routinely used as another link to the PT with the understanding that the EO and PO had interrupt privileges.

IRC is a flexible text based network communications program. Users connected to an IRC server can create and join channels, and select which

channels they will monitor. Communication is by selecting a particular channel and typing; traffic on channels being monitored scrolls in a window. To gain some experience with this tool, we used it locally between the PO and EO for one run; no significant problems were encountered.

### 4.2.1 Preliminary Visit to Remote Site

To evaluate our state of preparation and to settle any remaining issues, two C-Mod staff visited the LLNL site on March 13 and 14, two weeks before the run. The objective of the visit was to "eavesdrop" on the March 14 run, and verify operation of the various tools. Assorted problems were discovered and addressed; no serious impediment was discovered but some interesting difficulties emerged. The worst of these related to the heterogenous collection of computers at the remote site. Three different operating systems were in use there. Various incompatibilities were observed, primarily involving Motif and X; certain cut/paste Scope operations in particular tickled these bugs. Another minor source of annoyance was that the keyboards mappings on these machines varied from each other and from those used at C-Mod.

During the tests on March 14, diagnostics run at the remote site identified the cause of some erratic plasma startup behavior; recommendations to correct the problem were passed to the C-Mod control room via IRC. Control was given to the RES for three shots to implement the recommendations. These shots ran successfully; to our knowledge these shots were the first remote operation of any tokamak. Based on these results the decision was made to proceed with the full run on March 28.

## 5  Remote Operation of C-Mod

The remote run schedule called for the RES to take over operations on March 28 at 11:00 AM Eastern Time, or 8:00 AM Pacific time. Two C-Mod staff and the RES staff spent March 27 setting up the RES computers and running final tests. When the full C-Mod group arrived at 6:00 A.M on March 28, it was discovered that a local network problem had occurred, which required rebooting the RES computers. This delayed transfer of control by about an hour. The balance of the run went well, albeit slowly. Interactive response over the network link seemed subjectively slower than we expected, based

on our experience two weeks earlier. The time required either to bring up a Scope, or to make a change in PCS, was significant; we tended to rely on the physics team at MIT to diagnose problems, and to make changes involving PCS only when absolutely necessary. In spite of these difficulties we took 21 shots in about 6 1/2 hours; at our nominal rate we would have taken 26 shots.

During this period we demonstrated remote control of plasma shape, density, RF power, and the reciprocating probe; we also interacted with another remote collaborator, who controlled a spectroscopic experiment at C-Mod from his office at University of Maryland. Typically, at least one change in some parameter was required for every shot; sometimes several. Changes in density required the most coordination among all the parties. The feedback-controlled gas system does not, by itself, provide a sufficient particle source to access the required high-density operating regime. Additional gas is usually injected from a system of capillaries; control of this system remained with the physics team at MIT. In addition, the RF coupling varies with plasma density; tuning for a new density requires close coordination between the remote operator and the local team. Finally, as the plasma resistivity increases with increased fueling, thereby increasing the poloidal flux consumption, currents in some coils approach their engineering limits. This can cause subtle shaping errors which can be critical, however, because the Fast Scanning Probe, designed to diagnose the scrape-off layer, cannot survive even brief exposure to the confined plasma. Careful adjustments of the shaping system are required to control this effect.

Due to the late start and network slowdown, the run was several shots short of completion. It was decided to finish the physics portion of the run the following day; a configuration which showed both H-Mode confinement characteristics and divertor detachment was reloaded for further study. By monitoring the initial run, staff from the ESnet network operations center were able to trace the bandwidth limitations we experienced to network congestion on an ethernet segment at FNAL, and to route our communications path around the misbehaving segment. Interactive response improved significantly; the improvement in bandwidth combined with some minor tuning of the software resulted in a factor of five or better reduction in the time required to load a modification into the control system. The shot cycle time decreased to our nominal fifteen minutes per shot as a result of these improvements. and the program was successfully concluded.

11

## 5.1 Evaluation

The experiment succeeded in its immediate objective, which was to demonstrate that remote control of a tokamak over the Internet is possible with technology which exists today. The tools used routinely for C-Mod control and display of data were completely adequate for remote operation over the net. Real physics was accomplished even in the context of a demonstration project. We have also demonstrated that there is no obvious limit to the degree that a remote collaborator can become involved in a C-Mod experiment. The technology exists for a sufficiently authorized and knowledgeable person or group to interact with the C-Mod program up to the point of controlling any or all major systems remotely.

We believe that the experiment demonstrates the feasibility of remote operation for future fusion experiments both large and small. It is clear that with the pace of change in the computer and communications industry, the particular implementations that was used for the experiment described here will not be directly applicable past the near-term. However, similar principles will still apply. The improvements we are certain to see in technology will make remote operation of future projects easier, but perhaps only if we plan for it from the beginning.

### 5.1.1 Network Requirements

Despite the success of the experiment, it was clear that the network bandwidth was marginal and clearly inadequate for a "production" version of a remote control room. Available bandwidth limited the framing rate for the video and more importantly, the rate at which data was available at the remote site. The link speed for a full scale remote control room is likely to be substantial. We note that C-Mod LAN was upgraded several years ago from 10 to 100 Mbps to provide adequate throughput. This is for C-Mod with data rates of 60MB/shot and no more than 20-30 people looking at data. An ITER scale device would clearly make much greater demands. Video at 128 kbps per link per direction is just barely adequate even now. In a full scale deployment with a dozen links running both ways at several times this speed, we estimate the needed bandwidth to be in the 10's of Mbps with guaranteed throughput and latency.

## 5.1.2  Interpersonal Communications

"Sociological" issues – the dynamics of handling informal communications between people at two (or several) separate sites – may be more difficult to deal with than the purely technical issues involved in remote control. Subjective reactions to the experiment tended to correlate with the participants prior experience with the various communications tools – particularly the IRC program.

Communication between the remote physicists and local engineers was quite satisfactory; this was not surprising, since the VAT system had been arranged for this purpose, and those involved had practiced with the IRC program. However some members of the physics team at MIT felt "disconnected" from the run unless they were sitting in front of one of the VAT stations; they found the IRC program difficult to use. Our unfamiliarity with the tool was part of the problem; when multiple users are typing on a single channel, questions and answers become mixed in a confusing way. The fact must also be faced that not all the participants were blessed with fluent typing skills. Another problem comes from the lack of closed loop response. Response to typed messages was not always immediate; there was no way of knowing whether the target person was not watching, busy, or formulating a reply. In spite of these difficulties, the IRC carried the bulk of the routine communications, and logs of the conversations provided a useful supplement to the normal logbook entries.

VAT communication was a success. We were impressed with the way in which the VAT seemed to bring the C-Mod experiment to the remote site. The control room audio annunciator counting down the seconds to the shot, combined with the transmitted view of the plasma, gave real-time feedback which made the experiment feel immediately accessible. The frame rate available at the RES was sufficient to see several frames of the plasma evolution which provided immediate feedback to the remote staff as to the quality of a particular discharge. Figure 1 shows a frame capture made during a pulse; the plasma is seen to be lower diverted, since the visible emission, mostly due to recycling, is seen to be concentrated there. The small area of emission in the upper left corner is a "plume" due to gas being puffed from a capillary on the inner wall.

Although the experiment could have been made to work without this tool, its presence contributed an atmosphere to the remote site that would have

13

impossible to duplicate by other means.

It is clear that more VAT systems at each end would have been an improvement (assuming adequate bandwidth to support them). The RF group could have made good use of face-to-face communication and such capabilities might have vastly improved the communication between the session leader and the physics team. If all physics operations were carried out at one remote site, it would not be as crucial an issue.

## 5.2  New Tools Needed

As mentioned above, improvements in the desktop audio/video tools will be a great aid for remote operations. This will require more than an improvement in sound and picture quality or bandwidth. For a full-blown remote operation, we should assume that every participant would be sitting at a station with A/V capabilities. The issue becomes one of managing the resulting melange of conversations and conferences. The software tools should allow one to see a full "list" of participants, to indicate who is talking to who, to join conversations, to add or drop participants from conferences, to insure privacy when needed, to "broadcast" crucial announcements, etc. A wide variety of desktop audio/video tools are becoming available and commercial systems may soon provide the capabilities we need.

The complexity of the task of emulating a normal group discussion using software tools should not be underestimated. In the control room, small groups of scientists and engineers often sit in front of the same workstation and discuss results displayed there. They may take turns at the keyboard (or mousepad), manipulating the display or bringing up new data. Even with good a audio/visual connection, this is not easy to do over the net. Tools to allow any application to be distributed between sites need to be developed. Participants at both sites should be able to see the same display and they should be able to take turns in interacting with it. Distributed "whiteboards" are an example of such a tool, but the challenge is to make this work for any application that might be employed.
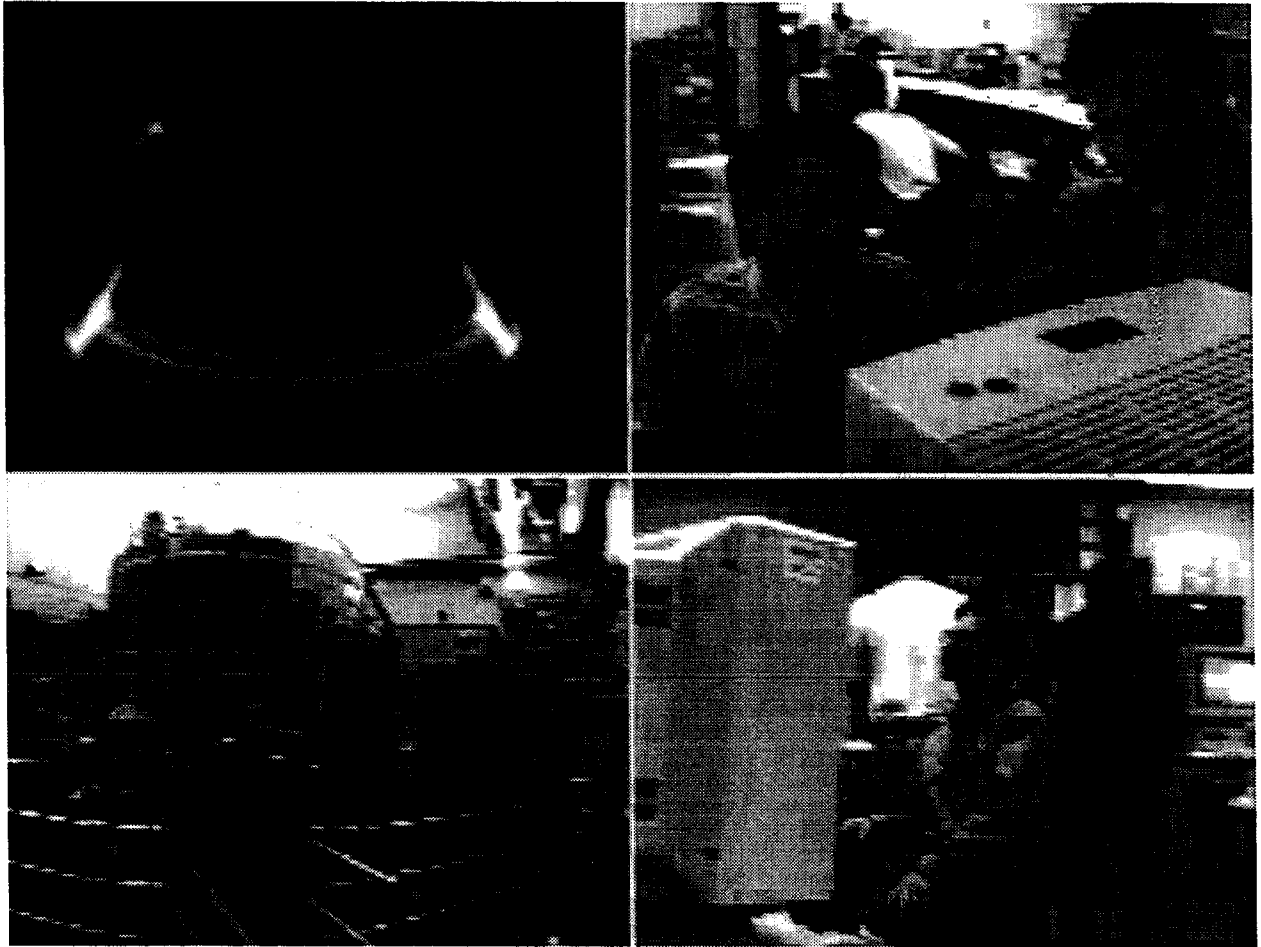
14

Figure 1: Frame capture of transmission from C-Mod to LLNL

15

# 6  Summary

In this paper we have reported on the first remote operation of a tokamak experiment. Although the C-Mod control system was not designed with remote control in mind, capabilities built in from the beginning enabled us to carry out this demonstration with relatively little effort. While making no claims for the future applicability of our own implementation, we believe our experience will allow a reasonable estimation of network bandwidth requirements and an assessment of the needed communications tools. We feel that this experiment was a "proof of principle" for routine remote operation of future machines such as ITER. The fact that the attempt required no special-purpose software or hardware to be developed exclusively for remote operation implies that including remote control capability in the design of a tokamak should not add significant costs.

# 7  Acknowledgments

16

# References

[1] R.J. Fonck, G. Cosby, R. Durst, T. Gibney, M. Thompson, and S.F. Paul. Remote Operation of the TFTR BES experiment from an offsite location. *Rev. Sci. Instrum*, 63(10):4803 – 4806, 1992.

[2] Benjamin Welch et al. Visible and Ultraviolet Light Measurments from Alcator C-Mod. In *Bulletin of the American Physical Society*, volume 38, page 1957. American Physical Society, November 1993.

[3] T.A. Casper and W.J. Lennon. Remote Experimental Site: A Command and Analysis Center for "Big Physics" Experimentation. In *Proceedings of the IEEE/NPSS 14'th Symposium on Fusion Engineering, San Diego, CA.*, October 1991.

[4] D. Barnes, S. Davis, P. Roney, P. Henline, T. Casper, and M. Greenwald. Collaborations in Fusion Research: A View to the Future Needs and Tools for Experiments. Technical Report PPPL Report 3025, Princeton Plasma Physics Laboatory, 1994.

[5] Darrel T. Emerson and Roger G. Clowes, editors. *Observing at a Distance*, Tucson Arizona, 1992. American Physical Society, World Scientific.

[6] B. Lipshultz, J. Goetz, B. LaBombard, G.M. McCracken, J.L. Terry, M. Graf, R.S. Granetz, Jablonski D, C. Kurz, A. Niemczewski, and J. Snipes. Dissipative Divertor operation in the Alcator C-Mod tokamak. In *Journal of Nuclear Materials*, 220-222, pages 50–61, 1995.

[7] G. Flor, G. Manduchi, T.W. Fredian, J.A.Stillerman, and K.A. Klare. MDS-Plus - A Comprehensive Data Aqusition and Analysis System. In *Symposium on Fusion Technology, London, U.K., 1990*, 1990.

[8] Joe Bosco and Steve Fairfax. The Alcator C-Mod Control System. In *Proceedings of the IEEE/NPSS 14'th Symposium on Fusion Engineering, San Diego, CA.*, pages 782 – 785, October 1991.

[9] Stephen Horne et al. Performance of the C-Mod Shape Control System. In *Proceedings of the IEEE 15'th Symposium on Fusion Engineering, Hyannis, MA.*, pages 242 – 245, 1993.

[10] P.F. Isoz, J.B. Lister, and Ph. Marmillod. A Hybrid Matrix Multiplier for Control of the TCV Tokamak. In *Symposium on Fusion Technology, London, U.K., 1990*, page 1264ff, 1990.

[11] Thomas Fredian and Josh Stillerman. An X-Windows Based User Interface for Data Acqisition and Display. In *Proceedings of the APS 18'th Topical Conference on High Temperature Plasma Diagnostics*, 1990.

[12] Research Sustems, Inc., Boulder, CO 80301. *IDL User's Guide*. IDL is a registered trademark of RSI, Inc.

[13] Michael R. Macedonia and Donald P. Brutzman. MBone Provides Audio and Video Across the Internet. *IEEE Computer Magazine*, pages 30–36, April 1994. Also available at ftp://taurus.cs.nps.navy.mil/pub/mbmg/mbone.html.

[14] Steve Casner. Frequently Asked Questions (FAQ) on the Multicast Backbone. ftp://venera.isi.edu/mbone/faq.txt, 1993.

# A Sample of IRC Conversation

To illustrate some of the strengths and weakness of communication via the Internet Relay Chat program, an extract from the IRC log is shown in figure 2.

The extract begins with the command from the PO to EO to take the next shot (shot 21), continues through discussion, confusion, extraneous conversation, to the command to take shot 22. The parties in this conversation are the P.O (horne) and S.L. (Brian) at the remote site, and the E.O. (daigle), and members of the P.T. (wolfe, Martin, blip) at C-Mod. The P.O. is using channel eo to communicate with the E.O., and channel cmod to deal with physics issues. Conversations between horne and wolfe, and between Brian and blip, and between Martin and horne, are interleaved in a manner which led to confusion at times. However a good deal of the confusion could have been eliminated by using separate channels for the different conversations, and assigning a separate window to each. The IRC program can do this, in principal; lack of familiarity with the tool combined with software incompatibilities prevented us from trying this mode of operation.

One development which significantly simplify the use of text-based communications tools is the emerging availability of speech-to-text and text-to-speech software. One could imagine an IRC-like tool which allowed one to speak into a microphone instead of type, and could be configured to convert a user-chosen channel back to speech. This would allow spoken conversation between two people with simultaneous scrolling of other conversations in text windows. A prototype of such a tool could be assembled today and might address some of the concerns raised in section 5.1.

```
<horne:#eo> Joe, we are ready
<#state> Tue Mar 28 13:52:20 1995 *** Entering INIT    ****** 950328021 *******
<#state> Tue Mar 28 13:54:17 1995 *** Entering CHECK   ****** 950328021 *******
<#state> Tue Mar 28 13:54:59 1995 *** Entering PULSE   ****** 950328021 *******
<wolfe:#cmod> Nice one steve
<#state> Tue Mar 28 13:55:14 1995 *** Entering RECOOL  ****** 950328021 *******
<wolfe:#cmod> You have better sources of information than I do.
<horne:#cmod> I see two basic problems here so far.  First, the low bandwidth
<horne:#cmod> makes all PCS actvity very slow.  Drawing, etc is acceptable,
<horne:#cmod> but loading etc causes all those labels to blink which takes lots
<horne:#cmod> of time.
<horne:#cmod> The other problem is that the computers/network here seems not to
<horne:#cmod> hold up well under the sort of hammering we're giving it.
<wolfe:#cmod> Another shot, another h-mode. Ho-hum!
<Martin:#cmod> Note that we are chewing up a lot of bandwidth with the video.
<horne:#cmod> I would like to back down the video frame rate.  Is that ok?
<Martin:#cmod> Sure - you have two channels going here. - they are apparently
<Martin:#cmod> using 240K apiece
horne:#cmod> ok -- vista is backed down about a factor of 2.
<horne:#cmod> Steve W, efc looks ok now.
<wolfe:#cmod> Yes, I don't know quite what the problem was before
<Martin:#cmod> Steve (H) - let us know if that helped the PCS responsivity.
<daigle:#eo>  Engineering is ready.
<horne:#cmod> steve w, we're tweaking upper xpoint to satisfy brian.
<wolfe:#cmod> Be careful. You are on the edge with EF1U
<horne:#cmod> we may be having another network problem.
<wolfe:#cmod> Try moving RXU out, not ZXU up.
<wolfe:#cmod> Note that we had put a ramp into RXU to try to accomodate what
<wolfe:#cmod> it wanted to do anyway. This is similar to what you and I were
<wolfe:#cmod> looking at last week
<blip> On the last shot we got 1.5e21 on probe 4.  There does not appear
    to be any detachment, the 'death ray' is still 'alive'.
<horne:#cmod> Next shot has zxu higher by 5 mm, starting at about .6 sec.
<wolfe:#cmod> Steve H-did you receive my previous message warning you about EF1U?
<blip> 'death ray' on probe 4 of inner divertor (rho=2 mm) as well on shot 21.
<Brian> Thanks, we are going to repeat this shot with the FSP.
<horne:#eo> yes, steve w
<wolfe:#cmod> Steve H-did you receive my previous message warning you about EF1U?
<blip> What about the oscilloscope
<horne:#cmod> steve w, I hear you.
<Brian> I sent you some mail
<blip> ok I will look
<wolfe:#cmod> Yes, I got that. Did you get mine?
<blip> Do you have a FAX machine near?  I could send you some copies of
the proccessed data.
<horne:#eo> physics is ready
```

Figure 2: Extract of IRC log.