GA–A25960

REMOTE PARTICIPATION OF LARGE-SCALE FUSION EXPERIMENTS

by
D.P. SCHISSEL, G. ABLA, M.J. GREENWALD,
T.W. FREDIAN, J. STILLERMAN

APRIL 2008

GENERAL ATOMICS

DISCLAIMER

GA–A25960

REMOTE PARTICIPATION OF LARGE-SCALE FUSION EXPERIMENTS

by
D.P. SCHISSEL, G. ABLA, M.J. GREENWALD,*
T.W. FREDIAN,* J. STILLERMAN*

*Massachusetts Institute of Technology, Cambridge, Massachusetts

This is a preprint of a paper to be presented at the International Conference on Accelerator and Large Experimental Physics Control Systems in Knoxville, Tennessee on October 15–19, 2007 and to be published in the Proceedings.

Work supported by
the U.S. Department of Energy SciDAC Program
and at General Atomics under
Cooperative Agreement No. DE-FC02-01ER25455

GENERAL ATOMICS PROJECT 30106
APRIL 2008

GENERAL ATOMICS

Abstract

This paper examines the past, present, and future remote operation of large-scale fusion experiments by large, geographically dispersed teams. The fusion community has considerable experience placing remote collaboration tools in the hands of real users. Tools to remotely view operations and control selected instrumentation and analysis tasks were in use as early as 1992 and full remote operation of an entire tokamak experiment was demonstrated in 1996. Today's experiments invariably involve a mix of local and remote researchers, with sessions routinely led from remote institutions. Currently, the National Fusion Collaboratory Project has created a FusionGrid for secure remote computations and has placed collaborative tools into operating control rooms. Looking towards the future, ITER will be the next major step in the international program. Fusion experiments put a premium on near real-time interactions with data and among members of the team and though ITER will generate more data than current experiments, the greatest challenge will be the provisioning of systems for analyzing, visualizing and assimilating data to support distributed decision making during ITER operation.

INTRODUCTION

For 50 years magnetic fusion energy has been an open collaborative science with worldwide participation. As the science has progressed, experimental facilities have grown in size and complexity. Concurrent with this growth has been a reduction in their number resulting in increased growth in collaborations. Today within the U.S., there are three large magnetic fusion facilities yet there are over 40 distinct organizations and ~1000 scientists engaged in fusion research.

The community's next major step is the construction and operation of ITER, a burning plasma magnetic confinement experiment to prove the scientific viability of controlled fusion as an energy source [1]. ITER will be located in France and will run as an international collaboration, with researchers from China, Europe, India, Japan, Korea, Russia, and the U.S. sharing operational and scientific responsibilities. It is expected that the full scientific exploitation of ITER will only be possible with substantial and efficient remote collaboration infrastructure.

Unlike very large experiments in other fields, the operation of fusion energy experiments put a premium on near-real-time interactions with data and among team members to support decision making during operations. In today's experiments, plasma pulses, or shots, are typically taken 2-4 per hour with a total of several thousand per year. The average cost of a plasma pulse is large (for ITER; total project cost amortized over all pulses is ~$1M) and therefore pulses required to carry out an experimental program must be minimized.

PREVIOUS WORK

The U.S. magnetic fusion community has a long and successful history of placing collaboration tools in the hands of real users. As early as 1992, a major diagnostic on the Tokamak, Fusion Test Reactor (TFTR) at Princeton Plasma Physic Laboratory was being operated from an off-site location [2]. Expanding beyond just diagnostic operation, just four years later both the Alcator C-Mod (Cambridge, MA) and DIII-D (San Diego, CA) tokamaks demonstrated full remote operation from a control room setup at LLNL [3,4].

Two important lessons were learned from these activities. First, the structure and efficiency of the data systems and its capability for transparent remote access were critical for successful remote operation. The MDSplus data management and data system [5] was initially developed in the late 1990s and is presently in use by over 30 experimental facilities worldwide including the two newest tokamaks EAST in China and KSTAR in South Korea. Based on a client/server model, MDSplus provides a hierarchical, self-descriptive structure for simple and complex data types and is used to store digitized, analyzed, and simulation code data. Second, it was recognized that effective tools for interpersonal communication in a geographically distributed environment were crucial and needed significant improvements.

These and other efforts were mostly carried out on an *ad-hoc* basis. To expand on these efforts, the National Fusion Collaboratory Project (2001-2006), part of the first round of SciDAC (Scientific Discovery through Advanced Computing) projects, consolidated previous work, developed new capability, and deployed collaboration tools to all major U.S. experimental facilities [6].

One of these new capabilities is the deployment of a national fusion energy sciences grid (FusionGrid) for securely sharing resources (data, codes, visualizations, communication) over the Internet. The goal of the FusionGrid is to allow scientists at remote sites to participate as fully in experiments and computational activities as if they were working on site. Access to data on FusionGrid uses a secure version of MDSplus that requires an x.509 identity certificate for authentication. Combining this with a unique distributed authorization scheme allows stakeholders to confidently grant restricted access to data, computers or codes as required. A web-based real-time monitoring system provides users and administrators with a detailed view of FusionGrid activity. Deployed initially within the U.S., it is now used across the world.

To assist in the human aspect of collaboration, the project deployed several novel technologies. Experimental control rooms are staffed by not several but by 20-40 individuals. Large screen shared display walls were installed along with customized software in the three main control rooms to facilitate collocated collaboration. Scientists have the ability to take a visualization displayed on their individual workstation and move it to the display wall for a group discussion. This same software can be used by remote scientists to either place their visualization on the display wall or to share with an individual scientist's display in the control room. Audio and video communication with remote scientists is handled through several different Internet based solutions.

Control of plasma pulses has long been performed digitally but with the increasing collaboration within the fusion community an increased effort has been undertaken to perform collaborative software development as well as collaborative usage. The Plasma Control System originally developed for the DIII-D tokamak [7] has been expanded in capability and is presently in use at six different experiments worldwide.

REMOTE OPERATION TODAY

Today, it is now routine for fusion scientists at institutions around the world to lead and participate in experiments on the major U.S. facilities without leaving their home institutions.

Data including raw, analyzed, and meta-data is securely available simultaneously to scientists worldwide. Utilizing MDSplus and relational database technology, data is available in a common API that allows interoperability of tools amongst experiments. Scientists support decision-making during experimental operations through the use of customized visualization applications to examine data.

Substantial data analysis, both automatic and user generated, is performed during experimental operation. All of this analyzed data is available to a participating scientist regardless of location. Off-site scientists as well as those individuals in the control room are expected to write their analyzed data back into the general data repository so that all can benefit. Analyzed data to support an operating experiment does not have to be generated on a local computer but can be located worldwide. Therefore, with this capability scientists can virtually bring their expertise and code to an experiment by placing it on FusionGrid. This eliminates the need to make codes portable and greatly reduces the work involved in "bringing" a code to an experiment. Another form of shared data analysis is human generated comments via an electronic logbook with a custom client. Here scientists worldwide can make written entries about the experiment that are instantaneously shared amongst all researchers. This rapid information sharing is very valuable when individuals are not all in the same room.

Interpersonal communication is supported to allow distributed scientists to discuss data analysis and changes to the experimental plan in pseudo-real-time. Since fusion experiments are not run in a batch mode, but rather in a very interactive mode, communication is critical. Access Grid, VRVS/EVO and commercial H.323 have been used to support voice and video over the Internet either person to person or group to group. When a leader of an experiment is off-site, their video is often placed on the display wall to bring a greater sense of presence into the control room.

The tools discussed above create an engaging and capable off-site environment, including streaming of control room audio and video, streaming status information and shared displays and applications. To support the collaborative data analysis and documentation during the operations, emerging web-based collaboration tools such as wiki and multi-user logbook are being deployed. The main purpose is to allow scientists to write and edit research analysis easily and collectively about ongoing experiments regardless of the geographical location. To support the required interpersonal communications, new approaches to telecommunications, including those based on Voice over IP (VoIP) and Instant Messaging (IM) technologies, are being tested. The aim is to enable interpersonal contact based on identity, location or role and to integrate a wide variety of media and data into a seamless communications fabric.

Existing tools work well for formal communications, but have not been effective for fostering the ad-hoc interactions that are critical to successful collaboration. Future experiments will be even more dispersed than the ones in service today. The main barriers to informal communication at a distance are the lack of good information about how busy someone is and what they are working on, and a lack of integrated tools that allow for communication with a range of urgency and intrusiveness. Tools to address both of these issues exist in modern VoIP telephony. Traditional presence tools combined with the ability to use a range of media, from instant messaging, email, telephone, and video go a long way towards lowering these barriers. These collaboration tools can be embedded into the web based run management and electronic logbook applications that most modern experiments are utilizing [8].

Recently, work has moved beyond remote operation to remote control. The Experimental Advanced Superconducting Tokamak (EAST) located in China is the world's first tokamak with all superconducting coils. It will contribute to the success of ITER in many ways including exploration of steady-state plasma operation in advanced confinement regimes. As part of an active collaboration between the U.S. and China, the digital PCS mentioned earlier has been installed on EAST and U.S. scientists were actively involved with the successful attainment of first plasma in late 2006 [9]. This collaboration included personnel both on and off-site and rudimentary control of EAST plasmas from the U.S.

A FUTURE VISION TOWARDS ITER

ITER is a nuclear licensed device and, therefore, its operation must be highly secure. In such an environment, capabilities for remote participation must be carefully examined. We note however, that even on site, the ITER control room will be physically separated from the machine building. From the machine's perspective, the physical location of a control signal's origin is immaterial. Before a control signal is transmitted to the machine, it will need to be verified from a safety standpoint. Therefore, whether the control is occurring nearby or halfway around the world is immaterial as long as the originator's identity and the integrity of the bits can be assured.

For ITER's successful operation, we envision enhancing the data handling, data analysis, and interpersonal communication performed today and adding to it a component of sending control signals to ITER from remote locations. To successfully accomplish such a task a remote gatekeeper needs to be implement that not only performs the required security checks but also performs checks on the validity of all commands. The preliminary design of the ITER computer systems describes such a gatekeeper.

To support a remote control room [10] and worldwide collaboration [11] for ITER is a very challenging task. Although much of the required functionality exists in individual pieces and are used on today's experiments, what ITER will require is a fully integrated solution that presents a unified working environment. Such an environment must present to the remote scientist an immersive environment that is every bit as productive as what is available onsite.

As the fusion community works toward a solution for ITER, collaborations with other communities who are attempting to solve similar problems should not be ignored. Closest to fusion, in terms of a scientific research project is the High Energy Physics community and the Large Hadron Collider (LHC) at CERN. Their need to support worldwide remote operation with near-real-time access to mission critical controls data is very similar to ITER's requirement. Additionally, the broader need to support remote operations of experimental facilities (distributed code development, computing, human interaction, and visualization) are also similar to ITER's requirements.

Although the operation of ITER will not commence until the middle of the next decade, the time for testing on present devices the general principals and system architecture required for remote operation and control is now. The complexity of the problem is substantially beyond what is presently addressed in today's experiments and therefore significant design, development, testing, and hardening must be accomplished before such systems are ready for use on ITER.

Work supported by the U.S. Department of Energy SciDAC Program and at General Atomics under Cooperative Agreement No. DE-FC02-01ER25455 and DE-FC02-99ER54512.

REFERENCES

[1] F. Perkins, D. Post, N. Uckan, M. Axumi, D. Campbell, N. Ivanov, N. Sauthoff, M. Wakatani, editors, "ITER Physics Basis," Nucl. Fusion, **39** (1999) 2137.

[2] R. Fonck, G. Cospy, R. Durst, T. Gibney, M. Thompson, and S. Paul, "Remote Operation of the TFTR BES Experiment from an Offsite Location," Rev. Sci. Instrum. **63** (1992) 4803.

[3] S. Horne, M. Greenwald, T. Fredian, I. Hutchinson, B. Labombard, J. Stillerman, Y. Takase, S. Wolfe, T. Casper, D. Butner, W. Meyer, J. Moller, "Remote Control of Alcator C-Mod from Lawrence Livermore National Laboratory," Fusion Technol. **32** (1997) 152.

[4] B.B. McHarg, T.A. Casper, S. Davis, D. Greenwood, "Tools for Remote Collaboration on the DIII–D National Fusion Facility," Fusion Engin. Design **43** (1999) 343.

[5] T. Fredian, J. Stillerman, "MDSplus Remote Collaboration Support-Internet and World Wide Web," Fusion Engin. Design **43** (1999) 327.

[6] D.P. Schissel, "Grid Computing and Collaboration Technology in Support of Fusion Energy Sciences," Phys. Plasmas **12** (2005) 058104.

[7] J.R. Ferron, A.G. Kellman, G.R. McKee, et al., "An Advanced Plasma Control System for the DIII–D Tokamak," Proc. 14th IEEE/NPSS Symp. on Fusion Engineering, San Diego, 1991 (Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey, 1992) p. 761.

[8] J. Stillerman, D. Baron, T. Fredian, M. Greenwald, H. Schulzirinne, "Communications Fabric for Scientific Collaboration," Proc. 6th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, Inuyama, June 4–8, 2007.

[9] B.G. Penaflor, J.R. Ferron, M.L. Walker, D.A. Humphreys, J.A. Leuer, D.A. Piglowski, R.D. Johnson, B.J. Xiao, S.H. Hahn, D.A. Gates, "Worlwide Collaborative Efforts in Plasma Control Software Development," Proc. 6th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, Inuyama, June 4–8, 2007.

[10] D.P. Schissel, "A Vision for a Collaborative Control Room for ITER," Proc. 6th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, Inuyama, June 4–8, 2007.

[11] M. Greenwald, D.P. Schissel, et al., "Visions for Data Management and Remote Collaboration on ITER," to be published in Proc. ICALEPCS05 (2005).