

The collaborative tokamak control room

D.P. Schissel for the National Fusion Collaboratory Project

General Atomics, P.O. Box 85608, San Diego, California 92186-5608, USA

Abstract

Magnetic fusion experiments keep growing in size and complexity resulting in a concurrent growth in collaborations between experimental sites and laboratories worldwide. In the US, the National Fusion Collaboratory project is developing a persistent infrastructure to enable scientific collaboration for all aspects of magnetic fusion energy research by creating a robust, user-friendly collaborative environment and deploying this to the more than one thousand US fusion scientists in forty institutions who perform magnetic fusion research. This paper reports on one aspect of the project which is the development of the collaborative tokamak control room to enhance both collocated and remote scientific participation in experimental operations. This work includes secured computational services that can be scheduled as required, the ability to rapidly compare experimental data with simulation results, a means to easily share individual results with the group by moving application windows to a shared display, and the ability for remote scientists to be fully engaged in experimental operations through shared audio, video, and applications. The project is funded by the USDOE Office of Science, Scientific Discovery through Advanced Computing (SciDAC) Program and unites fusion and computer science researchers to directly address these challenges.

Keywords: National Fusion Collaboratory; remote participation; tokamak control room, shared large display, grid computing

1. Introduction

ITER is a burning plasma magnetic confinement experiment, the next major step in an international program aimed at proving the scientific viability of controlled fusion as an energy source. As such, it will be the largest and most expensive scientific instrument ever built for fusion research. Unlike very large experiments in other fields, such as high-energy or nuclear physics that operate essentially in a batch mode, fusion experiments put a premium on near real-time interactions with data and among members of the team. Carrying out these activities in an international collaboration and on the scale of ITER is a technically demanding problem, requiring the provisioning of a working environment for off-site personnel engaged in experimental operations that is

every bit as productive as what is onsite [1]. Considering the scope of its mission and its increased pulse length, we can predict that ITER will generate significantly more data than the current generation of experiments, however, the greatest challenge will be the provisioning of systems for analyzing, visualizing and assimilating the data to support decision making during ITER experiments. ITER's scientific productivity will be inextricably linked to the capability and usability of its collaborative infrastructure. Such an effective infrastructure is required both for the success of the entire ITER project and will maximize the value of ITER to the home fusion programs as well.

The US National Fusion Collaboratory (NFC) project [2] was started with a long-term vision of creating a capability suite for an international device

(MS WORD TEMPLATE for Submission in Fusion Engineering and Design)

like ITER while at the same time testing and using these technologies to assist the present day US domestic program. Funded as a SciDAC Collaboratory pilot [3], the NFC Project builds on the past collaborative work performed within the US fusion community and adds the component of computer science research done within the USDOE Office of Science, Office of Advanced Scientific Computer Research. The project is a collaboration itself uniting fusion scientists and computer scientists from seven institutions to form a closely coordinated team.

The Project's vision is that experimental and simulation data, computer codes, analysis routines, visualization tools, and remote collaboration tools are to be thought of as network services. Users are to be shielded from software implementation details and instead allowed a sharper focus on the physics with transparency and ease-of-use being the crucial elements. In this environment, access to services is stressed rather than data or software portability. Significant progress has been made towards the project's long-term vision through the initial deployment of a Fusion Energy Sciences Grid (FusionGrid) offering computational, visualization, data storage, and remote collaboration services.

This paper reports on those aspects of the NFC project directly related to enhancing the fusion experimental science program through the development of the collaborative control room that combines Grid computing with collaboration technologies such as the Access Grid with application sharing and shared displays to dramatically improve the efficiency of experimental sciences. To be fully functional, the collaborative control room requires (1) secured computational services that can be scheduled as required, (2) the ability to rapidly compare experimental data with simulation results, (3) a means to easily share individual results with the group by moving application windows to a shared display, and (4) the ability for remote scientists to be fully engaged in experimental operations through shared audio, video, and applications.

2. FusionGrid security

FusionGrid security [4] employs Public Key Infrastructure (PKI) to secure communication on the Internet through the use of a public and private cryptographic key pair that is obtained and shared through a trusted authority. FusionGrid used the

X.509 certificate standard and the FusionGrid CA to implement PKI for secure communication. A scientist desiring to join FusionGrid generates a public/private key pair and applies to the FusionGrid CA for an X.509 certificate binding that public key with their name. The request goes to the registration authority who verifies their identity and the validity of their request. FusionGrid certificates are managed for the user by a myProxy online certificate repository securely installed at LBNL. In this system, the user's long-term certificate (private key+trust) is securely stored on the myProxy server and access to FusionGrid is done by issuing a *my-proxy-get-delegation* command along with their username and password. With this FusionGrid login, the myProxy server issues a short-term certificate that is used for authentication. By storing the users' long-term certificates on the myProxy server, users no longer have to manage their certificates, but instead delegate that task to a FusionGrid administrator.

Secure authenticated connections are accomplished using the Globus Toolkit. A user issues a *my-proxy-get-delegation* command once per day and types in the password for their private key. The single sign-on is accomplished behind the scenes using a short-lived proxy certificate derived from the user's long-term X.509 certificate. Centralized authorization of FusionGrid resources is accomplished through the Resource Oriented Authorization Management System (ROAM). This system allows a resource provider to implement either a simple or complex authorization policy using a web browser interface. System flexibility is maintained since the resource provider is allowed to either use existing permission levels or define their own as required. The system is implemented using an Apache-based web server and a single PHP (a recursive acronym for PHP: Hypertext Preprocessor) script. PHP is an open-source, server-side HTML embedded scripting language used to create dynamic Web pages (e.g. search results from a database). A dynamic Web page is a page that interacts with the user, so that each user visiting the page sees customized information. A PostgreSQL database is used to manage all the authorization information. Access to this information is via secure HyperText Transport Protocol (HTTPS) using either the user's certificate, if present, or by a myProxy login. Resources also check for authorization using HTTPS communication.

(MS WORD TEMPLATE for Submission in Fusion Engineering and Design)

3. Data and computational services

Data access on FusionGrid has been made available using the MDSplus data acquisition and data management system [5] combined with the relational database Microsoft SQL Server. Based on a client/server model, MDSplus provides a hierarchical, self-descriptive structure for simple and complex data types and is currently installed and used in a variety of ways by about 30 experiments, spread over 4 continents. MDSplus and the Globus Toolkit have been combined to create secure X.509 certificate-based client/server data access on FusionGrid using the standard MDSplus interface without any loss in speed or functionality. SQL Server is securely accessible via MDSplus since a production of release of Globus for Windows is not available. Presently, the three main MDSplus experimental data repositories at Alcator C-Mod, DIII-D, and NSTX are securely available on FusionGrid.

Data management by MDSplus of large datasets generated by simulation codes is presently being tested using results from NIMROD simulations. NIMROD is a 3D MHD simulation code that runs on very large parallel computers. Using the MDSplus server at DIII-D, output from NIMROD runs up to 100 GB have been stored and served to users for further data analysis and visualization. Although successful, this storage methodology proved to be inefficient. The installation of an MDSplus server on the NERSC LAN along with the high-performance computational servers has been undertaken to investigate increased throughput capability. Parallel network data transport are also being investigated in order to overcome TCP/IP flow control limits for high bandwidth, high latency connections. A satisfactory solution for efficient storage of simulation datasets is required to allow for rapid comparison of experimental and simulation datasets within the control room.

The code TRANSP, used for time dependent analysis and simulation of tokamak plasmas, was released as a service on FusionGrid in late 2002 along with supporting infrastructure development (data storage, monitoring, user GUI) [6]. This FusionGrid service has been so successful that it has become the production system for TRANSP usage in the US and is starting to be adopted internationally. Running on a Linux cluster at PPPL, over 5800 TRANSP runs from ten different experimental machines have been

completed within the FusionGrid infrastructure. European scientists use TRANSP on FusionGrid with approximately 40% of the runs performing analysis on data from European machines. This approach has drastically reduced the efforts to support and maintain the codes which were previously required of the developers and by users' sites. Recently the GATO ideal MHD stability code was released as a FusionGrid computational service running on a Linux computer at General Atomics. Following the same design as the TRANSP service, the time required to deploy GATO on FusionGrid was minimal. The result has given confidence that the design of FusionGrid will scale to the deployment of many services. Additional fusion physics codes will be released as FusionGrid services in the next year based on the number of users and their geographic diversity.

The deployment of FusionGrid's computational services has been initially for analysis of data after the experimental day has ended. It has been a long-term goal to try and use these services to perform between-pulse data analysis in support of the experimental day. Recent work has shown that this goal is possible and has the potential to dramatically increase the quantity of available information from which the experimental team basis their decision on the next pulse [7]. Agreement-based interactions on FusionGrid utilizing the Globus toolkit 3 (GT3) enabled fusion scientists to negotiate end-to-end guarantees on execution of remote codes between experimental pulses. This was demonstrated by having the FusionGrid TRANSP service run on the PPPL cluster to support DIII-D operations in San Diego. In preparation for the demonstration, significant work was done to reduce TRANSP run production time, through both software and hardware changes, to about 6 minutes, which was found to be acceptable for an experimental run. The actual TRANSP run execution was slightly over 3 minutes; the balance of the time was due to network transfers. These data transfer delays will be reduced through further optimization of the software.

In the first between pulse data analysis using TRANSP, only one timeslice of the experimental data was run. When implemented routinely within the next year, this capability will give scientists power balance and energy transport information that was previously not available between pulses. In principle it would be best to run a fully time-dependent

TRANSP simulation. At some point in the future if the TRANSP code is parallelized, this reservation system will scale to large runs on multi-node computational clusters. Furthermore, this computer science infrastructure has the potential to allow any FusionGrid computational service to support between pulse tokamak data analysis.

4. Shared display walls

Shared display walls were installed in the control rooms of the 3 largest US fusion tokamaks (C-Mod, DIII-D, NSTX) for enhanced large-group collocated collaboration (Fig. 1). The systems offer a high pixel count for displaying detail rich visual information that normally does not fit on one computer display [8]. For collocated scientists, the ability to share a scientific graph on the display wall results in easier and more rapid discussion compared to the old model of individuals gathering around a small computer monitor. The old model takes more time and accommodates fewer people than the large shared display. The ability of the shared display to accept video combined with audio has allowed remote scientists to interact and even lead an experiment.



Fig. 1. The control rooms of NSTX (a), DIII-D (b), and C-Mod (c) with shared display walls being used to enhance collocated collaboration. On the DIII-D display is also video from remote collaborators in Europe who were participating in that days experiment.

The project has created and deployed unique software that presents the scientists with a multi-user environment allowing them to simultaneously share data to the large display and simultaneously interact with the display to edit, arrange, and highlight

information. Examples of other unique visualization tools created for the shared display wall environment include the plasma shape movie player that presents a near real time visualization of the plasma's magnetic field topology. Records of formal discussions during an experiment are recorded in an electronic logbook. The electronic logbook ticker application reads the newest log data and displays it as horizontally scrolling text across the display wall. Using the shared display in this way saves display space on the individual's desktop, reduces network and server traffic, and allows the person just entering the control room to be rapidly brought up to date. The Data Analysis Monitoring system is used to monitor the many hundred between pulse data analysis routines that are necessary for a successful experiment. A new color-coded interface was created for the shared wall to rapidly alert the staff to any analysis errors.

Although the main purpose of the shared display is to support the interaction among scientists collocated in the control room, the shared display plays an important role in enhancing the collaboration between control room and remote scientists. For example, the video imagery of remote collaborators (see below) have been presented on the shared display when they have lead the DIII-D experiment. This has made the collaboration and interaction more natural, as the large video image has given the local impression that the remote scientist is sitting in the control room.

5. Remote participation

The Access Grid [9] is used by FusionGrid to create a service that enables group-to-group interaction and collaboration that improves the user experience significantly beyond teleconferencing. The Access Grid includes the ability to utilize for scientific research a complex multi-site visual and collaborative experience integrated with high-end visualization environments. Developed exclusively for a FusionGrid service, the personal interface to the Access Grid (PIG) and the Macintosh OS X version have been developed as low cost alternatives to a conference room size AG node. The Virtual Room Videoconferencing System (VRVS) [10] has been used on FusionGrid as a web-oriented, low-cost, bandwidth-efficient, extensible means of videoconferencing and remote collaboration. The small footprint of VRVS has allowed its usage from areas of limited infrastructure that has proved very valu-

able. Since March 2004 at DIII-D, VRVS has been accessed 3480 times by 341 unique users (Fig. 2).



Fig. 2. (a) A scientist from his institution in San Diego utilized FusionGrid services (AG, VRVS, and MDSplus) to be the Scientific Coordinator for the EFDA-JET experiment in England. Audio, video, shared data, and shared applications create the beginnings of the collaborative control room. (b) Enlargement of the video from the JET control room. (c) Enlargement of the data traces from JET.

Anticipating a need to use even less local infrastructure, a customized version of text Instant Messaging utilizing Jabber has been tested to provide a lightweight collaboration channel for fusion research [11]. Jabber is used not only for exchanging text messages between scientists but also for automatically posting textual information, triggered by MDSplus events, from the tokamak (pulse state), the between-pulse data analysis system (DAM), and from the electronic logbook. Data visualizations from the tokamak are automatically posted via hyperlinks in the Jabber client that invoke a CGI script to access the latest data from MDSplus.

6. Future work

Although the NFC project has made significant progress there is still a substantial amount of work that needs to be accomplished before an international device like ITER can be supported. First and foremost, the services and capabilities being deployed by FusionGrid need to be combined into one unified working environment. Presently, individual pieces are available to support the scientists but they are not

presented in one integrated framework. Emerging technologies for interpersonal communications based on Voice Over IP (VoIP) can possibly facilitate the convergence of physical and logical communications channels so that phone, audio, video, email, messaging, and data can be integrated into a common framework. Web-based directory services should be provided allowing people and data streams to be identified, located, scheduled and connected into a flexible communications fabric.

Second, the international ramifications of grid security (e.g. X.509) versus site security (firewalls) need to be understood and solved so that an organization such as ITER can truly exist as a virtual entity spread throughout the world. The ability to combine a physical token (e.g. SecureID) with X.509 certificates allowing for dynamically configurable site firewalls has shown early technical promise as well as the potential for being accepted within the international community.

Third, computational grid services need to be able to be scheduled as required to support an operating experiment to deliver analyzed data back into the control room within a guaranteed time frame. Individual pieces of this have been under investigation by a variety of projects (Network QoS, grid computational scheduling, large-scale data management, visualization, CPU scheduling) but they need to be unified for production use. The ultimate goal would be to have a massive supercomputer support quasi real-time analysis of fusion experimental data. Finally, a data management solution that embodies simulation and experimental datasets and satisfies the real-time needs of the control room must be deployed. Extending MDSplus using the underpinnings of GridFTP for parallel data transport has shown initial promise. Combining parallel I/O with caching and metadata management whether through MDSplus or a new system will be required.

These concepts are unified into what is called the collaborative control room and the resulting software infrastructure must be robust as down-time or wasted experimental plasmas will not be tolerated in a facility where it will cost ~\$1M for one plasma pulse when amortized over the life of the project.

7. Conclusions

The roadmap for developing these collaborative technologies can continue to involve testing and

deploying with the present US experimental community. Yet steps need to be taken beyond these borders as has begun with the present NFC Project. More direct involvement with existing facilities such as those in England (JET), Germany (ASDEX-U), France (Tore Supra), and Japan (JT-60U) will be clearly beneficial. Yet two new facilities that will begin operation in the next few years are in Korea (KSTAR) and China (EAST). These newer facilities, not burdened by existing infrastructure, can be excellent proving grounds for new ideas and will expand our work to two more ITER partners. Deploying new computational services in these countries will expand scientific collaboration and test new grid security techniques. Having some of these new computational services support tokamak operation will provide a proving ground for the variety of new technologies required to make this a success.

With the worldwide focus on ITER as the next generation machine, its success requires advanced remote collaboration capability. This capability and success needs to include more than just the experimental physics program. The final design, engineering, and construction phases will be worldwide collaborations as well and although they will not need the collaborative control room they will need the ability to richly interact with their distant colleagues. The ability to interactively share engineering drawings, conduct design reviews, and view 3-D mockups of machine components are all clearly required. Imagine an electronic tabletop display (Hitachi recently announced such a product), one can envision a drawing “rolled-out” electronically on a table. The designer picks up her pen and begins “pointing” to different areas of the drawing as the ad-hoc meeting begins during ITER construction to solve an unforeseen problem. Her counterparts spread around the world have a similar table, see the drawing and her “pointing,” and when they look up into the accompanying display device they see images of their colleagues and hear their words.

Our vision is that when ITER operation begins, the collaborative control room will allow scientists to share data and knowledge as readily as their engineering counterparts. Large shared displays will present information to the assembled team in the control room. Off-site scientists will share the results of their analysis to the large display as well as to individual small displays of their colleagues. ITER’s integrated data acquisition and data management system not

only allows for simultaneous data availability worldwide, but also the automatic starting of many data analysis tasks. Utilizing computational Grids, these tasks are dispatched to computer systems located worldwide and their results are rapidly integrated back into the ITER data management system. Some of these tasks are run on very large state-of-the-art supercomputers requiring network quality of service and CPU scheduling to support the time critical data analysis environment of the control room. Utilized in this way, the world’s power computer systems run data analysis and simulations in concert with the largest fusion device to advance fusion science in ways not previously envisioned.

Due to the complexity of these problems and the importance of reaching a satisfactory solution, design work and testing needs to start early in the ITER construction phase. Present day fusion energy science research provides an excellent proving ground for research to support the needs of ITER. Although ITER’s first plasma is over a decade away, the intervening time should be used to continue to develop the technology outlined in this paper.

Acknowledgments

This work was supported by US Department of Energy SciDAC program and under Cooperative Agreement DE-FC02-01ER25455.

References

- [1] D.P. Schissel, et al., Phys. Plasmas **12**, 058104 (2005).
- [2] D.P. Schissel, et al., Fusion Eng. & Design **71** (2004) 245, <http://www.fusiongrid.org>
- [3] <http://www.scidac.org/>
- [4] J.R. Burruss, et al., Security on the US FusionGrid, this conference.
- [5] T.W. Fredian, et al., Fusion Eng. & Design **60** (2002) 229.
- [6] J.R. Burruss, et al., Fusion Eng. & Design **71** (2004) 251.
- [7] K. Keahey, et al., Future Generation Computing Systems, Vol. 18 (2002) 1005.
- [8] G. Wallace, et al., to be published in IEEE Computer Graphics & Applications, Special Issue on Large Displays (2005).
- [9] L. Childers, et al., Proceedings of the 4th International Immersive Projection Technology Workshop (2000), <http://www.accessgrid.org>
- [10] D. Adarnczyk, et al., Proceedings of the Conference for Computing in High Energy and Nuclear Physics (2003), <http://www.vrvs.org>
- [11] G. Abla, this conference.