GA–A25814

A VISION FOR A COLLABORATIVE CONTROL ROOM FOR ITER

by

D.P. SCHISSEL

MAY 2007

GENERAL ATOMICS

DISCLAIMER

GA–A25814

A VISION FOR A COLLABORATIVE CONTROL ROOM FOR ITER

by

D.P. SCHISSEL

This is a preprint of a paper to be presented at the 6th IAEA Technical Meeting on Control, Data Acquisition, and Remote Participation for Fusion Research, Inuyama, Japan, June 4–8, 2007 and to be published in the *Proceedings.*

Work supported by
the U.S. Department of Energy
under DE-FC02-01ER25455

GENERAL ATOMICS PROJECT 30106
MAY 2007

GENERAL ATOMICS

ABSTRACT

A vision for an onsite control room for ITER that will support worldwide experimental collaboration and operation is presented. Fusion experiments place a particular premium on near realtime interactions with data and among members of the team. Enabling effective international collaboration on this scale is technically demanding, requiring powerful interactive tools and provision of a working environment for offsite personnel engaged in experimental operation that is every bit as productive as what is onsite. Expanding the view of the control room to include worldwide realtime resources, both computational, data, and human, allows for a collaborative design that will significantly benefit ITER's scientific productivity. While the worldwide fusion program has a significant track record for developing and exploiting remote collaborations, the community should recognize that the collaborative needs of other communities share some similarity and therefore joint or shared research into collaborative technologies will increase the benefit to all concerned.

1. INTRODUCTION AND CHALLENGES

The future visitor to the ITER control room will most likely observe a large room with individual and large wall computer displays being used by approximately 100 individuals to operate the world's largest, most advanced, and expensive scientific instrument ever built for fusion research. The naive visitor might assume that the scientific team leading the day's operations forms in the morning, works for the day, and then disperses at night. In reality, on ITER just as on today's fusion experiments, scientific teams coalesce well before they walk into the control room. To create a day's experimental plan, scientists review previous data, create a scientific hypothesis, debate ideas on how to best test the hypothesis, run theoretical simulations to test these ideas, write a detailed plan for the day's experiments, work with the plasma control team to insure that the plan can be realized, have the plan reviewed by peers, and then modify the plan as required. It is only at this point, with the plan in hand, that the team is ready to walk into the control room and perform the day's experiment.

Therefore, when the ITER control room is discussed, and the discussion turns to making it a collaborative enterprise allowing remote participation, the discussion needs to include the entire process, not just the day's events. A remotely distributed group of scientists must jell early on in the process that leads to the creation of the team that develops the experimental plan.

Returning to the control room, ITER will require near realtime interactions with data and among members of the team allowing for efficient between-pulse analysis, visualization, and decision making. This mode of operation places a large premium on rapid data analysis that can be assimilated in near real time. However these requirements align very well with the requirements of the team before they enter the control room. Activities such as debating scientific ideas and comparing results of data analysis are performed while creating a scientific plan; they are just not done in such a time critical environment. The one activity that is unique to the precontrol room work is the usage of very large-scale plasma simulations. It is clear that massive amounts of data analysis including simulations will be performed between ITER pulses. However, it seems unlikely that the largest whole device modeling simulations that will be running on petascale computers for experimental planning will also be used between pulses. The approximate 1 h between ITER pulses will not provide sufficient time for the largest simulations. Therefore, if the collaborative requirements of the control room are satisfied, including access to the largest scale simulations, the requirements of the entire process will be satisfied.

Carrying out these activities in an international collaboration and on the scale of ITER is a technically demanding problem, requiring a working environment for off-site personnel that is every bit as productive as what is onsite. One of the greatest challenges will be the provisioning of systems for analyzing, visualizing and assimilating the data to support decision making in support of 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 for the success of the entire ITER project and will maximize the value of ITER to the home fusion programs as well.

2. POLICY AND INFRASTRUCTURE

Implicit in the design of the collaborative control room are aspects associated with policy and infrastructure more than computational or collaborative capability.

First is the realization that the remote participant is as valuable as the local scientist. Therefore, ITER's policy must be to allow an individual scientist to retain the same rights and privileges as they move from on site to working remotely. A scientist should not be *de facto* penalized for being at a remote site.

ITER's data archives need to contain all raw, processed, and simulation data and should be available in a timely manner to all members of the scientific team. Thus, data analysis performed by off-site researchers should be written back into the ITER repository. For complex and highly coupled systems like ITER, the scientific team and program must not be fragmented or impeded by data access limitations.

Clearly the opening of the ITER control room and data repositories to off-site collaborators will require sufficient security to ensure protection of these valuable resources. However, at the same time, security must not be so onerous as to restrict the ability of remote scientists to contribute. ITER's cyberspace security needs to allow users from administratively and geographically distributed organizations to access resources (data, codes, visualizations, people). This imposes the need for users to have a unique cyber-identity and a means by which they can authenticate themselves as that entity at each site. Additionally, the user needs to be able to use a unique passphrase once, referred to as single sign-on and have authentications at other sites derived automatically from a limited lifetime token granted by sign-on. Additionally, the security system must enable multiple stakeholders for a single resource to set access policy for that resource.

Network connectivity to the ITER site will be an enabling technology and must be sufficient to allow for fully active remote participation. Estimates show that ITER will require network connectivity from the site in multiples of 10 Gb/s to fully support remote scientific collaboration [1]. Even with this network connectivity, ITER's local and wide area network should be able to support quality of service (QoS) guarantees so that outgoing and incoming time critical data is available to support experimental operations. Additionally, ITER will most likely need to adopt IPv6 due to the limitations of address space in IPv4 [2]. A number of methods have been invented to circumvent this limitation, with network address translation (NAT) being the most common. But NAT breaks the IP architectural model that assumes globally unique address for each host and breaks peer-to-peer connectivity symmetry required for many applications that support collaborative computing. The presence of NATs in the communication path hinders the ability to encrypt traffic at the network layer, which would otherwise provide security that is transparent to the application.

3. COLLABORATIVE CONTROL ROOM

The vision for ITER's onsite control room is one that supports worldwide experimental collaboration and operation. The concept is modular, in that any function performed onsite can also be performed remotely. Therefore, our vision supports one onsite control room and an arbitrary number of remote control rooms. The vision pushes in two directions. First, to view applications and data as services to be run on geographically dispersed resources thereby increasing the amount and detail of both analysis and simulation results that are made available to the scientific team. Second, by bringing in expertise from geographically remote teams of experts, the depth of interpretation can be increased leading to improved assimilation of those results. The following five subsections define the collaborative control room.

3.1. SECURE COMPUTATIONAL SERVICES

The model of computation for ITER is that of an application service provider (ASP) where computer-based data analysis and simulation services are provided to end-users over the wide area network [3]. This model does not preclude local scientists from writing and running their own applications locally. In fact, such codes can be offered in an ASP model to fellow local scientists as well as those stationed remotely. What the model does eliminate is the requirement for remote scientists to bring their application to the ITER site for it to be used. Instead, the scientist locally maintains the code as well as the resources that the code would execute on. This mode of operation has advantages for both the client and the provider: it frees the clients from maintaining and updating software and the providers from porting and supporting it on a wide-range of platforms. In this environment, access is stressed rather than data or software portability. In some cases, the ability to bring the code to ITER is not feasible. Imagine an integrated burning plasma simulation code that operates on one of the world's fastest supercomputers. Portability is not an option, yet by using the ASP model, the code and computer systems can be made available to the distributed ITER team.

This type of computing paradigm, the decoupling of production and consumption, is often associated with grid computing [4]. However, our usage of the term does not refer to computer cycle scavenging or distributed supercomputing. The proposed computational grid has been used in a variety of scientific disciplines and in the United States to support fusion research on FusionGrid [5]. The codes TRANSP, GATO, and ONETWO operate as FusionGrid computational services under the ASP model. The code TRANSP, used for time-dependent analysis and simulation of tokamak plasmas, was the first service released on FusionGrid and is presently used worldwide. As expected, this approach has drastically reduced the efforts to support and maintain the codes which were previously required of the developers and by users' sites.

The ASP model is appropriate for a wide variety of computational problems including those requiring a large supercomputer. Therefore, this model is appropriate for the vast majority of computational analysis required by ITER whether that be pre-experimental data analysis or between pulse analysis.

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 base their decision on the next pulse [6]. It should be noted that between-pulse analysis utilizing the ASP model requires the ability to guarantee computational time on demand as well as guaranteed network quality of service over the wide area.

The ASP model requires a strong security implementation that yields a reliable means of identification (authentication) and the ability to control access to resources (authorization). This security model, although discussed within the computational services section, applies to and works equally well with all aspects of the collaborative control room. The ITER project will cross numerous administrative boundaries with its resources such as networks and computers not owned or controlled by a single project or program. Nevertheless, ITER's security solution needs to provide for single sign-on. Users of present-day computer systems are familiar with signing onto a computer with a username and password yet if they move to another computer they need to re-authenticate. In a distributed computational environment, where computer systems are acting on behalf of the user, multiple authentication requests are not feasible.

A common solution to authentication in distributed environments is the usage of Public Key Infrastructure (PKI) to secure communication through the use of a public and private cryptographic key pair that is obtained and shared through a trusted authority [7]. The X.509 certificate standard along with a trusted Certificate Authority (CA) is one method to implement PKI for secure communication [8]. This security methodology is successfully utilized in a number of scientific grid projects. In some instances, certificates are managed for the user by a myProxy online certificate repository [9]. 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 an administrator. By using their username and password only once, the user is granted a certificate and this can be used by themselves, or by an application acting on their behalf, to reliably authenticate any number of times to any system in the world. Security could be a enhanced even further by requiring the usage of hardware-based one-time password (OTP) in concurrence with the username and password.

Once authenticated, the security system needs to determine what permissions a user or computer system are allowed for a particular resource. Whereas a central authority can accomplish management of certificates, authorization management of distributed resources is best accomplished in a distributed fashion since stakeholders need to be able to control access to their resources. An example authorization system in the fusion community is FusionGrid's Resource Oriented Authorization Management (ROAM) System [10]. 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. A limitation of ROAM is its inability to handle a scientist's dynamically changing roles. For ITER, since each experiment may have different key personnel, it is a requirement that permissions may change daily as different personnel take on different roles. Therefore, dynamic application of role-based authorization is a requirement.

3.2. COMPARE EXPERIMENTAL AND SIMULATION DATA

Data, just as a computational code, can be considered as a service to be provided. As such, data can be distributed and made available through the ASP model utilizing the same security mechanism previously discussed. Data should therefore be service rather than file oriented with accessibility being more important than portability. The data system should provide a coherent view of all data such as raw and processed as well as ancillary (calibrations, geometry, etc.) data. Metadata catalogues should allow for rapid searching such that the entire data system can be efficiently browsed and desired data can be located and collected [11].

The interface to the data system would be a simple single API and be the same for both on site and remote scientists. Data analysis, no matter where it was performed, will be written back into the data repository so that the entire ITER organization benefits. The ITER data repository will therefore also include simulation data. This coherent view and efficient browsing of all data will allow for rapid comparison of experimental with simulation data satisfying the requirements of those inside and outside of the control room.

The worldwide fusion community has significant experience in data management due to the numerous experimental facilities. One of the most common systems in use on over 30 experiments worldwide is the MDSplus data acquisition and data management system [12]. MDSplus has been combined with PKI security resulting in a secure X.509 certificate based client/server data access [13]. The usage of MDSplus as required by ITER gives credence to the proposed architectural design yet MDSplus itself in its present form is not sufficient for ITER's requirements.

With the data in hand, the actual comparison of data through visualization can be performed using a wide variety of methodologies. In today's experiments, there are usually one or two main visualizations programs maintained by a core group of software developers associated with the experiment. Since the data interface is an available standard, others are free to, and often do, create their own visualization applications. For ITER, it would be expected that a similar situation would arise with the possible exception of the need to remotely share visualizations (discussed below). In that case, the visualization application might be customized beyond the typical product produced by a member of ITER's scientific team.

3.3. SHARE INDIVIDUAL RESULTS WITH GROUP

Tokamak control rooms today have on the order of 20 to 30 scientists and engineers working in a coordinated fashion to safely and efficiently conduct physics experiments. Typically, these individuals interact verbally while within speaking range and visually by gathering around one small computer monitor to jointly examine data. These types of visual interactions are limited to a small number of people since they need to gather around a single monitor.

The ITER control room will be similar to today's control rooms but will be physical larger and used by on the order of 50 to 100 people. Utilizing the "gathering around the monitor" model that can at best accommodate 10% of the control room population is not acceptable. Instead, ITER will need to utilize large shared display walls for greatly enhanced collaboration within the ITER control room. Current research on tiled display walls is mostly aimed at utilizing commodity components to reduce cost and to visualize extremely high pixel count images [14]. However, for fusion and ITER, the benefit will not be in the ability to show one large pixel count dataset but rather many smaller ones (Fig. 1). Tokamak experiments do acquire a lot of data, but it is divided into many different measurements. A large-scale display system allows a large group in a collaborative fashion to assimilate data, much like a jigsaw puzzle, to tell a coherent physical story. The large collocated group will be able to work more efficiently as a team.

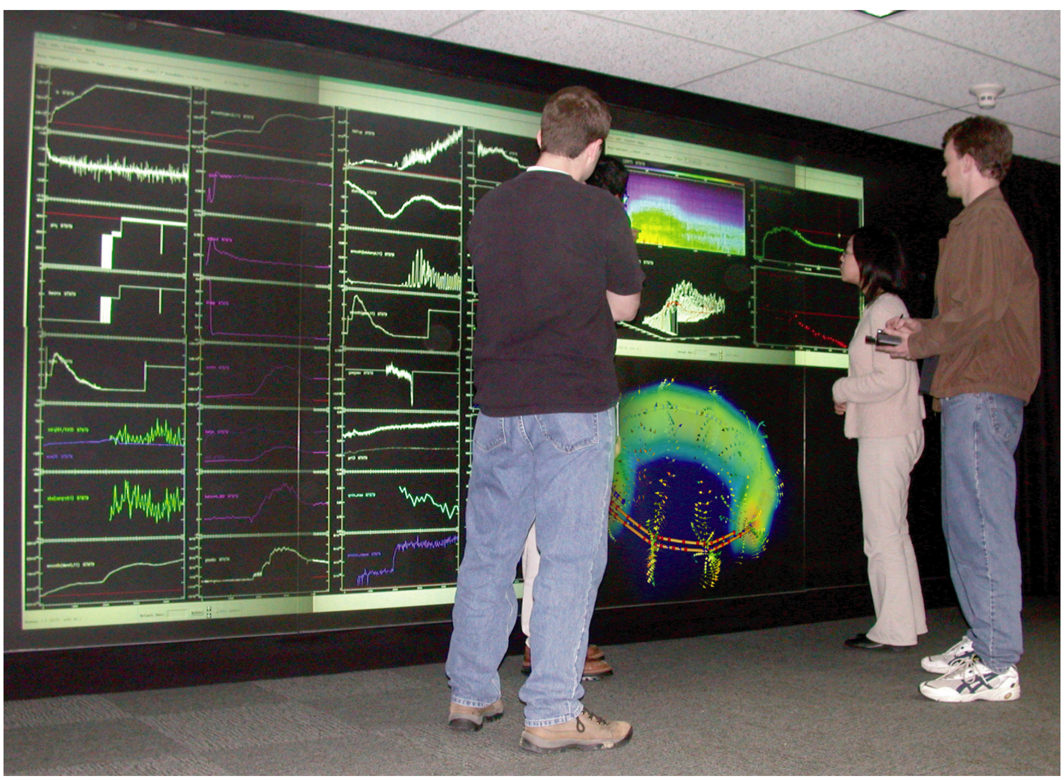

Fig. 1. An eight-projector, 10.5 MPixel tiled display wall at General Atomics is used to visualize and share data from the DIII-D National Fusion Facility.

Display walls of various size and shape have been placed in the control rooms of C-Mod, DIII-D, and NSTX [15]. Utilization of these systems has been aided by the deployment of unique software that presents a multi-user environment allowing for the simultaneous sharing of data to the large display and simultaneous interaction with the display to edit, arrange, and highlight information.

3.4. FULLY ENGAGED REMOTE SCIENTIST

To fully utilize the worldwide talent available for ITER, the control room has to allow for distributed collaborative problem solving. Ad hoc interpersonal communication is today, and will be for ITER, a fundamental requirement for successful operation. The ability to perform efficient ad hoc distributed scientific discussions should be an ITER requirement.

Videoconferencing, used frequently in the business world, has really solved the problem of holding distributed formal meetings such as technical seminars or management gatherings with shared audio and video [16]. Additionally, presentations are either distributed beforehand or shared in real time. But in such a meeting, if an audience participant wants to show a different version of a particular slide, the collaboration system typically breaks down. Yet, this type of ad hoc interaction is fundamental to the scientific process.

The ITER scientist, no matter where they are located, needs to have a simple answer to the question: “how do I show you this graph?” But the need goes farther, since scientific discussion often involve more questions about the data and how it is being interpreted. The graph the scientist desires to share must be interactive, both in changing the data’s representation as well as through annotations.

A number of systems have been created to try and solve the problem of remote interactive discussions involving audio, video, and shared applications. Two commonly used systems today are the Access Grid [17] and the Virtual Room Videoconferencing System (VRVS) [18] that are attempting to enable group-to-group interaction and collaboration that improves the user experience significantly beyond teleconferencing. Both systems have met with success yet are plagued by reliability issues (networks, firewalls, etc.) that limit their usability in the time critical environment of a tokamak control room.

Although the main purpose of the shared display discussed above is to support the interaction among scientists collocated in the control room, the shared display can also play an important role in enhancing the collaboration between control room and remote scientists. For example, the video imagery of remote collaborators can be presented on the shared display to make the interaction more natural, as the large video image gives the local impression that the remote scientist is sitting in the control room. The shared display can also allow the remote scientist to share visualizations with a large group of scientists in the control room [19].

ITER will also have the problems of time zones; it will be the project on which the sun never sets. With three team members, one each in Europe, Asia, and North America, one of them will always be awake. Thus, ITER will have to address the problem associated with persistent collaboration. That is, a distributed collaboration that persists over time. For example, the North American team member needs to be able to share the results of her day's work so that her European colleague knows where to begin.

3.5. PLASMA CONTROL

ITER's versatile plasma shaping, feedback, and control system will increase the flexibility of the device yet also increase the complexity of designing experiments. Therefore, plasma control experts will be integral members of the experimental team. Plasma control must, therefore, be considered within the context of remote collaboration.

Since ITER is a nuclear class device, its operation must be highly secure and the consideration of remote control might seem a bit preposterous. Yet, it makes no sense to have a remote scientist work for weeks preparing details of plasma pulse control only to be denied remote access during the actual experiment. Recent experience [20] on EAST has shown that transmitting the display from a local computer back to the remote location is a very inefficient method of operation. The preference would be to transmit data over the Internet rather than X Windows.

Therefore, for ITER's control room to be fully remote, a qualified remote plasma control specialist needs to be able to run the plasma control system on their local computer and export discharge control setup data to the ITER site for authentication and verification. The security requirements for this type of operation are severe yet the X.509 certificates mentioned earlier can satisfy a majority of these needs.

4. DISCUSSION

The world‘s fusion community has a rich history of collaborative research and therefore has a significant track record upon which to draw. Many of the required functions exist in individual pieces and are in use in today’s experiments. Yet, what is needed for ITER is an integrated solution that presents a unified working environment.

As the ITER community works towards this unified framework, it makes sense to examine what is being done in other areas to leverage ongoing activities. Closest to fusion, in terms of being a scientific research project, is the High Energy Physics community (HEP) and the Large Hadron Collider (LHC) at CERN. Their need to support remote operation of experimental facilities will require collaborative efforts involving over 2000 scientists worldwide. The HEP community also has a rich history of developing collaborative technology including the initial proposal for the World Wide Web and the VRVS system mention earlier. The overlap between the requirements and backgrounds of FES and HEP appear substantial and therefore joint research should benefit both programs.

Another area where the fusion community might benefit from a closer collaboration is the Collaborate Development Environments (CDE) for software [21]. These are being created so that distributed teams of software engineers can collaborate to solve problems. There has been increased emphasis on CDEs due to the increase in software outsourcing. Although at first blush, the work of a software engineer and an ITER physicist may seem different, their needs in the collaborative and persistent arena are very similar.

Lastly, the driving forces of the commercial world need to be careful examined as things evolve. An example here is the continued convergence of telecommunications and computing technology (e.g., Voice over IP), where integrated audio, video, messaging, email, and streaming of data should aid in the development of a productive remote collaboration environment for ITER.

As the ITER project moves forward, the technologies being developed should be prototyped and tested on the current generation of experiments and numerical simulation projects.

REFERENCES

[1] M. Greenwald, D. Schissel, private communication.

[2] W. Matthews and L. Cottrell, Achieving high data throughput in research networks, Proc. Computing in High Energy and Nuclear Physics, Beijing, China 2001.

[3] K. Keahey, et al., Grids for experimental science: the virtual control room, Proc. Challenges of Large Applications in Distributed Environments, Honolulu, Hawaii, 2004

[4] I. Foster, et al., IJSA 15(3) (2001) 200-222.

[5] D. Schissel, et al., Phys. Plasmas 12 (2005) 058104.

[6] K. Keahey, et al., Future Generation Computing Systems 18 (2002) 1005.

[7] W. Diffie, et al., IEEE Trans. Inf. Theory IT-22 (1976).

[8] ITU-T X.509 (03/00), ITU, Recommendation, 2002.

[9] J. Novotny, et al., An online credential repository for the grid: MyProxy, Proc. 10th IEEE Intl. Symp. on High Performance Distributed Computing, San Francisco, 2001 (Institute of Electrical and Electronics Engineers, Inc., 2001) p. 0104.

[10] J. Burruss, et al., J. Grid Computing 4(4) (2006) 413.

[11] M. Greenwald, et al., ICALEPCS 2004.

[12] T. Fredian, et al., Fusion Engin. Design 60 (2002) 229.

[13] J. Burruss, et al., Fusion Engin. Design 81 (2006) 1949.

[14] G. Wallace, et al., Tools and applications for large-scale display walls, IEEE Computer Graphics and Applications 25 (2005) 24.

[15] G. Abla, et al., Shared display wall based collaboration environment in the control room of the DIII-D national fusion facility, Proc. Workshop on Advanced Collaborative Environments, Redmond, Washington, 2005.

[16] H.323 reference.

[17] L. Childers, et al., Proc. 4th Int. Immersive Projection Technology Workshop, Ames, Iowa, 2000.

[18] D. Adamczyk, et al. Global platform for rich media conferencing and collaboration, Computing in High Energy and Nuclear Physics, La Jolla, California, 2003.

[19] G. Abla, et al., Advanced tools for enhancing control room collaborations, Proc. 5th IAEA Tech. Mtg. on Control, Data Acquisition, and Remote Participation for Fusion Research, Budapest, Hungary, 2005 to be published in Fusion Engin. Design.

[20] B. Penaflor, et al., Worldwide collaborative efforts in plasma control software development, this conference.

[21] G. Brooch, Advances in Computers 59 (2003) 2.

ACKNOWLEDGMENT

This work supported by the U.S. Department of Energy under Cooperative Agreement DE-FC02-01ER25455. The author gratefully acknowledges years of fruitful discussions with fusion-plasma and fusion-computer scientists worldwide.