

ENHANCED COMPUTATIONAL INFRASTRUCTURE FOR DATA ANALYSIS AT THE DIII-D NATIONAL FUSION FACILITY

by

**D.P. SCHISSEL, Q. PENG, J. SCHACHTER, T.B. TERPSTRA,
T.A. CASPER, J. FREEMAN, R. JONG, K.M. KEITH, B.B. MCHARG, Jr.,
W.H. MEYER, and C.T. PARKER**

AUGUST 1999

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ENHANCED COMPUTATIONAL INFRASTRUCTURE FOR DATA ANALYSIS AT THE DIII-D NATIONAL FUSION FACILITY

by

**D.P. SCHISSEL, Q. PENG, J. SCHACHTER, T.B. TERPSTRA,[†]
T.A. CASPER,[‡] J. FREEMAN, R. JONG, K.M. KEITH, B.B. MCHARG, Jr.,
W.H. MEYER,[‡] and C.T. PARKER**

[†]Princeton Plasma Physics Laboratory, Princeton, New Jersey

[‡]Lawrence Livermore National Laboratory, Livermore, California

This is a preprint of a paper to be presented at the 2nd IAEA Technical Committee Meeting on Control, Data Acquisition and Remote Participation on Fusion Research, July 19–21, 1999, in Lisboa, Portugal, and to be published in *Fusion Engineering and Design*.

**Work supported by
the U.S. Department of Energy
under Contract Nos. DE-AC03-99ER54463, W-7405-ENG-48,
and DE-AC02-76CH03073**

**GA PROJECT 30033
AUGUST 1999**

ABSTRACT

Recently a number of enhancements to the computer hardware infrastructure have been implemented at the DIII-D National Fusion Facility. Utilizing these improvements to the hardware infrastructure, software enhancements are focusing on streamlined analysis, automation, and graphical user interface (GUI) systems to enlarge the user base. The adoption of the load balancing software package LSF Suite by Platform Computing has dramatically increased the availability of CPU cycles and the efficiency of their use. Streamlined analysis has been aided by the adoption of the MDSplus system to provide a unified interface to analyzed DIII-D data. The majority of MDSplus data is made available in between pulses giving the researcher critical information before setting up the next pulse. Work on data viewing and analysis tools focuses on efficient GUI design with object-oriented programming (OOP) for maximum code flexibility. Work to enhance the computational infrastructure at DIII-D has included a significant effort to aid the remote collaborator since the DIII-D National Team consists of scientists from 9 national laboratories, 19 foreign laboratories, 16 universities, and 5 industrial partnerships. As a result of this work, DIII-D data is available on a 24×7 basis from a set of viewing and analysis tools that can be run either on the collaborators' or DIII-D's computer systems. Additionally, a Web based data and code documentation system has been created to aid the novice and expert user alike.

1. INTRODUCTION

The DIII-D National Team consists of about 120 operating staff and 100 research scientists from 9 U.S. National Laboratories, 19 foreign laboratories, 16 universities, and 5 industrial partnerships. This multi-institution collaboration carries out the integrated DIII-D Program mission which is to establish the scientific basis for the optimization of the tokamak approach to fusion energy production. Presently, about two-thirds of the research physics staff are from the national and international collaborating institutions.

As the number of on-site and remote collaborators continues to increase, the demands on the DIII-D National Program's computational infrastructure become more severe. The Director of the DIII-D Program recognized the increased importance of computers in carrying out the DIII-D mission and has focused resources to satisfy this increasing demand. Work has concentrated on both hardware [1] and software improvements to increase the DIII-D data analysis throughput and data retrieval rate. The underlying philosophy (Fig. 1) behind these development efforts is uniformity, both in terms of the look and feel of graphical user interfaces, in terms of access methods to analyzed datasets, and access to existing computer power. This paper presents the long term plan and progress to date.

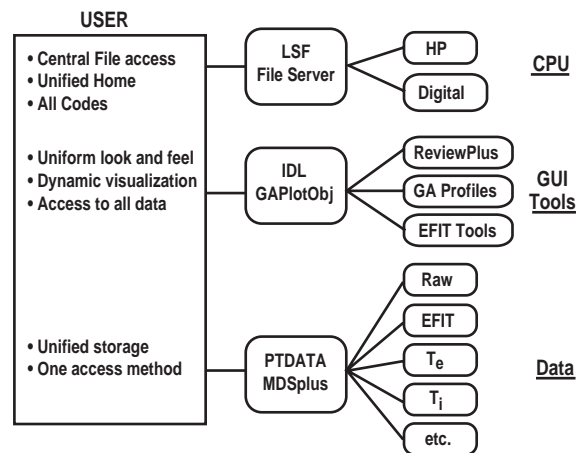


Fig. 1. Users benefit from the unification of a heterogeneous environment which includes access to computers, GUI tools, and data.

2. DATA STORAGE AND RETRIEVAL

The amount of raw digitizer data acquired per tokamak pulse has been increasing steadily since the start of the DIII-D Program. Presently a total of 1 TB of raw data, written into PTDATA format [2], has been acquired. A mass storage system with a total capacity of approximately 3 TB has been installed for storage of DIII-D raw digitizer data. Along with 100 GB of hard disk capacity, this system consists of a 600 GB HP 600fx Magneto-optical jukebox and a 2.2 TB ATL 7000 DLT tape library. The mass storage system provides data availability 24 hours a day and 7 days a week (24×7), which is demanded by a facility with collaborators in time zones that are up to 12 hours apart. The magnetic disk and Magneto-optical components interactively provide 700 GB of raw data that is the equivalent to approximately 9000 compressed tokamak discharges. The remaining 2.2 TB from the DLT system is available on a 3 to 5 minute time scale.

At the beginning of 1998 the MDSplus [3] data storage software was adopted for on-site storage of analyzed DIII-D data. Prior to MDSplus adoption, analyzed diagnostic data was written out in different formats and stored on a variety of computers (Fig. 2). The unified format of the Unix MDSplus system allows researchers to learn a few computer commands and read a vast amount of DIII-D data [4]. The diagnostic data stored in the MDSplus system continues to increase with presently 4700 archived shots representing 38 GB of analyzed data. The majority of MDSplus data is made available in between tokamak pulses giving the researcher critical information before setting up the next pulse. The benefits of Unix MDSplus are faster serving of data and easier integration into the DIII-D Unix analysis environment. Future storage plans include the integration into MDSplus of data from advanced simulation and modeling codes allowing for easy comparison to experimental data with existing visualization tools.

The increase in data acquired per year has been accompanied by an increased demand for data retrieval as the size of the DIII-D National Team has grown. This, in turn, has created a

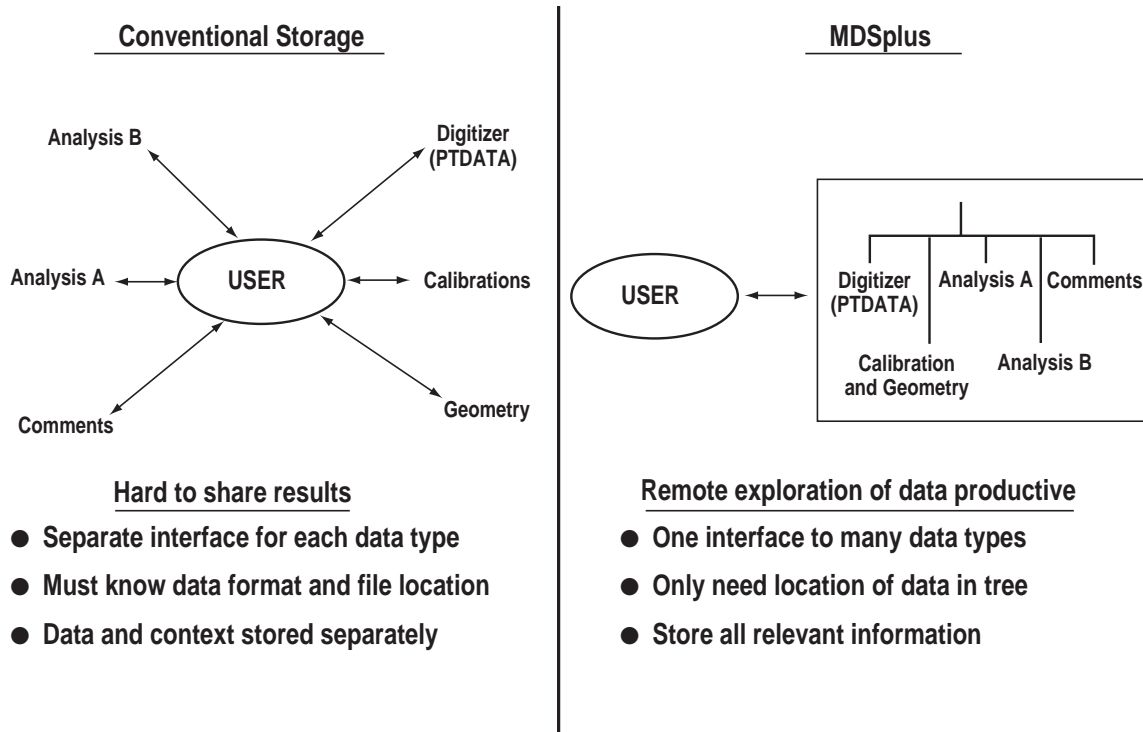


Fig. 2. MDSplus unifies the interface to analyzed data which is critical in a environment with a large number of off-site collaborators.

need for a substantial increase in network bandwidth. To satisfy this demand, the network within the DIII-D control room and to the offices was increased to switched Fast Ethernet at 100 Mbits/s at the start of 1999.

Relational data querying at DIII-D [5] has been accomplished over the past 10 years by using the S1032 relational database software that operates in the VAX/VMS environment. With this software, a series of time slice databases have been assembled that cover basic research topics such as confinement and divertor studies. Several advances in both diagnostic systems and in computer hardware have made this relational database obsolete. From a diagnostic standpoint, measurements at one point in time have given way to complete temporal evolution (e.g. Thomson scattering). Therefore, a single time point description of the plasma has less meaning. Computer hardware advancements include drastically reduced storage costs combined with greatly enhanced computational power. Therefore, greater amounts of analysis can be run and stored (e.g. magnetic equilibrium reconstructions). MDSplus is serving DIII-D's need for storage

of large amounts of information on the plasma's temporal behavior. What is presently lacking is a way to rapidly search the MDSplus repository. A relational database that summarizes MDSplus entries would allow the researcher to rapidly scan all DIII-D discharges and find that subset of special interest. This type of database, a broad and shallow data repository, will be one future evolution of the existing system. Additionally, a highly analyzed database will be created that will contain publication quality data to be loaded from MDSplus. DIII-D databases that are shared with the worldwide fusion community (e.g. previously assembled ITER confinement database) will be migrated from the old to the new system.

Presently, the DIII-D National Team along with MIT/C-Mod and PPPL/NSTX are investigating Microsoft's SQL Server 7 to replace existing commercial relational database software. The C-Mod Electronic Logbook has been adopted at DIII-D for use in the SQL Server environment and has been successfully used during DIII-D tokamak operations. Along with this database development will come GUI tools that are connected to the database engine, allowing relational queries from multi-platform clients and multi-software environments.

The DIII-D National Team is also looking beyond relational database technology to object-oriented database technology which has the potential to benefit both the theoretical and experimental researcher by solving problems where processes are so complex that relational databases cannot keep up. This new technology will work in concert with MDSplus, just as relational database technology has done in the past. Traditional relational database technology is being used successfully at C-Mod to compliment MDSplus storage for queries on scalar data from multiple shots; however, today's experiments and theory codes generate complex, multi-dimensional data that do not lend themselves to scalar quantification for querying. Object-relational databases should allow preservation of the complexity of the data and the relationships between them while not compromising query performance. For example, a fluctuation spectrum measured as a function of time over a wide frequency range could be represented as a spectrum object. A scientist could query the database to determine the shots for which the measured spectra are similar to one predicted from a simulation code, where the relationship between

objects that defines similar is stored either quantitatively or qualitatively in the database as part of the objects themselves. It is not yet apparent whether object-relational databases will perform better than the technologies already in use in the fusion community, both in terms of the speed at which large volumes of data can be queried, and in terms of the types of complex queries that are possible. Future work will aim to answer these questions.

At DIII-D, collaborators are offered full access to DIII-D raw and analyzed data as it is collected and analyzed. However, with these privileges comes the responsibility for collaborators to ensure that the data used are correct and are correctly interpreted and to ensure that appropriate credit for providing measurements and analysis are given. To help avoid misunderstandings over these responsibilities and to avoid potential loss of data access a DIII-D Data Usage and Publication Policy Agreement [6] has recently been formalized. Prior to gaining access to either the DIII-D computer systems or to MDSplus stored data, collaborators must read and sign the Data Usage Agreement. Briefly stated, this policy outlines a process for technical review that also assures publications are not unreasonably withheld or delayed. Publication includes posting the paper on either the DIII-D or collaborators web site.

3. DATA VIEWING AND ANALYSIS TOOLS

Work on data viewing and manipulation tools has focused on efficient and uniform GUI design with object-oriented programming for maximum code flexibility and access to both PTDATA and MDSplus data. The uniform GUI design, combined with a thorough documentation system (described below), decreases the non-productive time a new researcher must spend learning a new system. Also, existing users do not need to remember a new interface every time they switch analysis tools. GUI tools are being written in Interactive Data language (IDL) [7], a commercial product for scientific data manipulation and visualization.

A new object-oriented IDL based direct graphics library, GaPlotObj [4], has been created in collaboration with Fanning Software Consulting. This graphics library is a fundamental component of the new DIII-D viewing tools, providing a uniform GUI for graphical data manipulation. The GaPlotObj graphics library allows multiple 2D and/or 3D graphs with cursors for data readout, zooming, panning, slicing, and data selection for manipulation.

Two main viewing and analysis tools, EFITTools [8] and ReviewPlus [4], have been created that use the GaPlotObj graphics library. EFITTools combines the ability to perform an interactive EFIT, a kinetic EFIT, a time dependent EFIT, and the visualization of any EFIT calculation under one GUI umbrella. Visualizing EFIT results (EFITViewer) allows a researcher to examine a DIII-D magnetic equilibrium, the plasma profiles calculated during the equilibrium reconstruction, and the quality of the fit, including the magnetics and MSE data, in a more efficient method than was previously possible. The interactive EFIT capability dramatically reduces the time required to fine tune an MHD equilibrium reconstruction from hours to minutes by eliminating the unnecessary repeated initialization of the code. A kinetic EFIT typically took half a day, but with streamlining provided in the kinetic EFIT GUI, this time has been reduced to less than one hour.

The ReviewPlus tool is a general purpose data visualization program that provides interactive 2D and 3D graphs of data stored in either PTDATA or MDSplus. Researchers typically plot time histories, plasma profiles, or the temporal evolution of plasma profiles with this tool. As more analyzed data is stored in MDSplus it automatically becomes available for visualization in ReviewPlus. Besides adding more experimental data to the MDSplus archive, output results from theory/modeling analysis codes (e.g. ONETWO, GLF23, UEDGE) will also be added. Such storage will make it easier to compare experimental data with computer simulations.

The GAPprofiles analysis tool has been in existence for several years at DIII-D. This tool performs interactive bi-cubic spline fitting of experimentally measured plasma profiles including T_e , T_i , n_e , V_r , and P_{rad} . Graphical manipulations include the elimination of spurious data points as well as knot placement. This type of graphical data manipulation allows for rapid interactive fitting. Compared to the old batch mode type of analysis, this interactive work is much more efficient allowing the researcher to examine a larger number of discharges than was previously possible. These fitted profiles are used as inputs into energy transport and MHD stability codes.

Accompanying these new analysis tools is an HTML based searchable documentation system (<http://fusion.gat.com>) that brings critical knowledge to the on-site and remote collaborator via their web browser. The documentation focuses on computer codes and tokamak data including PTDATA and MDSplus. The data component of this documentation system is also available from the main analysis tools. A researcher simply needs to click on a data quantity of interest and the researcher's web browser will be started, or brought to the forefront, with the requested HTML page.

4. COMPUTER POWER FOR DATA ANALYSIS

Computer power for data analysis at DIII-D is provided by an HP 9000 Model T-600 3 processor server and a load balanced cluster of three HP and five Digital Unix workstations. Interactive computer load balancing (Fig. 3) is accomplished with LSF Suite 3.2 [9]. This software operates in a heterogeneous computational environment thereby combining all of the newer Unix based computers into one CPU cluster. The benefits of such a cluster are that all computers are easily and transparently available to all researchers, that CPU upgrades are as simple as removing one workstation and adding another, and a new on-site collaborator can easily add their own computer to the CPU cluster. Such an implementation has been possible because of the fast network connecting the workstations, the central file server that is available from all workstations, and the unified data access methodology. This cluster has more than tripled the Unix computer power available to the researcher.

Equilibrium reconstructions (EFIT) are performed in between tokamak pulses in a distributed computing environment consisting of workstations at General Atomics and LLNL. This system allows for two types of EFIT reconstructions to be run between the approximate 10 to 15 minute tokamak pulse cycle with a time resolution of 25 ms. Porting of EFIT to the Linux operating system running on PC hardware is presently underway. It is our intention to study the cost to performance ratio of EFIT running on a present day Linux PC versus our existing workstation environment. Assuming favorable results, a dedicated Linux PC cluster of up to 16 processors would be created for between shot equilibrium reconstructions. Preliminary estimates are that this type of cluster would produce EFIT results closer to 5 minutes after the conclusion of a tokamak pulse thereby giving the experimental team critical knowledge of the last pulse on a more rapid time scale.

The recent port of the MDSplus IP library to the MacOS allows all MDSplus and PTDATA stored data to be retrieved directly to the Macintosh for viewing and manipulation. Presently,

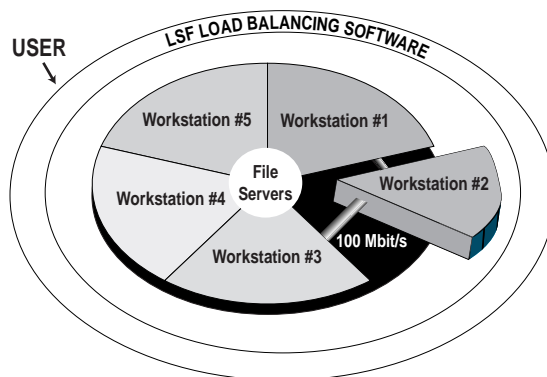


Fig. 3. LSF Suite 3.2 provides load balancing in a heterogeneous environment resulting in more efficient usage of CPU cycles and an easier upgrade path for computer power.

both EfitViewer and ReviewPlus run on the MacOS with IDL. Usage of our existing G3 Macintosh computers for more intense IDL based data analysis and visualization is being investigated; this has the potential to greatly expand existing CPU power. Additionally, this capability allows a researcher who is visiting another laboratory or attending a workshop to bring a PowerBook G3 computer, establish an IP connection, and perform data analysis as if they were in their office.

5. SUPPORT OF OFF-SITE DATA ANALYSIS

Remote or temporarily on-site collaborators typically receive an account on the HP T-600 server with the total number of user accounts now exceeding 300. To alleviate the ever increasing load placed on our CPU resources by off-site collaborators our analysis environment encourages usage of off-site computers. Such analysis is simplified by the availability of raw and analyzed DIII-D data via the MDSplus client/server interface. Additionally, the new IDL based viewing and manipulation tools are being distributed to remote collaborators either in the form of compiled binary executables or from a source code management system (CVS). Computer code management allows the interested researchers to modify existing tools and merge their changes back into the main repository. Creating tools in IDL has the added benefit of being able to move among different operating systems with minor modifications. These tools have presently been installed at C-Mod, NSTX, SSPX, and JET, as well as DIII-D.

Another aspect of remote data analysis is the ability to hold meetings to discuss on-going analysis. The DOE/OFES has facilitated off-site communication by providing funds to improve remote audio/video capability for both DIII-D and our collaborators. At the DIII-D facility, our current capability includes two conference rooms near the staff offices that have been equipped to share a ShowStation IP [10]. This device acts as a viewgraph machine for the researcher in the conference room and a Web server for those not in the conference room. The off-site collaborator can see the viewgraphs via a Web browser or, if their remote conference room is equipped with another ShowStation, the viewgraphs can be projected on their screen by their ShowStation. The off-site collaborator also has the flexibility to project from their Web browser viewgraphs on another ShowStation thereby allowing one researcher to remotely present viewgraphs to a large audience. Audio is presently handled via a telephone conference call which works fine for small meetings although these rooms will most likely be upgraded to full-duplex audio-conferencing. For larger meetings more sophisticated audio technology is presently being

investigated. The conference room next to the DIII-D control room will be equipped with a similar capability before the start of next year's tokamak operations which should allow routine enhanced broadcasting of the morning operations meeting. We anticipate the installation of ISDN based video-conferencing in these meeting rooms to enhance communication. Combinations of all of these technologies will be experimentally investigated in the DIII-D control room to understand what combination is best suited for use during tokamak operations.

6. SUMMARY

A long term plan has been formulated and is being implemented to increase the DIII-D data analysis throughput and data retrieval rate. The basic component of this plan is to create a uniform interface to a very heterogeneous environment. Present software enhancements have focused on streamlined analysis, automation, and GUI systems to accommodate a larger user community. Future software work will include establishing computer science collaborations among the different fusion facilities worldwide to avoid unnecessary duplication of effort and to benefit from a larger base of programmers. Such projects include database tools, a code run management database, signal name translation capability, and a general visualization tool for examining results of power balance analysis. Additionally, as the breadth of the DIII-D research team continues to expand, future work will examine the best technology for remote participation in both meetings and tokamak operations.

REFERENCES

- [1] B.B. McHarg Jr., “The DIII-D Computing Environment: Characteristics and Recent Changes,” these proceedings.
- [2] B.B. McHarg, Proc. 15th IEEE/NPSS Symposium on Fusion Engineering, Hyannis, Massachusetts, Vol. 1 (Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey, 1994) p. 123.
- [3] J. Stillerman, *et al.*, Rev. Sci. Instrum. **68**, 939 (1997).
- [4] J. Schachter, “Data Analysis Software Tools for Enhanced Collaboration at the DIII-D National Fusion Facility,” these proceedings.
- [5] D.P. Schissel, *et al.*, Rev. Sci. Instrum. **57**, 1932 (1986).
- [6] <http://Fusion.gat.com/Comp/Analysis>.
- [7] Interactive Data Language, Research Systems, Inc. (<http://www.rsinc.com>)
- [8] Q. Peng, *et al.*, “IEFIT – An Interactive Approach to High Temperature Fusion Plasma Magnetic Equilibrium Fitting,” to be published in Transactions on Nuclear Science.
- [9] Platform Computing Corporation (<http://www.platform.com>)
- [10] Polycom, Inc. (<http://www.polycom.com>).

ACKNOWLEDGMENTS

This is a report of work supported by the U.S. Department of Energy under Contracts No. DE-AC03-99ER54463, w-7405-ENG48 and DE-AC02-76CH03073. The authors acknowledge stimulating discussions with T. Fredian, M. Greenwald, and J. Stillerman from Alcator C-Mod.