

GA-A23536

**STATUS REPORT:
DATA ANALYSIS INFRASTRUCTURE
AT THE DIII-D NATIONAL FUSION FACILITY**

**by
D.P. SCHISSEL and B.B. McHARG, Jr.
for the DIII-D NATIONAL TEAM**

**Prepared under
Contract No. DE-AC03-99ER54463
for the U.S. Department of Energy**

OCTOBER 2000

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.

GA-A23536

**STATUS REPORT:
DATA ANALYSIS INFRASTRUCTURE
AT THE DIII-D NATIONAL FUSION FACILITY**

by
D.P. SCHISSEL and B.B. McHARG, Jr.
for the DIII-D NATIONAL TEAM

Prepared under
Contract No. DE-AC03-99ER54463
for the U.S. Department of Energy

**GENERAL ATOMICS PROJECT 30033
OCTOBER 2000**

EXECUTIVE SUMMARY

This paper summarizes the architecture, including hardware and software, which is used at the DIII-D National Fusion Facility for analyzed data acquisition, analysis, and archival/retrieval. A number of the new features of the present data analysis environment have been developed since the formation of the Data Analysis Applications Group in September 1997. This group was formed by the Director of the DIII-D National Fusion Facility to increase the data analysis throughput and data retrieval rate for the DIII-D National Team. This status report provides a concise summary of the work performed, in concert with the Computer Systems Group, by the Data Analysis Applications Group during its first three years.

The DIII-D system supports both near real-time on-line monitoring and direction of the operation of the DIII-D magnetic fusion (tokamak) experiment and longer-term detailed analysis of the resulting plasma physics data. This comprehensive data analysis capability facilitates efficient conduct of plasma science experiments using DIII-D, under the direction of an operation and science staff comprising more than 200 individuals drawn from both General Atomics and worldwide fusion laboratories, universities, and industries. 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. Since a large fraction of the DIII-D team is off-site, DIII-D data and applications are readily available at their various home institutions via high-speed network connections.

The DIII-D analyzed data acquisition, analysis and archiving system is fully modular and network-based, and is configured to provide a uniform data processing and display environment supporting both interactive and automatic analysis by local and remote users. Figure E-1 shows the overall configuration of this data analysis environment that handles storage for both raw (as digitized) and analyzed data from each DIII-D plasma operation pulse (commonly termed a "pulse" or "shot"). Data from each pulse comprises a time-sampled array of information from the various plasma diagnostic systems that view and record the plasma, plus corresponding records of the various DIII-D device (tokamak) operation settings that are used to produce the pulse. The overall duration of each DIII-D operation pulse is approximately 10 s, with the duration of the plasma portion typically being from 3 to 7 s. Each pulse typically produces 150 MB of compressed raw digitized data. Once the data is acquired, specialized analysis routines are automatically initiated on a variety of computers. As soon as any raw or analyzed data is acquired, it can be graphically displayed by the DIII-D operations and experiment staff. Typically, about 4 minutes after the pulse ends, a complete temporal history of the pulse is available including magnetic equilibrium and profiles of density, temperature, current, and

plasma rotation. The DIII-D team, both on and off-site, uses this near real-time “control-room” data to monitor the results of the pulse and modify the parameters for the next. Figure E-2 shows a typical example of a control-room presentation.




Fig. E-1. Data flow at the DIII-D National Fusion Facility.




Fig. E-2. Interactive visualization is available near real-time during DIII-D operations .

Data is archived in three formats, with increasing degrees of analysis and reduction to higher level plasma parameters and attributes: raw (PTDATA, presently 1.4 TB), analyzed (MDSplus,

presently 70 GB) and analysis highlights (DIII–D Relational Database or “D3DRDB”, presently 0.2 GB). The relationship among these three archives is illustrated schematically in Figure E–3. The D3DRDB contains the most refined and abstracted distillation of a pulse and can be used by the DIII–D team to investigate similarity and correlations among subsets of the data. The MDSplus archive contains all analyzed data and allows the staff to conduct detailed studies of candidate pulses identified by inspection or D3DRDB searches. The PTDATA archive contains all raw or as digitized DIII–D data.




Fig. E–3. The DIII–D relational database, D3DRDB, works in concert with PTDATA and MDSplus by storing highlights of the data.

The raw data acquisition system (PTDATA) presently supports the acquisition of approximately 300 GB of compressed raw data per year (~2000 pulses) and also provides for cumulative storage and timely access to the 1.4 TB of data that has been acquired since DIII–D operation began in 1986. A 3 TB hierarchical data storage system comprising 200 GB of hard disk storage, 600 GB of magneto-optical storage and 2.2 TB of tape storage provides unattended access to all raw data with a delay that does not exceed 5 minutes. Data from the most recent or frequently accessed high interest pulses are available instantaneously from the hard disk. Data from less frequently requested pulses are available from the magneto-optical and tape archives on correspondingly slower but still interactive time-scales. Access to the data is completely transparent to the requester with only the time to retrieve the data varying. As the total amount of analyzed data is much less than raw data, both the analyzed data repository (MDSplus) and the DIII–D relational database (D3DRDB) are contained only on hard-disk.

As in most present and scientific and computer related endeavors, there is an overriding need to accommodate rapid and continued data growth. Figure E-4 shows the growth of the cumulative compressed raw data (PTDATA) since the start of DIII-D operations. This growth is presently scaling as $t^{2.7}$, where t is the time since the inception of DIII-D operation. If the present growth of raw data continues, this data archive will reach approximately 4 TB by 2004. It is anticipated that the increasing use of multi-dimensional (2D and 3D) experimental and simulation data will expand the analyzed data repository (MDSplus) to the 5 TB range by 2004.




Fig. E-4. There has been a quadratic increase in total raw data (PTDATA) acquired over the lifetime of DIII-D operations.

The ability to update the system processing power to handle the increasingly sophisticated processing of 2D and 3D analysis, and to handle high time resolution analysis of longer duration pulses, is critical to the continued success of the DIII-D program. Support of general data analysis by both local and off-site staff is handled by a powerful and readily expandable load balanced Unix workstation cluster. Additionally, specialized workstations and a 12-processor Linux parallel processing cluster provide local users with near real-time automatic processing of key data used in the pulse-to-pulse conduct of experiments. Such recent increases in processing power have yielded near-real-time analyzed data feedback to the experimental team. Certain key pieces of data can now be compared between pulses without having to make allowances for approximations or interpolations. The combination of high time resolution and high spatial resolution with enhanced data visualization has resulted in more efficient utilization of DIII-D experimental time.

From the user's point of view, all of the DIII-D data, regardless of where it resides, can be accessed and/or analyzed using a common and easily learned set of standard software tools and commands. This commonality of software and procedures, coupled with universal network access, facilitates rapid and efficient utilization of the DIII-D data by a national and international team of operations and research staff. Data visualization tools include specialized applications and one general purpose program that provides interactive 2D and 3D graphs. As the amount and complexity of data continues to increase, applications are being developed that use advanced graphics techniques that take advantage of specialized hardware for rapid visualization of 3D images.

The DIII-D analyzed data acquisition, analysis and archiving system had its origins in the now rudimentary data acquisition system that was initially set up in 1975 to support the operation of the Doublet III tokamak experiment, DIII-D's direct predecessor. Since then, hardware and software configurations of the DIII-D data system have evolved over the course of time into what is now a fully network-based and modular system predicated on standard computer industry data storage and processing protocols. There are many "lessons learned" behind the present configuration. Briefly put, the most important of these lessons learned are:

- The need for continuous expandability in data capacity and analysis computing power
- The need to be able to incorporate new hardware and software innovations without the need to modify the user interface and operation aspects of the system
- The importance of configuring the system in a modular fashion to be able efficiently use industry standard hardware and software
- The need to make both current and past/historic data, both raw and analyzed, available to present users in a timely and transparent fashion
- The need to provide a common and effective and user-friendly set of data retrieval, analysis and viewing/presentation tools
- The need to provide immediate and high-quality presentation of raw and initially analyzed data to the control room and remote-site operating staffs

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 in September 1997 formed the Data Analysis Applications Group whose mission is to increase the DIII-D data analysis throughput and data retrieval rate. This status report presents the results of work performed, in concert with the Computer Systems Group, over the past three years [1,2,3] and includes the long-term plan, summarizes progress to date, and outlines near term work. The underlying philosophy 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.

Figure 1 illustrates the flow of data at the DIII-D National Fusion Facility. In the lower left-hand side of Fig. 1 the raw data from DIII-D is acquired and stored in PTDATA (Section 2). As soon as the data is available the automatic between pulse analysis begins on dedicated Unix systems and on the multi-processor Linux cluster (Section 7). This analyzed data is stored either in MDSplus (Section 3) or in the DIII-D relational database D3DRDB (Section 4). The onsite user enters the diagram from the top right and has access to all data from the LSF Unix cluster (Section 7). The remote user (Section 8) has access either to the DIII-D data directly via the Internet or by first login into the DIII-D LSF Unix cluster.

As stated earlier, the underlying philosophy behind these development efforts is a uniform interface to the user community that results in more efficient data analysis. Figure 2 illustrates the software that is placed on the hardware of Fig. 1. Starting from the top of Fig. 2 and reading left to right, the user has transparent access to many Unix workstations via the LSF Unix cluster (Section 7). Thus, the user never needs to know what workstation they are using since there is a unified file system with equal access to all codes. The interactive visualization tools (Section 6) that run on the LSF Unix cluster have a uniform look and feel so that a user does not need to remember a new interface every time they switch tools. Finally, there is a unified data storage mechanism (Sections 2, 3) for all DIII-D data.




Fig. 1. Data flow at the DIII-D National Fusion Facility.




Fig. 2. Users benefit from the unification of a heterogeneous environment that includes access to computers, GUI tools, and data.

2. PTDATA: RAW TOKAMAK DATA STORAGE AND RETRIEVAL

The DIII-D tokamak operates in a pulsed mode producing plasmas of 5–10 s duration every 10–15 min, with 25–35 pulses per operating day. The number of operating days is usually set by the available operating funds and is typically around 50 to 75 days per year. The PTDATA format and associated code is the mechanism for archival and retrieval of DIII-D raw digitizer data since DIII-D began operation in 1986 [4]. Raw data is stored on disk in a compressed format and is decompressed by PTDATA when accessed. Thus this allows for more efficient data storage as well as a decrease in network traffic. PTDATA uses TCP/IP networking for enhanced connectivity thereby yielding data access from anywhere on the Internet.

The amount of raw digitizer data acquired per tokamak pulse has been steadily increasing with the largest pulse to date being 250 MB compressed. Figure 3 illustrates the cumulative total amount of PTDATA raw compressed data acquired versus fiscal year (October–September) since 1986. Figure 4 illustrates the quantity of raw compressed data acquired per fiscal year. These increases in raw data are driven primarily by: (1) digitization of data at faster rates and for longer times, (2) longer plasma pulses, and 3) more diagnostics on the tokamak. Presently a total of over 1.4 TB of compressed raw data (PTDATA) has been acquired over the lifetime of DIII-D.




Fig. 3. There has been a quadratic increase in total raw data (PTDATA) acquired over the lifetime of DIII-D operations.




Fig. 4. Raw data (PTDATA) acquired per fiscal year continues to increase.

A mass storage system with a total capacity of approximately 3 TB has been installed for storage of DIII-D raw digitizer data [1]. Along with 200 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. This hardware, along with the associated Hierarchical Storage Management (HSM) software, provides unattended raw data availability 24 hours a day and 7 days a week (24x7) for every pulse ever taken. This 24x7 raw data access 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 800 GB of raw data that is the equivalent to approximately 8000 compressed tokamak pulses. The remaining 2.2 TB from the DLT system is available on a 3 to 5 minute time scale. After the tokamak pulse and compression of the data, the PTDATA data is moved to the pulse server system with the 200 GB disk array where it becomes part of the HSM system. If not accessed, data will eventually migrate to optical media, and later to tape media. Another benefit of the HSM system is the automation of data archiving. Data is copied to a different HSM volume from which it is automatically copied to DLT tape for permanent archiving. Two DLT tape copies are made, one to remain on site, and one for offsite storage in the event of the loss of onsite data.

3. MDSplus: ANALYZED TOKAMAK DATA STORAGE AND RETRIEVAL

As discussed in the last section, DIII-D uses the PTDATA system for acquisition, storage, archival and retrieval of raw data. Prior to the adoption of MDSplus [5] for DIII-D in September of 1997, no equivalent system existed for results of analyses performed on the raw data. Each analysis code or diagnostic therefore developed its own system, writing their results in differently formatted files, sometimes hundreds per plasma pulse, scattered across the disks of multiple computers. These data were often stored without any contextual information or comments. Therefore, in addition to the confusion of having to use a different code to access each analyzed dataset, it was not possible for scientists to explore their colleagues' results without first meeting to determine the context of the data. The unified format of the MDSplus system allows researchers to learn a few computer commands and read a vast amount of DIII-D data [6]. Figure 5 graphically illustrates the two different storage philosophies.


MDSplus also simplifies the tasks of the computer support personnel at DIII-D. Before MDSplus, each data access routine for each analyzed dataset had to be maintained on each of the operating systems in the heterogeneous DIII-D computing environment. Now, only the MDSplus libraries need to be maintained. Data in MDSplus is also compressed, saving disk space, and only a handful of files are necessary to store the results from an entire plasma pulse. The flexibility of MDSplus allows addition of datasets to the existing archive and storage of multiple versions of analysis.

A key benefit to MDSplus is its ability to interface with other data systems. For example, the interface to the PTDATA system allows users to transparently retrieve the raw data stored within through MDSplus. This is accomplished by writing MDSplus functions containing calls to the PTDATA shared library routines. In situations where analyzed data are simply calibrations applied to raw measurements, MDSplus will store the calibrations and an expression multiplying them by the raw data stored in PTDATA. This reduces the amount of disk space used, and allows calibrations to be changed independently of the raw data, while automatically providing the most up-to-date results to the users.

The DIII-D analyzed diagnostic data stored in the MDSplus system continues to increase with presently 8900 archived pulses representing 70 GB of data. Unix MDSplus data service was chosen at DIII-D with the benefits of faster data service and easier integration into the DIII-D Unix analysis environment. MDSplus data is served from a Compaq DEC AlphaServer 800 5/500 with 1 GB of RAM running Compaq Tru64 Unix, with an nSTOR GigaRAID AA system

Conventional Storage: Hard to Share Results

- Separate interface for each data type
- Must know data format and file location



MDSplus: Remote Exploration of Data Productive

- One interface to many data types
- Store all relevant information




Fig. 5 MDSplus simplifies the interface to analyzed data making it easier for remote collaborators to explore.

containing 284 GB. This system provides a flexible upgrade path for future expansion in both CPU power and magnetic storage space. Given the projected growth rate of MDSplus storage needs and the falling cost of hard disk storage, the long term plan is to store all MDSplus data on magnetic disk, rather than involving slower, more complex optical or tape mass storage systems.

Currently, up to 21 separate analyzed datasets are present for each pulse, including EFIT magnetic equilibrium results, kinetic T_e , T_i , n_e , V_r , and n_i profiles from the Thomson scattering and CER diagnostics, and spectroscopy and divertor imaging results. Presently in MDSplus,

there are usually 21 MB of analyzed data stored per pulse. As more diagnostic data are added, this number will continue to increase to an anticipated 100 MB per pulse. More analyzed datasets are being added both to new pulses and retrospectively to old ones. During experiment operation, most of the data are loaded automatically between plasma pulses, giving the researcher critical information before setting up the next pulse. Between pulse analysis jobs and data storage are managed by the MDSplus dispatching system based on events declared when the requisite raw data are available after acquisition. MDSplus loading can occur at any time after a plasma pulse so that the results of more detailed analyses can be stored in the days following an experiment.

Longer term, datasets from microturbulence and macrostability simulation codes will be added to the DIII-D MDSplus repository. This unified storage of simulation results will greatly facilitate the comparison of experimental and theoretical data.

4. RAPID DATA MINING WITH THE DIII-D RELATIONAL DATABASE

A relational database is a collection of data items organized as a set of formally described tables from which data can be accessed or reassembled. The tables are related to each other by at least one common column and indexed by the relational database software for rapid data retrieval and updating. The new DIII-D relational database, D3DRDB, works in concert with PTDATA and MDSplus by storing highlights of these data repositories (Fig. 6). Users can rapidly search through the data highlights stored in the relational database and find the subset of pulses that have special interest. D3DRDB allows survey type questions such as “What pulses have a plasma current greater than 1.5 MA?” to be answered much more rapidly in comparison to scanning all 50,000 pulses stored in the PTDATA HSM system. The subset of pulses returned from a survey can then be analyzed in more detail by using the complete dataset stored in PTDATA and MDSplus.




Fig. 6. The DIII-D relational database, D3DRDB, works in concert with PTDATA and MDSplus by storing highlights of the data.

In early 2000 the switch was made at DIII-D from the S1032 relational database software that operates in the VAX/VMS environment [7] to the Microsoft SQL Server 7 software running in the Windows NT environment. This change has the benefit of providing faster data service,

easier integration into the DIII-D Unix analysis environment, easier creation of GUI based tools that allow for relational queries from multi-platform clients and multi-software environments, and reduced hardware and software costs associated with maintaining and serving the data. A Pentium III 550 MHz computer with 512 MB of RAM and a 9 GB RAID array is the hardware for serving data from the DIII-D relational database (D3DRDB).

In the old DIII-D relational database, 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. EFIT magnetic equilibrium reconstructions). MDSplus is serving DIII-D's need for storage of large amounts of information on the plasma's temporal behavior and is therefore replacing a large amount of the functionality of the old relational database system.

DIII-D experimental summary information was the first data added to D3DRDB. The summary information in the database can be queried directly using the relational database industry standard Structured Query Language (SQL). Such SQL capability is available from any web browser [8], from a variety of programming languages in the Unix environment, and from a variety of both PC and Macintosh based application programs like Excel. Experimental summary information contains the daily text summaries from the Session Leader, Physics Operator, Chief Operator, as well as the Physics Miniproposal number, experimental title, any electronic logbook comments, key experimental personnel, and approximately 50 scalar quantities that characterize every DIII-D pulse. Entries into the electronic logbook can be made by any staff member and are divided into broad headings covering the experiment, diagnostics, and engineering. The logbook, adopted from C-Mod, has many advantages over a traditional paper based system. The two primary advantages are the ability to see comments from others in real time and the ability to rapidly query past entries. The electronic logbook allows the DIII-D National Team to monitor, from anywhere, the pulse-to-pulse progress of the experiment. Where possible, old historical data has been entered into the new database including the last 10 years of text based experimental summaries. For present and future run-days, a series of automatic procedures and new GUI tools have been developed to load the summary data into D3DRDB.

A very powerful feature of the DIII-D relational database system is the ability for users to create their own tables of data that can be relationally joined to the existing D3DRDB structure. This system simultaneously benefits the individual contributor and the entire DIII-D National Team. The contributor benefits since their data can be easily combined with the entire D3DRDB data repository. The National Team benefits since the contributor's data is available to all team

members. Data stored in user tables include information on plasma impurities and pellet injection. Another form of user table includes the DIII-D data that is shared with the international fusion community (e.g. previously assembled ITER databases). These tables originated in the old relational database system and are presently being migrated to the new infrastructure.

Documentation of data stored in PTDATA, MDSplus, and D3DRDB is stored in the new relational database environment. For each data value, information is stored on the responsible officer, the type of data, a brief description of the data, a longer description of the data, where the data is stored, data units, which diagnostic created the data, the last calibration change, the diagnostic revision history, and any sources of further documentation. A simple web interface allows queries based on data name, by a string in its short description, by diagnostic type, by responsible officer, or by any user-created SQL statement [9].

The relational database system is also being used to monitor data analysis at DIII-D to support future decisions for improving the data analysis infrastructure. Decisions such as adding more computer power, memory, or data storage are based on usage statistics. Storing this information in the relational database is replacing writing the information to a text file. This has the benefit of rapid data querying and easy data display. A code usage system tracks when one of the general analysis tools is started, from what platform it is started from, and who is using the tool [10]. Information on computer usage, memory usage, disk space usage, and PTDATA or MDSplus data retrieval speeds are all presently being transferred to the relational database environment. Data on the usage of MDSplus will be added to the relational database soon.

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. Object relational database management systems (ORDBMS) add new object storage capabilities to the relational system. They integrate management of traditional data with complex objects such as time-series and geospatial data and diverse binary media such as audio, video, images, and applets. By encapsulating methods with data structures, an ORDBMS server can execute complex analytical and data manipulation operations to search and transform multimedia and other complex objects. 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 to compliment MDSplus storage for queries on scalar data from multiple pulses; 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 pulses 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 during calendar year 2001 will aim to answer these questions.

5. DATA ACCESS POLICY

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 [11] has 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.

6. DATA VISUALIZATION AND ANALYSIS TOOLS

Work on data visualization and manipulation tools has focused on efficient and uniform GUI design with object-oriented programming for maximum code flexibility and access to PTDATA, MDSplus, and D3DRDB data. The uniform GUI design, combined with a thorough documentation system, 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), a commercial product for scientific data manipulation and visualization from Research Systems Incorporated [12]. IDL combines a rich, easy-to-use language with extensive and powerful interactive graphics, and can run on a variety of CPU architectures with minimal porting. One of the strengths of IDL is its ability to couple with the window management system to provide dynamic plotting; that is, the ability for the user to change how the data is displayed, and even the data themselves, with a mouse-driven, “point and click” interface.

Dynamic plotting must be included in an application by the developer since it is not automatically available in IDL. Rather than coding the same interface over again in each application, a new object-oriented IDL based direct graphics library, GAPlotObj [6], has been created that can be used in all applications for dynamic plotting. This graphics library is a fundamental component of the new DIII-D viewing tools, providing a uniform GUI for graphical data manipulation. From the user’s point of view, this means that in every application the same action, either a keystroke or a mouse click, gives the same result. Users therefore need to learn the interface only once. The GAPlotObj graphics library allows multiple 2D and/or 3D graphs with cursors for data readout, zooming, panning, slicing, and data selection for manipulation (Fig. 7). 3D data items can be shown as surfaces, images, or contour plots. Slicing means that 3D data can be plotted as a function of one of their two independent dimensions. Furthermore, 2D and 3D plots can be linked so that actions performed on one plot affect others (Fig. 8).

Two main viewing and analysis tools, EFITTools [13] and ReviewPlus [6], 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. EFIT is the MHD equilibrium fitting code that has been in use at DIII-D for over 15 years. Visualizing EFIT results (EFITViewer) allows a researcher to examine a poloidal cross section of the DIII-D magnetic equilibrium (Fig. 9). This display is dynamic as users can zoom in on features of interest, turn geometry overlays on and off, and overlay cross sections of multiple pulses for quick comparison. Other plots available in EFITViewer include the plasma profiles calculated during the equilibrium reconstruction, and




Fig. 7. The GAPlotObj IDL-based graphics library allows dynamic plotting including select, zoom, edit, cursor, and slice.

how well the calculation reproduces the magnetics and Motional Stark Effect (MSE) data. At the press of a button users can map the profiles to any of three possible coordinate systems: major radius, toroidal flux, or poloidal flux. The interactive EFIT capability dramatically reduces the data. This results in a more accurate equilibrium reconstruction of the plasma current and pressure profiles, which is critical to transport and MHD stability analysis. Previously, 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.




Fig. 8. The GAPlotObj IDL-based graphics library allows interaction between plots which is illustrated here between a 3D contour plot and 2D slice plots of a temperature profile.

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 (Fig. 10). It is the main data viewing tool for both experimental operations and for more in-depth analysis afterwards. Researchers typically plot time histories, plasma profiles, or the temporal evolution of plasma profiles with this tool. ReviewPlus is also able to plot any mathematical combination of data signals from MDSplus and PTDATA, including allowing built-in or user-written IDL functions




Fig. 9. The EFITViewer tool provides rapid assessment of an EFIT equilibrium fit.

to operate on the data. Menu options are provided for the most common operations including smoothing, integration, differentiation, and fast Fourier transforms. As more analyzed data are stored in MDSplus they automatically become available for visualization in ReviewPlus. Besides adding more experimental data to the MDSplus archive, results from theory/modeling analysis codes (e.g. ONETWO or TRANSP for transport analysis/simulation) will also be added. Such storage will make it easier to compare experimental data with computer simulations.

ReviewPlus




Fig. 10. The ReviewPlus tool provides general 2D and 3D data visualization.

The GAPfiles 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 processing method 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.

Other analysis tools include the Entry_Display program, adopted from C-Mod, which serves as an electronic logbook. The program interface is written in IDL and the text comments are stored in the DIII-D relational database system as part of the experimental summary information. Simple GUI based tools that do not require plotting are written in Java so that only needed IDL licenses are used. The Entry_Display program is being ported to Java.

Available from these new analysis tools is the data documentation system, described in Section 4, that brings critical knowledge to the user via a web browser. From within a tool, a user can either request information on a specific piece of data or search the documentation repository.

Either way, a web browser is automatically started and displayed on the users' window management system. This system has numerous advantages. First, the web browser is remotely started on a Linux PC thereby saving expensive HP or Compaq workstation memory and CPU resources. Second, the user interacts with the familiar browser GUI that they have seen when starting from their own web browser.

As the amount and complexity of data continues to increase some time has been spent investigating methods for displaying this data by means that go beyond our present day usage of IDL. Specifically, we have investigated techniques for utilizing hardware graphics rendering with OpenGL. Rendering graphics in hardware rather than software has the advantage of much greater speed and therefore the ability to handle much larger amounts of data. IDL implements hardware rendering and OpenGL via their Object Graphics Library. Prototype development in this environment has lead us to believe that IDL in its present form has a limited capability for volume rendering, shading, and opacity compared to other software packages. In contrast to IDL, the freeware package Visualization Toolkit (VTK) has a very powerful visualization capability. Prototype applications developed in VTK on DIII-D data show good promise in the ability to visualize higher dimensionality datasets. Figure 11 shows sample output from an OpenGL based 3D and 2D interactive visualization tool which is being used to examine magnetic islands. Figure 11(a) shows in 3D an island in yellow surrounded by neighboring field lines. The magnetic field line of this island is shown exclusively in 3D in Fig. 11(b). Poloidal cuts through these 3D images are shown in Figs. 11(c) and 11(d) where the individual field line is colored yellow alongside its neighbors and also exclusively. The close-up view in Fig. 11(a) and the multiple images in Fig. 11(c) illustrate the zoom feature of the tool.

Combined with the software work is the corresponding effort for deploying a low cost visualization workstation on the scientists' desktop. A Linux PC workstation with a graphics card that supports OpenGL hardware graphics rendering is being investigated for the desktop. This graphics station would cost approximately \$1k each and can therefore be deployed in large numbers to the onsite DIII-D National Team. This solution has the added benefit of being a computational workstation for a large number of our existing scientific analysis applications. Finally, to make systems management of these workstations tolerable, remote network booting of the system is being investigated. This type of situation allows a system administrator to maintain only one system while the network automatically manages all the other systems.




Fig. 11. Interactive 2D and 3D visualization of magnetic field line data using hardware graphics rendering with OpenGL.

7. COMPUTER POWER FOR DATA ANALYSIS

Similar to the chaotic situation described in Section 3 before the adoption of MDSplus for organizing analyzed data, there used to be no coordinated plan to share compute resources among data analysis workstations. Individual groups of users used to buy their own workstations to perform data analysis with very limited sharing of computer resources among groups. This situation resulted in a large under-utilization of computer resources and increased costs associated with commercial software licenses. For example, if one of the groups departed for a week-long meeting, their workstation would remain idle while other groups that remained behind continued to work. From the software standpoint, each workstation was purchased fully equipped even though the software was utilized for a limited amount of time.

The commercial software LSF Suite (Load Sharing Facility) from Platform Computing [14] has been adopted to perform distributed load sharing among the many Unix workstations at DIII-D. This software operates in a heterogeneous computational environment thereby combining all of the newer Unix based computers into one CPU cluster. Presently there are six HP and five Digital Unix workstations in the cluster (Fig. 12). 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. Additionally, LSF Suite manages and understands the location of commercial software licenses so that a user on any workstation can transparently access any commercial software package on a different workstation. LSF Suite's management of commercial software has resulted in a significant cost reduction by reducing the amount of licenses we have needed to purchase.

Computer power for data analysis at DIII-D still includes one large central system, an HP 9000 Model T-600 3 processor server. This system can easily handle 100 simultaneous logins and acts as a client in the LSF cluster. A client LSF license allows users to be load balanced away from the client and never to the client. Therefore users are encouraged to login to the large HP system to start their work. In contrast, a typical HP or Compaq workstation could never handle so many simultaneous logins. However, since most of the intense work is load balanced to the workstations, the large server does not need to be upgraded as often as would typically be required and therefore results in a significant cost savings.

Such a load balanced cluster 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. The fast network includes switched Fast Ethernet (100 Mbits/s)




Fig. 12. LSF Suite 3.2 provides load balancing in a heterogeneous environment of an arbitrary number of workstations resulting in more efficient usage of CPU cycles and an easier upgrade path for computer power.

between workstations, to the onsite offices of the National Team, and within the DIII-D control room. Central file serving in the LSF cluster is accomplished with a 200 GB Network Appliance F520 network data storage system. This cluster has increased the Unix computer power available to the researcher by over a factor of ten.

EFIT magnetic equilibrium reconstructions are performed in between tokamak pulses on a new 12 processor (Pentium III 550 MHz with 256 MB RAM) PC Linux cluster [15]. This work required modifying EFIT to use the parallel message passing interface (MPI) library so that multiple independent equilibria can be rapidly generated from experimental data. This new system calculates equilibria eight times faster than the previous system yielding a complete equilibrium time-history on a 25 ms time-scale two minutes after a pulse ends (Fig. 13). Additionally, the PC-based hardware for this increased analysis capability costs significantly less than a comparable HP or Compaq Unix workstation solution. A new interface has also been created so that the physics operators and session leaders can directly control the EFIT setup including the time resolution and the input parameter file being used.

These between-pulse EFITs have been integrated into the Fault Identification and Communication System (FICS) for rapid error detection during experimental operations [16]. FICS executes automatically after every plasma pulse to check dozens of subsystems for proper operation and then communicates the test results to the experimental team. This system has been used routinely during DIII-D operations and has led to an increase in tokamak productivity. It is planned in the near term to add over a dozen between pulse diagnostic analysis codes to FICS.




Fig. 13. Total computation time versus the time resolution of EFIT equilibria calculation for different number of processors. All analysis covers the 6 s length of pulse 103770 which for the 25 ms resolution case requires 240 equilibria. The data point represents the actual result for between pulse calculation during experimental operation. The EFIT equilibria are calculated for a 65x65 grid size at double precision.

Long term it is anticipated that a complete time dependent power balance analysis will be performed between pulses. FICS will then allow a complete consistency check of many diagnostic systems during operations allowing rapid identification and repair of any problems. This is to be contrasted with the present situation where higher level diagnostic consistency checks, like a power balance analysis, are performed after the experimental day has concluded when it is too late to fix the error and take more data.

8. SUPPORT FOR 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 [17], and JET, as well as DIII-D. Utilization of off-site hosts benefits both the remote collaborator and the DIII-D computing infrastructure when the bandwidth required for interacting with the tools exceeds that required for transmitting the data. Collaborators would experience faster response from their operations on the data, as all interaction with the window management system is local, and the DIII-D network and computers would be less burdened. In situations where remote collaborators cannot obtain IDL licenses, or where the bandwidth required for data retrieval is greater than interactive display, users are able to log into the DIII-D computing network to use the local tools on DIII-D computer's, setting their display back to their own, off-site host.

ReviewPlus is one of the tools currently in use in the U.S. fusion community that is already capable of viewing data from multiple sites simultaneously. For each signal of interest, a user can specify the IP address of the MDSplus server from which to retrieve data. Since nothing in the code is specific to DIII-D, it can be easily setup and used at any remote site with MDSplus and an IDL license. Users familiar with ReviewPlus therefore do not need to learn how to use other tools to view data from other sites.

Another aspect of remote data analysis is the ability to hold meetings to discuss on-going analysis. The DIII-D National Team has worked to improve off-site communication by upgrading our remote audio/video hardware. At the DIII-D facility, our current capability includes two conference rooms (Fig. 14) near the staff offices that have been equipped to share a ShowStation IP from Polycom [18]. 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




Fig. 14. The DIII-D National Fusion Facility has several electronic conference rooms equipped for multi-site remote meetings.

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.

Video conferencing is presently handled with a ViewStation MP from Polycom [18] in both conference rooms. Broadcast is accomplished over ISDN at 384 kbits/s. Mutli-point video conferences are handled by an ESNET bridge allowing an unlimited number of remote sites to

conduct a meeting. Therefore, a complete working meeting can involve multiple sites sharing viewgraphs with the ShowStation, and video and audio with the ViewStation.

The DIII-D tokamak is physically located approximately 1.5 miles from the National Team offices. A conference room, located next to the DIII-D control room, is similarly equipped as the conference rooms described above. However, this room and the DIII-D control room are equipped with the ability to broadcast audio and video over the web [19]. Therefore the morning operations meeting, which is held prior to every experimental day, is available for monitoring via any web browser. At the conclusion of the meeting the web broadcast is switched to monitor the control room so that the days experimental operation can be monitored.

REFERENCES

- [1] B.B. McHarg Jr., "The DIII-D Computing Environment: Characteristics and Recent Changes," Proc. Of the 2nd IAEA Tech. Com. Mtg. On Control, Data Acquisition and Remote Participation on Fusion Research, to be published in Fusion Eng. And Design (1999).
- [2] D.P. Schissel, *et al.*, Proc. of the 26th Euro. Conf. On Contr. Fusion and Plasma Phys., Maastricht, Netherlands, Vol. 23J (European Physical Society, Nieuwegein, 1999) p. 1225.
- [3] D.P. Schissel, *et al.*, "Enhanced Computational Infrastructure for Data Analysis at the DIII-D National Fusion Facility," Proc. of the 2nd IAEA Tech. Com. Mtg on Control, Data Acquisition and Remote Participation on Fusion Research, to be published in Fusion Eng. and Design (1999).
- [4] B.B. McHarg, Proc. of the 15th IEEE/NPSS Symposium on Fusion Engineering, Hyannis, Massachusetts, Vol. 1 (Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey, 1994) p. 123.
- [5] J. Stillerman, *et al.*, Rev. Sci. Instrum. **68**, 939 (1997).
- [6] J. Schachter, "Data Analysis Software Tools for Enhanced Collaboration at the DIII-D National Fusion Facility," Proc. of the 2nd IAEA Tech. Com. Mtg on Control, Data Acquisition and Remote Participation on Fusion Research, to be published in Fusion Eng. and Design (1999).
- [7] D.P. Schissel, *et al.*, Rev. Sci. Instrum. **57**, 1932 (1986).
- [8] <http://d3dnff.gat.com/summaries/>
- [9] <http://d3dnff.gat.com/documentation/>
- [10] http://d3dnff.gat.com/code_usage/
- [11] <http://fusion.gat.com/comp/analysis/policy/data-policy.shtml>
- [12] Research Systems, Inc. (<http://www.rsinc.com>)
- [13] Q. Peng, *et al.*, "IEFIT — An Interactive Approach to High Temperature Fusion Plasma Magnetic Equilibrium Fitting," to be published in Transactions on Nuclear Science.
- [14] Platform Computing Corporation (<http://www.platform.com>)
- [15] Q. Peng, *et al.*, "A Linux Cluster for Between-Pulse EFIT and Other CPU-Bound Analyses at DIII-D," General Atomics Report GA-A23473 (2000).

- [16] M.L. Walker, *et al.*, “Automated Fault Detection for DIII-D Tokamak Experiments,” to be published in the Proc. of the 18th IEEE/NPSS Symposium on Fusion Engineering (1999).
- [17] E.B. Hooper, L.D. Pearlstein, R.H. Bulmer, Nucl. Fusion **39**, 863 (1999).
- [18] Polycom, Inc. (<http://www.polycom.com>).
- [19] <http://d3dpc1.gat.com/d3dtv>

ACKNOWLEDGMENTS

This is a report of work supported by the U.S. Department of Energy under Contract No. DE-AC03-99ER54463. Computer science work at the DIII-D National Fusion Facility is coordinated between the Computer Systems Group managed by Bill McHarg (mcharg@fusion.gat.com) and the Data Analysis Applications Group managed by David Schissel (schissel@fusion.gat.com). We would like to acknowledge stimulating discussions with T. Fredian, M. Greenwald, and J. Stillerman from Alcator C-Mod and with D. McCune from Princeton Plasma Physics Laboratory.