Data acquisition system for Alcator C-Mod

T. W. Fredian, J. A. Stillerman, and M. Greenwald
MIT Plasma Fusion Center, Cambridge, Massachusetts 02139

(Presented on 15 May 1996)

The Alcator C-Mod experiment requires an efficient data handling system to acquire and distribute the large number of measurements recorded during each pulse of the tokamak. Over 2000 signals amounting to over 80 megabytes of data are stored and distributed to 40 workstations for display and analysis. With machine pulses occurring every 10 to 15 min, most of the information must be distributed within 2 to 3 min after the pulse to enable the experimentalist to adjust the control settings for the next pulse. In addition to the local control room data distribution requirements, the system must also provide data access to remote sites which monitor diagnostics installed on the Alcator C-Mod experiment. This article will describe the hardware and software used to accomplish data handling tasks. © *1997 American Institute of Physics.* [S0034-6748(97)64801-7]

I. INTRODUCTION

The Alcator C-Mod tokamak fusion experiment generates over 80 megabytes of information for each pulse of the machine which occurs approximately four times each hour during normal operation. All of this information must be retained for future analysis and most of it must be delivered to scientists and engineers quickly after each pulse so they can perform analysis on the data and make decisions about operational parameter changes prior to the next machine cycle. The actual machine pulse lasts for only a few seconds but changes to any parameters must be completed 2 to 3 min prior to the next machine pulse. This leaves about 12 min for the information to be collected, distributed, and analyzed. Optimally, all data should be collected and distributed in less than 5 min with the more important signals becoming available for analysis within seconds after the machine pulse. After each machine cycle, several thousand plots of raw and analyzed data are generated on more than 30 computer workstations.

During the operation cycle of the experiment, several sequential operations must take place. Prior to the machine pulse, machine settings selected by the engineers and scientists must be loaded into the hardware that controls the operation of the machine. Also, all measurement recording devices (predominantly CAMAC) must be prepared to acquire data during the machine pulse. This stage of the cycle is referred to as initialization and generally takes about 1 to 2 min. Once this stage is complete, the machine goes through a check phase during which the power equipment is energized for the pulse and all systems are checked for readiness. This check phase takes 45 to 60 s. Next the machine pulse occurs, at which time all measuring devices begin recording various measurements at a fast rate, from a few thousand samples per second to as many as several million samples per second. This data is recorded in memory local to the measurement device and is not available for analysis until it is gathered by computers in the next stage of the cycle. The pulse stage lasts for only a few seconds. The next stage of the operation cycle is where the bulk of the data handling is performed. During this stage, referred to as the store phase, all the measurement data is collected from the measurement devices and made available for analysis. The large quantity of measurements must be collected, stored, and distributed quickly to enable sufficient time for analysis prior to the next machine cycle. After the data is analyzed, the scientists and engineers can make changes to settings to alter the machine operation of the next pulse. Once these changes are complete, another cycle can begin.

This article will describe the computer hardware and software used to collect, store, and distribute the information generated by the Alcator C-Mod experiment.

II. DATA ACQUISITION COMPUTERS

To collect and distribute the large quantity of data to more than 30 scientists and engineers, we utilize a cluster of high performance workstations linked with fast networking hardware. The majority of the systems are Digital Equipment AlphaStations[1] and VAXstations[1] using the open VMS[1] operating system. There are currently 24 workstations used for data analysis and 10 infrastructure systems which provide file, print, database, data acquisition, archival, and retrieval services. In addition, there are a dozen Xterminals which connect to the data analysis workstations providing additional graphic displays. Nearly all of these systems are on an uninterruptable power supply to maintain high availability.

The main computer system providing most of the data acquisition and file serving functions is an AlphaServer system with 128 megabytes of RAM and 12 gigabytes of magnetic disk storage. The disks are arranged in a redundant array of independent disks, level 5 (RAID5) which provides high speed input/output (I/O) ports and protection from a drive failure. Also attached to this system are two Kinetic Systems SCSI CAMAC serial highway drivers providing fast access to the CAMAC based measurement devices. The entire data acquisition process is controlled by applications running on this server.

Two slower VAXserver systems with Q-bus based CAMAC serial highway drivers are used to acquire data from two more CAMAC highways. Two additional servers, one DECstation and one VAXstation, perform as boot servers and distribute the operating system and applications to the other servers and workstations. Another VAXserver system is used to manage an optical archive of previous Alcator C-Mod pulses. One more VAXserver system is used to communicate with a hybrid plasma position control system via bitbus.[2]

III. CAMAC SYSTEM

There are currently four CAMAC serial highways in use on the C-Mod experiment. Two of the highways are SCSI based and are attached to the main data acquisition AlphaServer. Two smaller highways are attached to slower VAXserver 3200 systems. In total there are approximately 50 CAMAC crates on the four highways containing over 350 modules. Of these 350 modules, 200 are digitizers, 75 are timing and clock modules, and the remaining are a mix of various special purpose modules. Over 1600 signals are recorded for each pulse of the machine.

IV. DATA ANALYSIS COMPUTERS

Twenty-four analysis workstations access the experimental data soon after each signal is written to disk by the data acquisition servers. These workstations are typically running MDSplus[3] and IDL[4] applications which display the signals to the user and/or analyze the data and store the results of the analysis back in the same disk files as the raw data. A typical workstation may display 50–100 plots within 2 to 3 min after the pulse. These workstations are a mix of VAXstations and AlphaStations running the OpenVMS[1] operating system. They are generally configured with 90–160 megabytes of RAM. Each system has a magnetic disk which is used for virtual memory paging.

In addition to the analysis workstations, there are 16 Xterminals and over 100 personal computers used for communicating with the machine and diagnostic control systems and for accessing the experimental data.

V. NETWORKING HARDWARE

All of the workstations, Xterminals, personal computers, and server computers are connected via ten megabit per second thinwire Ethernet. In addition, a 100 megabit per second FDDI optical network is installed connecting the AlphaStations, AlphaServer, and approximately half of the VAXstations (not all of the VAXstations are compatible with FDDI). To further improve the network bandwidth, a 12 port Ethernet switch was added to divide the Ethernet load among several segments. The personal computer systems used to communicate with the machine control subsystems are connected via a separate Ethernet (see Fig. 1).

The Ethernet attached to the analysis and server systems, connects to the MIT FDDI backbone which in turn connects to the ESnet T3 backbone. This provides a high speed wide area network connection for collaboration and remote control of diagnostics on the experiment. The remote control of the entire experiment was demonstrated using this wide area network when a group of C-Mod physicists traveled to Lawrence Livermore National Laboratory in California and remotely operated the Alcator C-Mod experiment.[5]

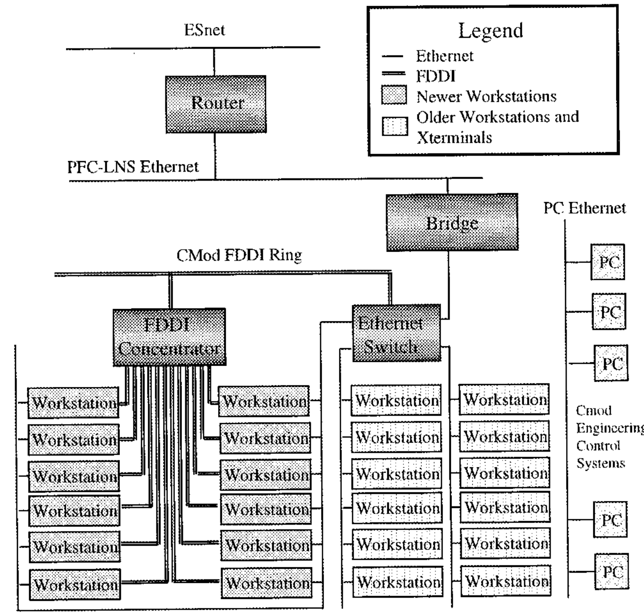

FIG. 1. C-Mod computer network.

VI. SOFTWARE

The main operating system used by the data acquisition and analysis system is Digital Equipment Corporation's open VMS. The predominant reason for selecting this operating system is that open VMS provides a robust mission critical environment. It is not unusual for these systems to run without rebooting for several months. In addition, open VMS provides record management facilities which allow a high level of file sharing. During the experiment cycle, programs on any of the data acquisition systems or data analysis systems can read and write data to the same files at the same time.

The workstations all use standard Xwindows and Motif as their predominant graphical user interface. Xwindows was designed with networking capabilities built in. This enables an authorized user anywhere on the Internet to run applications on any of the C-Mod open VMS systems and have the graphics appear on his workstation.

The data acquisition and analysis software used is the MDSplus Data Acquisition System.[3] MDSplus provides tools for acquiring data from sources such as CAMAC and storing the information in a hierarchical database. All of the information regarding data acquisition settings, calibration constants, and scheduling of data acquisition and analysis tasks are stored in the same database as the data acquired from measurement devices and the results of analysis. All the information is organized in a tree structure. The current C-Mod experiment tree contains over 30 000 nodes of structure and data. MDSplus also provides a complete set of graphical tools for configuring the data base, describing the data acquisition tasks to be performed and plotting the resulting information.

Next to the MDSplus applications, the most important tool used for analyzing the experiment data is interactive data language (IDL) produced by Research Systems, Inc.[4]

IDL provides interactive data manipulation functions with extensive graphics capabilities including $X-Y$ plotting, surface and contour plotting, and numerous image processing procedures. It also includes procedures for creating widget based user interfaces with menus, pushbuttons, and other controls normally found in windows applications. The IDL application has been closely integrated into the MDSplus system. MDSplus data can be read directly in an IDL application, then manipulated, displayed, and written back to the MDSplus database. Users can also write analysis applications with no user interface using its high level language and have these analysis applications automatically executed by MDSplus. A large percentage of the analyzed data generated from the experiment is computed in this manner.

In addition to the data stored in the MDSplus hierarchical database, a subset of analysis results, comments, and measurements is stored in a relational database. The database software used is Oracle RDB (formerly Digital Equipment RDB). Loading a subset of the data into a relational database permits the user to perform quick queries to find all the pulses meeting some selection criteria. In addition to storing analysis results, RDB is used to manage a database of pulse file locations in an optical archive which is discussed in more detail below.

Another important application developed as part of the MDSplus system is an electronic logbook. This application integrates MDSplus, RDB, and IDL into a widget based application for entering and viewing textual information about the machine operation. Users can enter comments about certain aspects of the machine setup and performance and categorize the comments by subject matter.

VII. DATA ARCHIVAL SYSTEM

The data associated with the experiment pulses is stored in a three tier storage system. As soon as the data is acquired, it is stored on a RAID set of magnetic disks. This provides high speed read write access to the most recently acquired data. The data remains on this RAID set until it is archived later in the evening. Approximately 3 gigabytes of fast RAID storage are available for the current day's pulse data.

Each evening after the machine operation is complete, the pulse files created that day are copied to both a magnetic disk staging area and an optical archive area and removed from the RAID storage area. The magnetic disk staging area, with a capacity of 15 gigabytes, is used to house the most recently accessed pulse files. The optical archive area contains all of the pulse data collected during the life of the experiment. It currently consists of two optical jukeboxes with a total storage capacity of 330 gigabytes. As this storage area begins to fill, it will be upgraded to increase the capacity. Currently a little over half of the optical storage is used containing pulse data from the beginning of operation in 1991 through the last machine operation in March 1996. It is estimated that the remaining storage will be consumed in one to two years.

The archival and retrieval of pulse data is performed automatically. An archive process runs each evening and moves shots from the fast RAID storage to both staging storage and optical storage. When free space in the staging area falls below a specified threshold, an automatic skimming process begins to remove pulse data from the staging area. A relational database is maintained containing pulse file location and access statistics to provide information to guide the retrieval and skimming operations. The skimmer looks at this access information and determines which files on the staging area were accessed the furthest in the past and whether or not they have been modified. The oldest accessed files are removed from the skimming area to free up space. If the data had been modified, the modified data is written to the optical storage area. Otherwise it is just deleted. If a user attempts to access data which is not on magnetic storage, a server process is alerted to retrieve the data from optical storage and place it in the staging area. The requested files are retrieved automatically, usually within 15–60 s. There is no operator intervention and the user does not need to explicitly request the retrieval.

VIII. COLLABORATION SUPPORT

The C-Mod data acquisition and analysis system is particularly suited for outside collaboration. As mentioned previously, there is a high speed Internet connection to the data storage, and the Xwindows utilities can be used by authorized users anywhere in the world. In addition, the MDSplus data acquisition system allows easy integration of new data acquisition equipment whether it be new types of CAMAC modules, other hardware architectures such as bitbus, GPIB, or even other computers such as Macintosh's or PC's. MDSplus also provides program libraries for UNIX, Windows, and Mac which use TCP/IP sockets to connect to server processes on the OpenVMS server systems. Users can develop programs on their choice of system that can store and retrieve data directly into or from the MDSplus database residing on the open VMS cluster. One of the most used applications in MDSplus, dwscope, has been ported to UNIX and accesses the data using the TCP/IP sockets interface. The dwscope program is a tool for plotting $X-Y$ traces and has features such as panning and zooming.

IX. DATA ACQUISITION/ANALYSIS PERFORMANCE

When the data acquisition system for C-Mod was originally designed in the 1987–1988 time frame, it was estimated that C-Mod would eventually generate 10–15 megabytes of data per machine pulse. This ''maximum'' rate was exceeded shortly after startup in 1991 and has continually grown to over 80 megabytes per pulse in 1996 (see Fig. 2).

The distributed nature of the system has allowed us to utilize new networking and computer technology without changing much if any of the software developed for the system. Even though the data throughput has grown to 800% of the original design goal, the system remains responsive. As the data handling requirements have increased, we have been able to add new network hardware, faster workstations and file servers, and faster disk subsystems to maintain a constant level of responsiveness. Figure 3 illustrates how some of these hardware improvements have impacted our data acquisition rates.

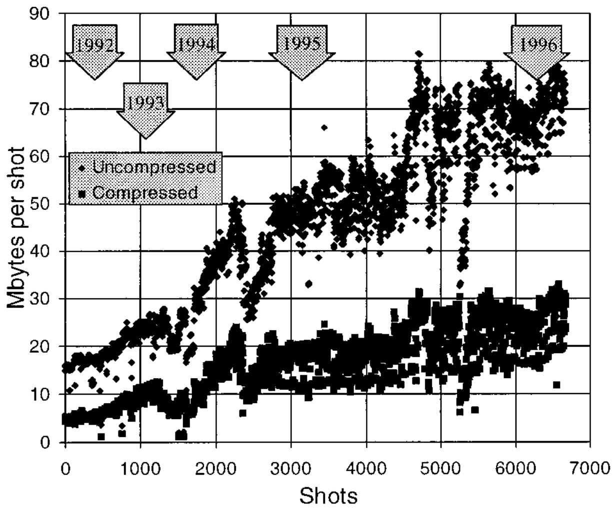

FIG. 2. C-Mod data acquisition quantities.

The system has also provided sufficient performance in distributing the data to the physicists and engineers. Typically there are 50–100 traces displayed on each of the 40 workstations and Xterminals within a few minutes after the pulse. Critical signals are displayed within several seconds.

X. FUTURE PLANS

Some of the next steps planned in the software area are centered around the enhancement of MDSplus. A port of MDSplus to both Windows and UNIX is planned. This would allow a broader selection of hardware and software platforms for future systems. In addition it will allow remote collaborators to have faster access to C-Mod data when using nonopen VMS systems.

We also hope to enhance the World Wide Web interface to the C-Mod operation. This would include Java display applets that would allow authorized collaborators to plot and manipulate C-Mod data directly from their web browser. The addition of web pages containing live updating graphics, video, and audio of the C-Mod experiment is also planned.

To maintain our current data acquisition and analysis performance levels with ever increasing amounts of experimental data, we will be gradually upgrading our older computer equipment with more powerful systems. Additional FDDI networking will be added when the older computers are replaced. Additional disk servers with RAID disk sets will be added as needed. The q-bus CAMAC systems will be replaced with faster SCSI based CAMAC systems.

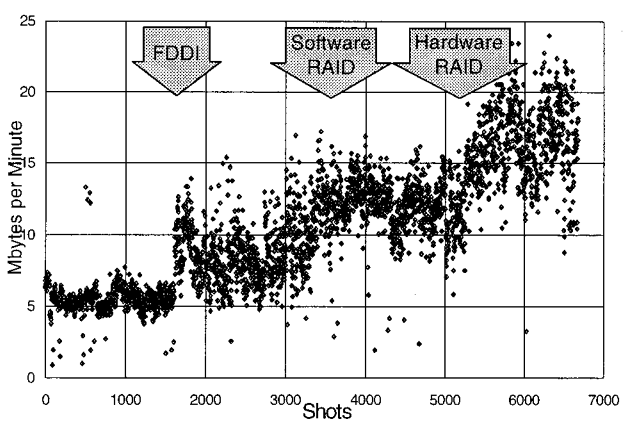

FIG. 3. C-Mod data acquisition rates.

XI. SUMMARY

The data acquisition and analysis system implemented for Alcator C-Mod records and distributes the experimental data in an efficient and timely manner. The system design has been expandable to handle the increased data loads which far exceed the original expectations. The entire collection of data acquired, both measurements and analysis results, during the lifetime of the experiment will be available online. Both the hardware and software chosen have been very reliable with most systems providing uninterrupted service for several months at a time. It is anticipated that the system can be expanded to handle any additional data loads that may occur from various sources such as new diagnostics or new analysis programs.

[1] VAX station, AlphaStation, and OpenVMS are registered trademarks of Digital Equipment Corporation.
[2] P. F. Isoz, J. B. Lister, and P. Marmillod, in Proceedings of the Symposium on Fusion Technology, 1264 London, 1990 (unpublished), p. 1264.
[3] J. A. Stillerman, T. W. Fredian, K. A. Klare, and G. Manduchi, Rev. Sci. Instrum., these proceedings.
[4] IDL is registered trademark of Research Systems Inc.
[5] S. Home, V. Bertolino, J. Daigle, T. Fredian, M. Greenwald, S. N. Golovato, I. Hutchinson, B. LaBombard, J. Stillerman, Y. Takase, S. Wolfe, T. Casper, D. Butner, W. Meyer, and J. Moller, Fusion Technol. (submitted).