CONF-831203--140

UCRL--89222

DE84 012601

The Local Area Network for the Plasma Diagnostics System of MFTF-B

Norman H. Lau
Earl G. Minor

Proceedings from 10th Symposium
on Fusion Engineering
Philadelphia, PA

December 5-9, 1983

DISCLAIMER

THE LOCAL AREA NETWORK FOR THE PLASMA DIAGNOSTICS SYSTEM OF MFTF-B

Norman H. Lau, Earl G. Minor
Lawrence Livermore National Laboratory
P.O. Box 5511, L-634
Livermore, CA 94550

Abstract

The MFTF-B Plasma Diagnostics System will be implemented in stages, beginning with a start-up set of diagnostics and evolving toward a basic set. The start-up set contains 12 diagnostics which will acquire a total of about 800 Kbytes of data per machine pulse; the basic set contains 23 diagnostics which will acquire a total of about 8 Mbytes of data per pulse. Each diagnostic is controlled by a "Foundation System" consisting of a DEC LSI-11/23 microcomputer connected to CAMAC via a 5 Mbits/second serial fiber-optic link and connected to a supervisory computer (Perkin-Elmer 3250) via a 9600 baud RS232 link. The Foundation System is a building block used throughout MFTF-B for control and status monitoring. However, its 9600 baud link to the supervisor presents a bottleneck for the large data transfers required by diagnostics. To overcome this bottleneck the diagnostics Foundation Systems will be connected together with an additional LSI-11/23 called the "master" to form a Local Area Network (LAN) for data acquisition.

The Diagnostics LAN has a ring architecture with token passing arbitration. It has a hardware data rate of 10 Mbits/second. A "wire center" is used to create a star shaped ring so that all nodes connect at a central point. Each node must actively join the ring or it is automatically bypassed by the wire center. The ring connects to the supervisory computer through a high speed parallel DMA link in the master LSI-11/23. A program in the master imposes network-level arbitration to control the transfer of data to the supervisor. The data transfer must satisfy three requirements; it must complete within 60 seconds, a sub-set of the data, called quick-look data, must be given priority, and the transfer must be fault-tolerant, so that if some data is unavailable the remainder is still transferred normally. In addition, the Diagnostics LAN must provide the supervisor with a way to redefine the data set to be transferred during the four minutes from the end of one transfer to the next machine pulse.

"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48."

Introduction

The Plasma Diagnostics System (PDS) is one of many major engineering systems of MFTF-B, such as the Magnet System, the Neutral Beam Power Supply System, the Cryogenic System, the Local Control and Instrumentation System (LCIS), the Supervisory Controls and Diagnostic System (SCDS), etc. The complexity of the PDS requires a decentralized multi-level control system, just as other major systems of MFTF-B. The higher level control and data acquisition for the PDS is handled by a distributed computer network described in detail in reference 1. The architecture of the lower level or local control level is based on the Foundation System (reference 2) which forms the fundamental data gathering and control system of MFTF. The Foundation System was successfully used in the Technology Demonstration which was completed in February, 1982.

The Foundation System consists of a Local Control Computer (LCC), a DEC LSI-11/23 microcomputer, connected to CAMAC crates via a 5 Mbits/second serial fiber-optic link and connected to the Diagnostic Data Processor (DDP), one of the SCDS computers (Perkin-Elmer 3250), via a 9600 baud RS232 link. The RS232 link is used for carrying control commands, acquiring data, and down-loading programs from SCDS. However, it creates a bottleneck for the large data transfers required by diagnostics. To overcome this bottleneck a Fast Link is added to the diagnostics Foundation Systems to be used for data transfers, while the RS232 link remains for carrying control commands and down-loading programs from the DDP.

System Requirements

There are 12 diagnostics in the start-up set of the PDS. Each diagnostic requires at least one LCC. Two of the twelve diagnostics do not have data transfers between shots, while the other ten have a total of about 800 Kbytes of data between shots. Only one DDP will be used for the start-up set. In the basic set the number of diagnostics will gradually increase to 23 with a total data of about 8 Mbytes per shot. Two DDPs will be used in the basic set, each acquiring 4 Mbytes of data. There are two different classes of data, quick-look data and regular data. The quick-look data has higher priority for display than the regular data and must be displayed within about 15 seconds. After the shot, the system will acquire quick-look data before regular data. The quick-look data is a small subset of diagnostic data (about 50 Kbytes) and requires little processing while the regular data are for archiving and display. All quick-look data and regular data must be moved from CAMAC crates to the memory of the DDP within 60 seconds after the end of the shot, corresponding to an overall data transfer rate of 66.7 Kbytes/second per DDP.

The 9600 baud link in the Foundation System is too slow for data transfer for diagnostics. It would require more than 3.5 minutes to move all start-up data to the DDP, and more than 28 minutes to move all basic set data to the DDP. Therefore, a Fast Link between the DDP and the LCCs for data transfer is needed. The general requirements for the Fast Link are:

- The hardware must be compatible with the LSI-11 Q-bus
- The overall effective data throughput must not be lower than 66.7 Kbytes/second. (Raw data rate more than 5.3 Mbits/second).
- Bit error rate must be less than 10E-9 at maximum data rate.
- A commercially available system is preferable
- Cost effective

Design Approach

Development of the architecture for the Fast Link has spanned twelve months of study, during which two workshops and a preliminary design review were held. The first workshop (November 10, 1982) considered several design topologies such as a star network in which the DDP connects to all LCCs with either a parallel or serial link, and a linear bus/multi-drop system in which the DDP and all LCCs attach to the same parallel-line or single-line bus. The former scheme was rejected because of potential reliability problems and large software overhead in the DDP. The parallel bus multi-drop was also rejected due to low reliability and flexibility, especially when the system expands to the basic set. As a result of the first workshop the architecture of the Fast link was established as shown in Figure 1. A Master LCC (MLCC) was added as a host

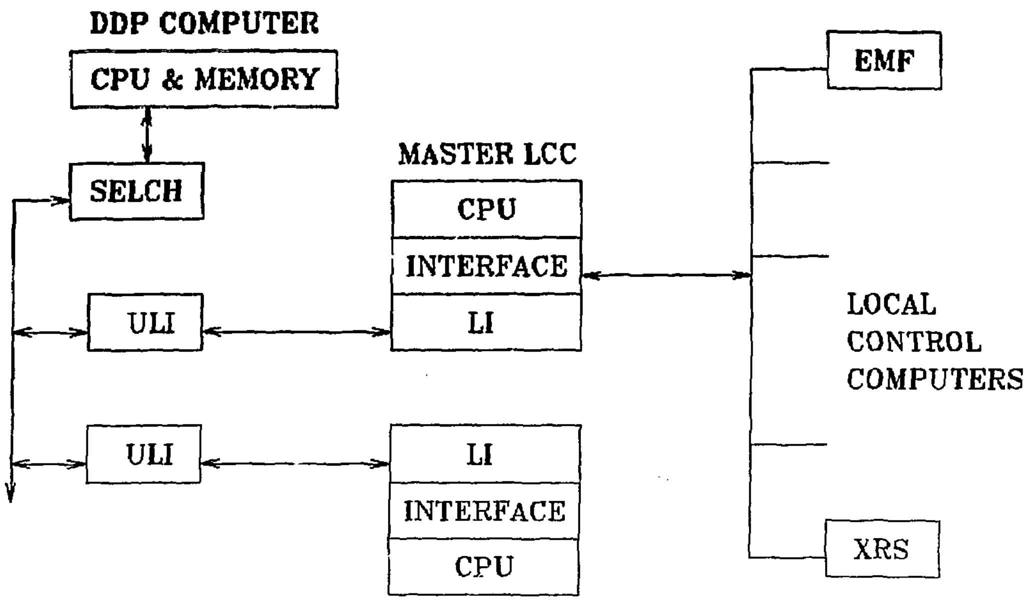

Figure 1. Multi-drop System

interface to the DDP. The MLCC forms a serial bus/multi-drop system with all other LCCs. A second off-line MLCC was added to the system to serve as a back-up in case the on-line MLCC fails. The search for a commercial system for the serial bus/multi-drop system for LSI-11/23 was on.

The second workshop (March 17, 1983) reviewed Local Area Network (LAN) technology to-date and compared several LANs on the market. As a result of the workshop a detailed study was done of the Pronet Token Passing system and of an Ethernet system. The Pronet Token Passing system was selected for the following reasons:

- Roughly equivalent software effort for low level driver
- Higher throughput for Pronet because of larger packet size and smaller header.
- Ring protocol is well suited to a polling design.
- It can be interconnected via fiber optic cable.
- Pronet power consumption is lower.
- The ring network is passive so that a node failure will not bring the whole network down.

A prototype for the Fast Link using the Pronet LAN was built, and test results were presented in a Preliminary Design Review (September 28, 1983). Favorable comments were received and final design is in progress.

Description of the LAN

Figure 2 shows a block diagram of the Fast Link prototype, with two MLCCs connected to the DDP with parallel links, and forming a LAN with all diagnostic LCCs. RS232 links between the DDP and each of the LCCs, not shown in the diagram, are used for program down-loading and control command transfer. There are also RS232 links between the DDP and the MLCC which will be used for program down-loading and sending error messages. The LAN as shown in Figure 2 is used for data transfer from the LCCs to the DDP. The DDP and all nodes connecting to the LAN are located in the MF1F Control Room of Building 439. Data from diagnostic instruments is buffered in CAMAC modules located in the machine area of Building 431. These two areas are separated by about 75 meters and communicate via fiber-optics. After a shot each LCC will acquire data from the CAMAC modules and transfer it over the LAN to the MLCC. The MLCC will then transfer the data to the DDP.

The interface between the DDP and the MLCC is a parallel link comprised of a Universal Logic Interface (ULI) module for the Perkin Elmer 3250 computer and a Link Interface (LI) module for the LSI-11 microcomputer. The ULI and the LI are 16-bit by 128 words First-In-First-Out buffers plus some control lines in both directions. They support DMA transfers and they were designed in-house.

The Pronet LAN, manufactured by Proteon Associates, employs ring network architecture and token arbitration. It links the two MLCCs and all LCCs through a wire-center. The wire-center creates a star-shaped ring which is a physical rearrangement of the classic ring architecture such that all nodes are interconnected at a central point. The network data transmission rate is 10 Mbits/second and it supports up to 255 nodes per ring. Rings may be interconnected via gateways to further expand the capability of the network. Each node consists of two hardware modules: the Ring Control Board and the Host Specific Interface Board. Together these modules provide

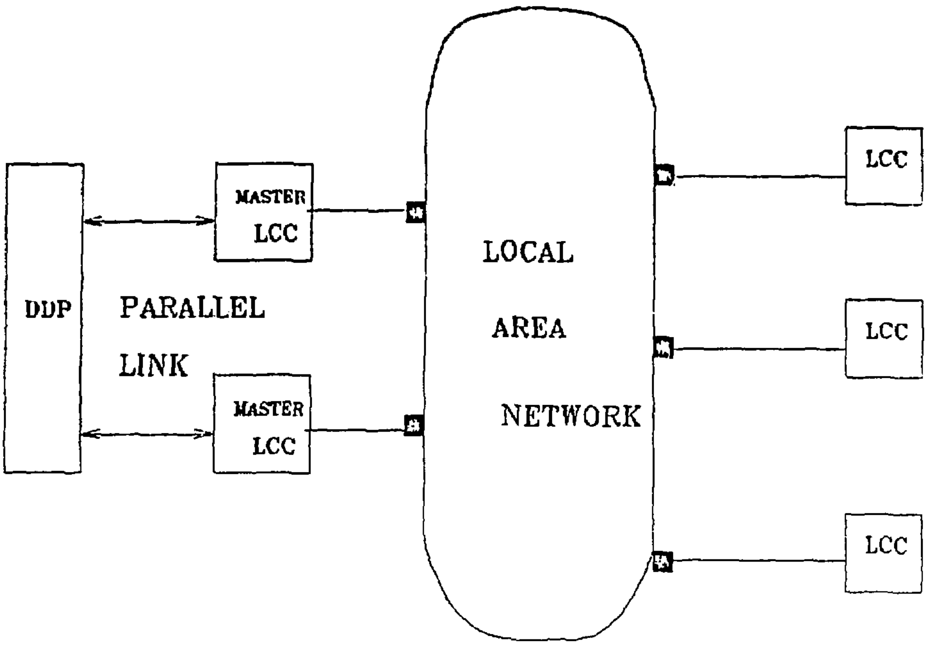

Figure 2. Fast Link System Overview (Per DDP)

the two lowest layer protocols (Physical Layer and Link Layer) as defined by the International Standards Organization Open Systems Interconnection reference model. They support LSI-11 Q-bus, 22-bit extended address, DMA and full duplex.

The Pronet's data format consists of packets of data up to 2044 bytes long. Each packet contains a Beginning of Message character which marks the start of a new message, destination address, source address, a variable number of data bytes, the End of Message character which marks the end of data within the message, the status bits "link parity" and "refused", and a token. The parity bit for each message is checked by the copying node. The refused bit indicates that a message destined for a particular node was not copied into its packet buffer. The originating node may try to retransmit the message. When idle, the ring network circulates a token control character around the ring. If a node has a message to send, it will look for the token and convert the received token into a Beginning of Message control character by changing its last bit, marking the beginning of a message. The conversion of the token into Beginning of Message ensures the orderly queueing among nodes wishing to place messages on the ring, thus avoiding collision and increasing network bandwidth. A token sent at the end of the message will allow nodes downstream of the originating node to place messages onto the ring if desired.

Test Results

A ring network of four nodes was purchased for the prototype study. By using vendor software test programs and handler (operated under RT-11) the interface boards ran successfully under analog loopback and digital loopback tests. One of the vendor's test programs, Special Transmit/Receive, was modified so that it could run continuously and check for error rate. This test was used to exercise the hardware using no communication protocols. The transmitting node sent a fixed pattern of data, and the receiving node checked all the incoming data for errors. There was no checksum generated and checked, except a single parity check in the data frame. With all the tests run so far, few errors have been detected. The longest run between two nodes was about 90 hours for over 100 billion bytes of data transfer without error. Data throughput between the DDP and the MLCC was limited by the LSI-11 Q-bus which has been estimated at about 400 Kbytes/second. Data throughput between two nodes in the LAN has been calculated at about 269 Kbytes/second. When two pair of nodes exchange data simultaneously, it has been found that the overall data throughput is almost double. This result was expected.

Software Design

The software for the PDS Fast Link has the job of coordinating the many computers and communication links in the network to form an integrated system. It must manage the LAN, providing for efficient data transfer from the LCCs to the MLCC. It must handle the parallel link to the DDP, implementing both the low level control required by the parallel link hardware and the high level protocol required by programs on the DDP. It must arbitrate the data transfer so that quick-look data is given priority over regular data. It must be flexible enough to allow data collection lists to be changed easily, and fault-tolerant enough that hardware failures cause only limited damage. Most important, it must meet the basic throughput rate of 66.7 Kbytes/second.

The design of the Fast Link software was shaped by a few basic decisions. The first decision was made by considering how and when data transfers would occur. MFTF will operate in pulses or "shots", and data will be transferred at the end of each shot. Thus all of the LCCs will be ready to start sending data over the LAN at the same time. Furthermore, they will all want to transfer to the same two nodes, the two MLCCs. It appeared that this type of loading would be highly prone to contention, so the decision was made to avoid contention by imposing a master/slave protocol between the MLCC and the LCCs. This decision then influenced the choice of a LAN; a master/slave protocol could be imposed on a contention type network, but it is more natural to a token ring.

A second important decision concerned how to implement the quick-look prioritizing of the data acquisition. It was decided to divide this task into two parts, with one part carried out at the LCC level and the other part assigned to the MLCC. Each LCC would be responsible for prioritizing it's own data acquisition, always making quick-look data available for transfer over the LAN in preference to regular data. The MLCC would then poll the LCCs (in keeping with the master/slave protocol mentioned above) and give preference to quick-look data. With the quick-look task partitioned in this way it is not necessary for the MLCC to have a list of data to be acquired. Instead the MLCC can act almost as a 'channel', simply moving the highest priority data first. The data list is distributed among the LCCs, contributing to overall system flexibility and fault tolerance.

A final important design decision concerned how the two MLCCs in the network would coordinate their operation. The original concept for the second MLCC was purely as a hardware backup; if the primary MLCC failed the Fast Link would go down briefly while the backup was brought on line. But a much more desirable option would be to design the second MLCC into the system as a fully redundant component. Then if one MLCC failed the other could take over it's job without any interruption or human intervention. To achieve this goal it was decided to have both MLCCs perform data acquisition during normal operation. For the times when both MLCCs are working this approach has the benefits of providing considerably higher data throughput, and avoiding single point failure.

The basic decisions described above provide a broad framework for the Fast Link software design; the next level of design detail is contained in the high level LAN communications protocol and in the MLCC polling algorithm. The LAN communications protocol is summerized in Table 1. The number of different types of messages which the protocol allows is very small. A transaction consists of an MLCC query for data of a given priority, followed by an LCC reply indicating either that no

MLCC queries	LCC replies
Send {Regular / Quick-Look} Data	No data ready My priority is {Quick-Look / Regular / Idle}
	Data of length x follows. My priority is {Quick-Look / Regular / Idle} + variable length data message
Reset	Reset acknowledged My priority is {Quick-Look / Regular / Idle}

Table 1. High level LAN protocol

such data is available, or that the data will follow in a variable length transmission. Every reply message also updates the MLCC as to the LCC's "priority level". The priority level is defined as the highest priority data still remaining to be acquired by the LCC. Note that it is possible for an LCC to be at, say, quick-look priority level but still not have quick-look data available for transfer. This could happen if, for example, the data had not yet completed transfer from CAMAC.

The MLCC polling algorithm uses the LAN communications protocol to coordinate data collection from the LCCs. The polling algorithm is outlined in Figure 3. The ring is a data table in the MLCC containing one entry for each LCC. Each entry contains the LCC number (node address in the LAN) and the LCC priority most recently received in a

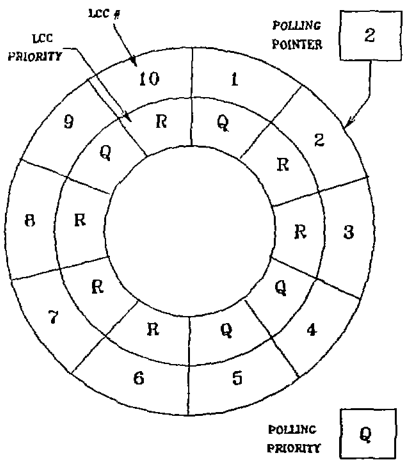

BEGIN: START AT POLLING POINTER

POLL LCC's AT "PRIORITY" OR ABOVE FOR DATA OF LEVEL "PRIORITY"

IF NO DATA, REDUCE "PRIORITY" AND BEGIN

IF DATA, TRANSFER IT AND MOVE POINTER TO ONE PAST THAT LCC. REASSIGN "PRIORITY" AND BEGIN.

Figure 3. Polling algorithm

reply message. A pointer (called the polling pointer) indicates where in the table the next polling cycle begins. A global variable, called the polling priority, is used to select which LCCs will be queried for data in a polling cycle. At the beginning of each polling cycle the polling priority is reassigned to equal the highest priority LCC in the table. The pseudo-code gives details of how the polling cycle works.

The software design for the Fast Link recently passed a Preliminary Design Review and is now being refined to a final design level. This software should be completed and running by August 1984.

Summary

The implementation of a LAN simplifies the system design of the Fast Link which provides a fast data transfer path from the LCCs to the DDP. It has been estimated that by using the system as described above the overall data throughput from the CAMAC crates to the memory of the DDP is 95-123 Kbytes/second for one MLCC, and 174-216 Kbytes/second for two MLCCs. This system not only has met all the basic requirements, but can accommodate some additions of instruments or diagnostics beyond the basic set.

References

1. R. H. Wyman, et al., "An Overview of the Data Acquisition and Control System for Plasma Diagnostics on MFTF-B", paper in this conference.

2. N. H. Lau, "Foundation System of the Local Control and Instrumentation System for MFTF", Proceedings of the 8th Symposium on Engineering Problems of Fusion Research, pp. 2009-2010, November 13-16, 1979.