LEGIBILITY NOTICE

A major purpose of the Technical Information Center is to provide the broadest dissemination possible of information contained in DOE's Research and Development Reports to business, industry, the academic community, and federal, state and local governments.

Although portions of this report are not reproducible, it is being made available in microfiche to facilitate the availability of those parts of the document which are legible.

3

LA-UR -87-1497 (FINAL)

CONF-8711552--26

LA-UR--87-1497-Final

DE88 000520

TITLE: THE MEGA DATA ACQUISITION SYSTEM

AUTHOR(S): Michael A. Oothoudt, James F. Amann, Richard D. Bolton, Martin D. Cooper, Will M. Foreman, Gary E. Hogan, Thomas Kozlowski, Franklin J. Naivar, Wayne Smith, E. Barrie Hughes, Charles C. H. Jui, John N. Otis, Peter S. Cooper, Kevin Hahn Rochelle B. Lauer, John K. Markey

SUBMITTED TO: 5th Conf. on Real-Time Computer Applications in Nuclear Particle and Plasma Physics,
May 12-15, 1987
San Francisco, CA

THE MEGA DATA ACQUISITION SYSTEM

Michael A. Oothoudt, James F. Amann, Richard D. Bolton,
Martin D. Cooper, Will M. Foreman, Gary E. Hogan,
Thomas Kozlowski, Franklin J. Naivar, Wayne Smith
Los Alamos National Laboratory*
Los Alamos, N.M. 87545

E. Barrie Hughes, Charles C. H. Jui, John N. Otis
Stanford University**
High Energy Physics Laboratory
Stanford, Ca. 94305

Peter S. Cooper, Kevin Hahn, Rochelle B. Lauer, John K. Markey
Yale University***
Department of Physics
New Haven, Ct. 06511

The MEGA experiment [1] will acquire 3 MB/sec of data into a 9-crate Fastbus system controlled by a single Fastbus master. A second-level trigger will be implemented in a Fermilab Advanced Computer Program (ACP) [2] farm of 32 Motorola 68020 microprocessors. Output from the ACP farm at a rate of 24 KB/sec will be sent to a MicroVAX II for taping.

Overview

The MEGA experiment (Muon decays to Electron and GAmma) at the Clinton P. Anderson Meson Physics Facility (LAMPF) will attempt to improve the limit on the branching ratio for the decay of the muon to an electron plus gamma ray from 5*10**-11 to 1*10**-13. Decays of 3*10**7 muons/second will be observed for 1.2*10**7 seconds. Due to the LAMPF beam structure (120 beam bursts per second with a 6 percent duty factor), this corresponds to an instantaneous decay rate of 5*10**8 muons per second. A hardware trigger will reduce the average rate to 2400 events/sec.

During a 500 microsecond beam burst, events are buffered in memories of Fastbus ADCs, TDCs, and latches. Deadtime is reduced by double buffering Fastbus modules where necessary. Each event will average 1400 bytes, with an average of 20 events acquired during each beam burst. Figure 1 shows an overview of the data paths in the experiment.

In the 7.8 milliseconds between beam bursts, a Fastbus master (the CERN-designed GPM [3]) will read the Fastbus memories and dump all data from a beam burst into one node of a farm of 32 Motorola 68020 microprocessors running the Fermilab Advanced Computer Program (ACP) software. The farm will reconstruct the events sufficiently so that 99.5% of the least promising candidate events can be discarded. The raw data for the remaining events plus related calculated quantities will be sent to the host MicroVAX II from the ACP system at a rate of 24 KB/sec. In the host, the LAMPF standard Q Data Acquisition system [4] will write the Fastbus data to tape along with low rate data (scalers) from CAMAC and from an Environmental Monitor system. Approximately 2000 6250-bpi tapes will be produced for further off line analysis.

Manpower for development of the Fastbus, ACP, and MicroVAX II hardware is limited to less than one man year; manpower for related software development is limited to less than five man years. Thus the system was designed to use commercial hardware and existing software where possible. Software is being reused from Fermilab (for software management tools and for the ACP system), CERN (for some of the Fastbus related software), and LAMPF (for the MicroVAX II host and CAMAC).

Fastbus Hardware

The Fastbus system for MEGA consists of 9 crates of ADCs, TDCs, and latches controlled by a

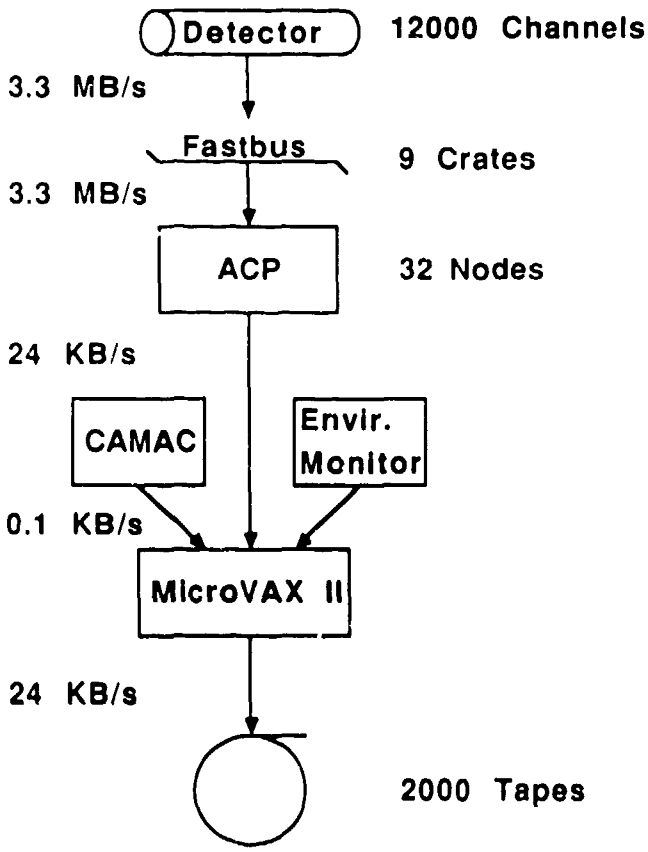

Figure 1 Overview of Data Paths in MEGA Experiment.

CERN designed Fastbus General Purpose Master (GPM). A Fermilab designed Fastbus to Branch Bus Controller (FBBC) provides a data path from Fastbus to nodes in the ACP farm. Table I lists measured and estimated delays and response times associated with components of the Fastbus system. A diagram of the Fastbus system is shown in Figure 2.

Data Acquisition Modules

The MEGA detector consists of Multiwire Proportional Chambers, Drift Chambers, and Scintillators for a total of approximately 12000 data channels. We have adopted a proposal by Phillips Scientific [5] for 1000 channels of Fastbus ADCs, 1200 channels of Fastbus TDCs, and 9500 channels of Fastbus latches.

The TDC proposed by Phillips Scientific contains

Table I - Fastbus Component Timing

Module	Operation	Time (ns)
GPM	Block write	135
	Block read (overlapped)	110
	Block read (no overlap)	250
Latch, ADC, or TDC	Block read	<110*
FBBC	Block write	<190*
SI + 3 meter cable	Signal delay	190
SE + 1 meter cable	Signal delay	70
Cable Segment	Signal delay per meter	~5
Crate Segment	Signal delay	**

-----
*Estimate based on information from module designer.
**Included in measured times for modules.

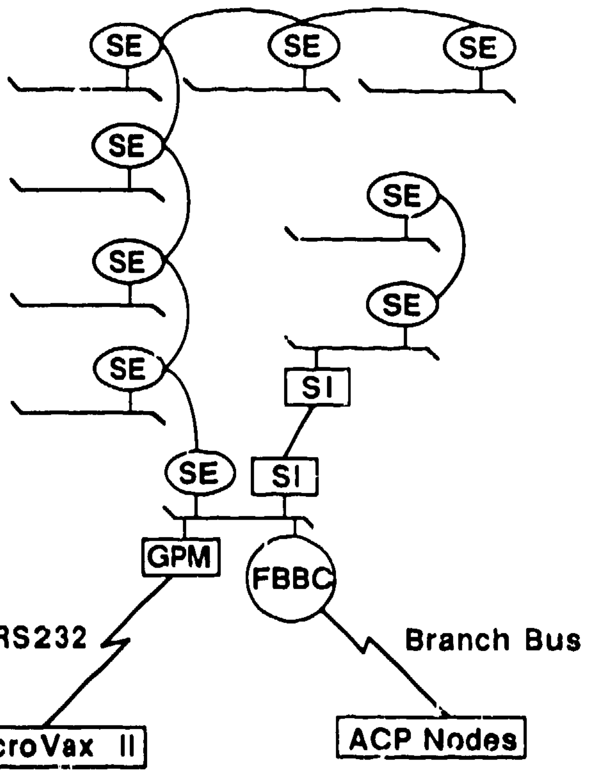

Figure 2 Layout of the MEGA Fastbus System. Only control modules are shown. ADCs, TDCs, and latches may be located in any of the crates.

32 channels with a range of either 128 ns or 1024 ns and a resolution of 0.1 percent. The proposed ADC contains 32 channels with 12 bit resolution and 512 pC full scale. Both modules digitize an event within 7 microseconds into LIFO memories 256 events deep with automatic pedestal correction on a channel by channel basis. Upper and lower cuts on a channel by channel basis may be used to perform a sparse data scan compaction.

The latches proposed by Phillips Scientific are 128 channel latches with one type providing sparse data compaction. The latches convert and store data within 0.5 microseconds in LIFO memories 2048 events deep (nonsparse latch) or 1024 events deep (sparse scan latch).

Readout of the ADCs, TDCs, and latches from Fastbus is done by the GPM with Fastbus block reads with handshake. However, the experiment has a relatively small amount of data (28 KB/beam burst) spread over approximately 200 modules. Setup time by the Fastbus master to address each module would consume an appreciable fraction of the time between beam bursts (more than 2 ms out of 7.8 ms available). Introduction of the "MEGAblock" readout scheme by Phillips Scientific, however, greatly reduced this overhead. Modules in a single crate are daisy-chained together using the Fastbus daisy-chain lines DLA/DRA and DLB/DRB. Readout is started with the first module in the chain by the standard primary and secondary address cycles. When readout of the first module completes, a "token" is automatically passed to the next daisy-chained module and it responds to the next DS cycle from the master. The token is passed through all modules until the the last module in the chain signals "End Of Block" when data is exhausted. Thus it is possible to read the entire MEGA Fastbus system with only one setup overhead per crate. This reduces the setup overhead from more than 2 ms to less than 0.1 ms.

A prototype Phillips nonsparse latch has been received for testing. Production units of this latch and the TDC are expected to be ready by Fall 1987.

Crate Connection Architecture

The nine crate Fastbus system is connected by a combination of Segment Extenders (SEs) and Segment Interconnects (SIs). SIs [6] are very flexible but are relatively slow (see Table I) and expensive. SEs [7] are cheaper and faster, but only allow chaining 7 crates together with only the first crate allowed to contain a Fastbus master. Figure 2 shows the architecture selected to provide a nine-crate system with the minimal cost and readout overhead. A Master crate contains the Fastbus master, the FBBC link to the ACP farm, an SE-based "extended backplane" of six Fastbus crates, and an SI connected to a secondary crate. An SE-based "extended backplane" goes from the secondary crate to provide the ninth crate. ADCs, TDCs, and latches may be in any of the 9 crates. Modules expected to have the most data per beam burst (such as nonsparse latches) will be placed in the Master crate or the SE crates to reduce readout time.

GPM

The GPM is a Motorola 68000-based Fastbus master designed at CERN and manufactured by CES [8] and DR Struck [9]. It is programmable in both assembler and higher level languages (FORTRAN and PASCAL) using CERN M68000 Cross Software Tools [10]. Two RS232 ports are provided for communication at speeds up to 19200 baud. The "termline" is intended for communication with the resident monitor software, MONICA [11], while the "hostline" is provided for communication with the host MicroVAX II computer. (A high speed parallel port on the GPM is not being used.)

In the MEGA experiment the GPM will be interrupted via a front panel NIM signal at end of beam burst. An optimized interrupt handler will read data from the Fastbus modules into local memory using the MEGAblock scheme. An idle node will be found (as described below) and the memory will be block written to the idle ACP node through the FBBC. Table I shows the speeds for reads and writes using the GPM along with the delays associated with SIs, SEs, and other elements. Based on these numbers, a typical beam burst of data (20 events of 1400 bytes each) will require approximately 2.5 ms to read data from memories in Phillips modules into GPM memory and approximately 2.4 ms to dump the data to an ACP node. This leaves roughly 2.9 ms per beam spill for polling the ACP nodes, performing run control commands, and other housekeeping chores.

A set of Fastbus standard subroutines has been written [12] to run in the GPM. Based on these routines and the CERN Remote Procedure Call software [13], Fastbus operations may be carried out from the MicroVAX II host at low speed. A detailed design of the high-speed data acquisition software for the GPM has been completed and implementation started. A design for Fastbus diagnostics is being worked on.

FBBC

The Fermilab Fastbus to Branch Bus Controller (FBBC) [14] was designed for use by the CDF experiment at Fermilab to interface Fastbus to the ACP microprocessor farm. Branch Bus is an RS485 bus used in the ACP system to connect multiple VME crates. The FBBC is a Fastbus slave and Branch Bus master (only a single master is allowed on the Branch Bus). The FBBC is not a programmable device and thus requires an intelligent master in Fastbus, like the GPM. As a Fastbus slave the FBBC responds to block read and write requests from the GPM to the node selected in an FBBC register. As a Branch Bus master, the FBBC initiates arbitration cycles for use of the VME bus in the crate containing the specified node and performs the requested read or write in blocks of 256 bytes. The FBBC is specified to have a transfer rate of 20 MB/sec, but we have as yet been unable to measure its speed in our system.

The FBBC is not commercially manufactured. We have fabricated two wire-wrap modules and with assistance from Fermilab are in the process of testing them.

Non-Fastbus Data Sources

In addition to the high-speed data stream described above, low-speed data will be acquired from CAMAC into the host microVAX II. The majority of the low-speed data will be scalers read every few seconds. In addition, a stand-alone computer will be used to monitor voltages, currents, and temperatures in the system. This Environmental Monitor will be periodically read through CAMAC. Since scaler and evironmental data is not associated with specific Fastbus events, the data is simply written to tape by the MicroVAX II as it is received from CAMAC.

ACP System

The Fermilab Advanced Computer Program [2] has developed hardware and software to allow inexpensive use of powerful 32-bit microprocessors for off-line analysis of high energy physics data. The system exploits the "trivial parallelism" of high energy physics data: events are independent of each other, allowing all calculations on a single event to be done completely within a single node. Thus while the ACP system is a "parallel processor" system, no communication between the nodes is required.

In a standard ACP system the user partitions his program into an input/output section to run in the host VAX and a compute intensive section that runs in the nodes. Code that runs in the nodes must be pure FORTRAN77 or be specially written in a language acceptable to the nodes. Code in the VAX reads data from tape, passes events to nodes, processes returned data, and (at end of run) sums any counters or histograms that may be in the nodes. Node code performs calculations on the buffers of data sent from the VAX and may return buffers of calculated quantities.

Hardware

The MEGA experiment uses the standard ACP hardware as manufactured by Omnibyte [15]. Figure 3 shows the layout that will be used.

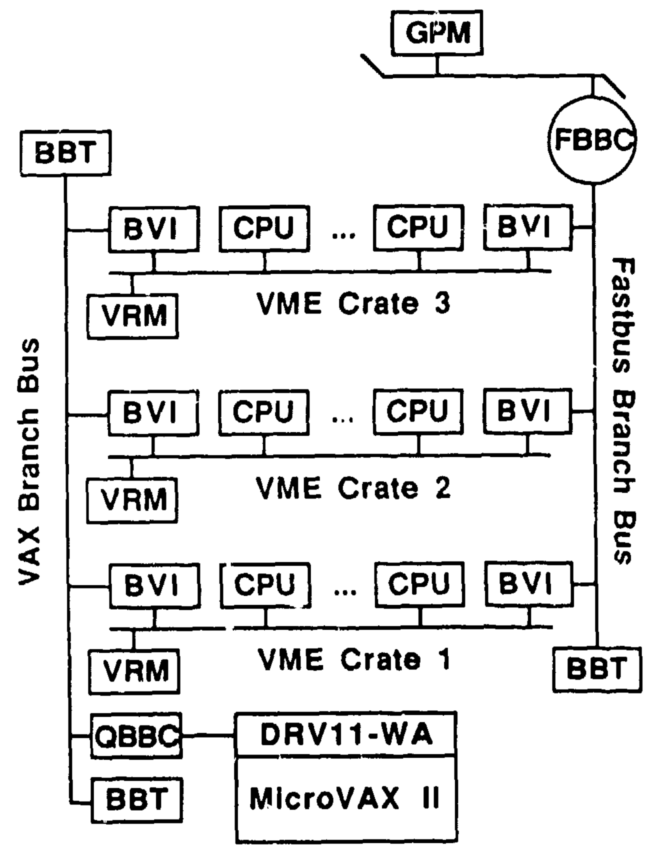

Figure 3 - Layout of the ACP system for MEGA. The QBBC and BBTs are in the VME crates, but only to obtain power.

The nodes consist of 16 MHz Motorola 68020 microprocessors with 68881 floating point coprocessors and 2 MB of memory. Up to 15 such nodes may be placed in each VME crate. A BVI (Branch Bus to VME Interface) in each crate serves as a "crate controller" from the Branch Bus. Two BVIs are required in each crate, since one Branch Bus connects to the FBBC and the other connects to the MicroVAX II. The MicroVAX connection also requires a QBBC (Q-bus to Branch Bus Connector) and a DRV11-WA. VRMs (VME Resource Module) are required in each crate to allow for MicroVAX polling of the nodes. BBTs are Branch Bus terminators.

We have received control modules and ten nodes from Omnibyte. The hardware has been exercised extensively and has worked very well.

Software

The ACP system supplies a library of software to support off line analysis of data, but does not directly support data acquisition. To learn the ACP system, we have converted the off line analysis code for Experiment 777 at Brookhaven to run on a standard ACP system. Based on this code we have preliminary performance figures for the processor boards in our ACP system. For the track finding code, which is dominated by integer operations, we measure 73% of the performance of a VAX 780. For fitting procedures, which are dominated by floating point operations such as matrix inversion, we measure 46% of a VAX 780's speed. No attempt has been made yet to optimize the code to take advantage of ACP node properties.

For data acquisition, the obvious solution of sending data from the GPM through the MicroVAX II to the ACP system at 3 MB/sec will not work because of bus saturation in the MicroVAX. However, ACP system software may still be used without modification because of the low rate of filtered events in the MEGA experiment from the ACP to the MicroVAX (24 KB/sec).

At start of run the MicroVAX sends each node a "dummy" event, causing MEGA control code in the node to run. The control code clears a mailbox word in the node's memory and loops, waiting for that word to change in value. When the GPM has a beam burst of data, it reads the mailbox words in the nodes' memories through the FBBC to find a ready node. It dumps the beam burst of data into a ready node and then changes the mailbox word. The node then processes the buffer, performing track reconstruction. If any good candidate events are found, the node returns them to the MicroVAX and waits for a dummy event. If good candidate events are not found, the node clears its mailbox word and waits for more GPM data.

This scheme works only because of the low rate into the MicroVAX. The MicroVAX must poll the nodes to determine that they have data to be returned. However, at 12 events per second (24 KB/sec) the polling overhead in the MicroVAX is acceptably low. GPM polling of the nodes is estimated to take 5-20 microseconds [16] per node or 160-640 microseconds to poll all 32 nodes. Since the GPM will have approximately 2.9 ms of idle time per beam burst, this is acceptable.

The control software in the ACP nodes builds events from the block of data sent by the GPM and passes them to the physics event reconstruction code. The reconstruction code attempts to find an electron and gamma ray in the detector with the appropriate energies and geometry. The electron of interest is tracing out a helix in a 1.5 Tesla magnetic field. In addition there are typically nine other electrons circulating in the chamber from events that did not fire the hardware trigger. Thus the event reconstruction is a complicated combinatoric problem. The goal of the event reconstruction is to reduce the rate of data to tape by a factor of at least 200. Reconstruction of events generated by Monte Carlo simulations show that the current algorithms can meet this goal in the time available in the ACP nodes, but only if beam intensity is reduced to half the desired level. A second generation of development is taking place to speed up the algorithm and reduce its memory requirements.

The control software that runs under the ACP system software has been written. This software handles the data flow from the GPM through the ACP node to the MicroVAX II. It is being tested now without direct input from the GPM. When the FBBC is fully debugged, testing will be continued with data from the GPM at the expected experimental rates.

MicroVAX II System

Software in the MicroVAX II provides the user interface for run control and monitoring the data. Interfaces are also necessary to the GPM, ACP, and CAMAC systems.

GPM Related Software

The MicroVAX II can only access the GPM and Fastbus at low speeds due to the RS232 interface. Thus the MicroVAX will generally ship commands to the GPM, for execution by GPM code. Run control and diagnostic commands from the MicroVAX will be implemented via the CERN-developed Remote Procedure Call (RPC) package [13], which is currently being adapted for use.

ACP Related Software

A single process is required in the MicroVAX II to communicate with the ACP nodes. This process is used to send dummy events to the nodes and receive good candidate events back for taping and distribution to other processes that wish to inspect the events. In addition this process is responsible for properly starting and ending runs in the ACP system.

This software has been written and is being tested with the ACP node software described above. We have measured the maximum tape writing speed to be 90 KB/sec for 1600 bit per inch tapes at 75 inches per second tape speed and estimate the speed will be in excess of 140 KB/sec for the 6250 bpi, 45 ips tape drives we plan to use. The maximum speed of interprocess event distribution has been measured on a MicroVAX II to be 350 KB/sec. These bandwidths are well in excess of the expected data rate of 24 KB/sec.

Q System

The Q system [4] is the standard data acquisition package at LAMPF used for CAMAC data acquisition. In the MEGA experiment Q is used as the Master run controller, the data acquisition system for CAMAC data, and as the manager of output tape writing, merging CAMAC and Fastbus data into a single data stream.

No modifications to the Q system were required to support the MEGA experiment. A few minor modifications to Q may be made to simplify some operations for MEGA.

Summary

An engineering run with a partial detector is scheduled for Fall of 1987. The run will test the response of Phillips TDCs and nonsparse latches, the performance of SEs and SIs, and the data path from Fastbus through the ACP system to the MicroVAX II. Debugging of the GPM and FBBC is currently in progress.

Software development for the GPM, ACP nodes, and MicroVAX II host is in progress and will be ready for use by the Fall run.

Production data taking will start in late 1988.

*Supported by U. S. Department of Energy Contract W 7405 ENG 36.
**Supported by U. S. NSF Grant PHY 84 11168.
***Supported by U. S. Department of Energy Grant DE FG02 87ER40313.

References

[1] The MEGA experiment is LAMPF Proposal 969 and is a collaboration of Los Alamos National Laboratory, Princeton University, The University of California at Los Angeles, The University of Chicago, The University of Houston, The University of Virginia, The University of Wyoming, Stanford University, Texas A & M University, Valparaiso University, and Yale University.

[2] I. Gaines, et. al., "Use of the Fermilab Advanced Computer Program Multi Microprocessor as an On Line Trigger

Processor," IEEE Transactions on Nuclear Science, vol. NS-32, pp. 1397-1404, August 1985.

[3] H. Muller, "Online Fastbus Processor for LEP," Proceedings of the Conference on Computing in High Energy Physics, Amsterdam, Netherlands, 1985.

[4] M. A. Oothoudt, et. al., "Design of the RSX-11M Q System," IEEE Transactions on Nuclear Science, vol. NS-26, pp. 3724-3730, October 1981.

[5] Phillips Scientific, 305 Island Road, Mahvah, New Jersey 07430.

[6] IEEE Standard FASTBUS Modular High-Speed Data Acquisition and Control System. New York: IEEE, 1985, Ap. E.

[7] D. R. Machen, "The Fastbus Segment Extension," in Proceedings of the Fastbus Software Workshop, 1985, pp. 145-156.

[8] Creative Electronic Systems S.A., Case Postale 122, 1213 Petit-Lancy, Switzerland.

[9] DR. B. Struck, Hauptstrasse 95, D-2000 Tangstedt, Hamburg, Federal Republic of Germany.

[10] J. D. Blake, Use of Microprocessor Cross Software under VAX VMS, CERN Internal report, CERN/DD/SW-LH-[illegible].

[11] H. Von Eicken, MONICA a Symbolic Debugging Monitor for the M68000. (User manual in preparation), DD-Division, CERN, 1986

[12] T. Kozlowski and W. M. Foreman, "An Implementation of the New IEEE Standard Routines for Fastbus," presented at the IEEE Conference on Real-time Computer Applications in Nuclear, Particle and Plasma Physics, San Francisco, Ca., May 12-14, 1987.

[13] T. J. Berners-Lee and A. Pastore, RPC Users Manual. Geneva, Switzerland: CERN DD/OC, TBL-RPCCMAN, March 31, 1987.

[14] M. Larvill, FASTBUS To Branch Bus Controller (FBBC) Preliminary Specification. Fermilab, CDF453, Sept. 4, 1986.

[15] Omnibyte Corporation, 245 West Roosevelt Road, West Chicago, Il. 60185.

[16] M. Larvill, Fermilab, private communication.