Fermi National Accelerator Laboratory

FERMILAB-Conf-91/172

PAN-DA and BEYOND
Data Acquisition for the Next Generation Experiments

R. Pordes, J. Anderson, D. Berg, E. Berman, D. Brown, T. Dorries, B. MacKinnon, J. Meadows, C. Moore, T. Nicinski, G. Oleynik, D. Petravick, R. Rechenmacher, G. Sergey, D. Slimmer, J. Streets, M. Vittone, M. Votava, N. Wilcer, and V. White

Online Support Department

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

June 1991

* Presented at the *7th Real Time Conference*, June 25-28, 1991, Julich, Federal Republic of Germany.

PAN-DA AND BEYOND
Data Acquisition For the Next Generation Experiments

R.Pordes, J. Anderson, D. Berg, E. Berman, D. Brown, T. Dorries,
B. Mackinnon, J. Meadows, C. Moore, T. Nicinski, G. Oleynik,
D. Petravick,R. Rechenmacher, G. Sergey, D. Slimmer, J. Streets,
M. Vittone, M. Votava, N. Wilcer, V. White

Online Support Department
Fermi National Accelerator Laboratory
P.O. Box 500

Batavia, Il 60510, USA

Sponsored by DOE Contract no DE-AC02-76CH03000

Abstract

We report on the status of the PAN-DA data acquisition system presented at the last Real Time Conference. Since that time, PAN-DA has been successfuly used in the fixed target program at Fermilab.

We also report on the plans and strategies for development of a new data acquisition system for the next generation of fixed target experiments at Fermilab.

I. INTRODUCTION

The PAN-DA[1] system provides for acquisition of data from FASTBUS, CAMAC and VME, software data filtering and monitoring in VME based processor nodes, logging of data from VME to 9 track tape or 8 mm tape cassettes, and low rate ancillary distribution of data over Ethernet through VME to backend VMS and Unix systems. The system is centrally controlled from a VAX, which also provides message and statistics reporting for monitoring the divers data acquisition components .(See Figure 1).

While PAN-DA can sustain front end rates of 7 Mbytes/sec and logging rates of up to 1 Mbyte/sec, these rates fall short of the requirements of the next generation of fixed target experiments at Fermilab, which will be taking data in 1994. These experiments require front end rates of 100 Mbytes/sec, parallel filtering at a rate of 10KHz, and logging rates of 5-10 Mbytes/sec. We report on our work towards design and implementation of this new system.

II. STATUS OF PAN-DA

The full PAN-DA data acquisition system has been successfully used by one experiment, E687[2] during the 1990 fixed target run. It has since been adopted by a second experiment for the upcoming run. Components of PAN-DA are used at six other experiments.

As implemented, PAN-DA supports a FASTBUS "event builder" host crate. A Struck General Purpose Master (GPM) reads out sub-events from multiple LeCroy 1892 buffer memories and builds events within its memory. Buffers of built events are written through branch bus to Fermilab ACP single board computers[3] residing in VME. All data flows over the FASTBUS backplane twice, and the data rate from FASTBUS to the VME buffers (including the buffer handling protocol overhead) gives an overall throughput of 7 Mbytes/sec between FASTBUS and the VME memories.

The "Enhanced Event Pool" consists of a central "event flow control" processor - an embedded 68020 VME single board computer. Multiple 2-6 Mbyte single board VME 68020 processors provide for event buffering and possible software filtering. Intelligent I/O controllers - which relieve the burden of execution of the protocol of the devices from the main processor board - are used to log and distribute the data. The control processor manages the parallel logging of data and distribution of events, and provision of available buffers to the front end data system.A logging rate of up to 960 Kbytes/sec to 4 exabytes in parallel is sustained. An event distribution rate of 150-180 Kbytes/sec to the analysis computers is achieved. (see figure 2)

Since the last data taking run, E687 have added 3 Sun workstations to their Online Computing system. PAN-DA distributes events directly from the VME buffers to the Suns in parallel with distributing events to the existing VMS monitoring computers. E687 ported their complete offline reconstruction code to SunOS to have sophisticated monitoring of the physics of the experiment in real time.

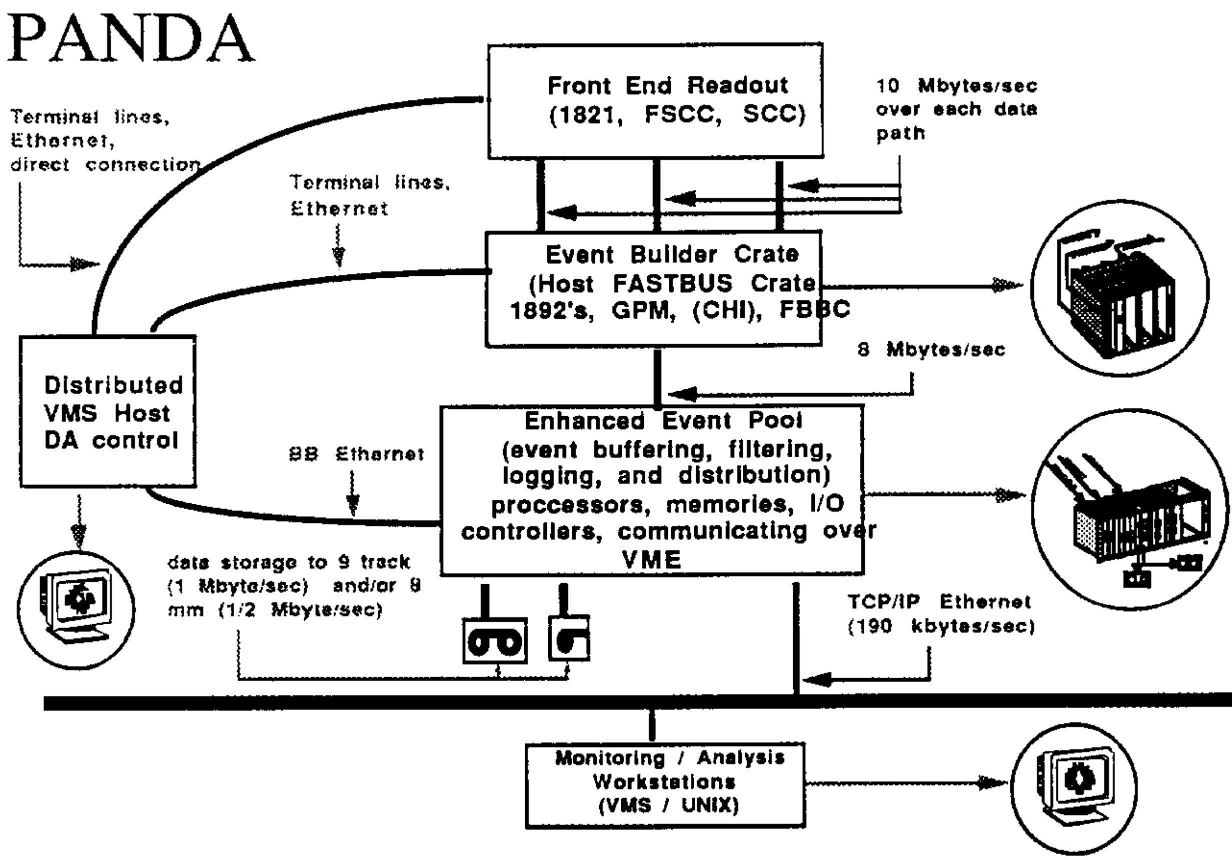

Figure 1

III. Experiences Using PAN-DA

After initial commissioning, PAN-DA has turned out to be a very stable system. E687 reported minimal problems during the 6 months run - most of them caused by bad exabyte media or drives. The deliberate effort we expended on system diagnostic and monitoring tools - both hardware and software - has proven its worth. These tools include message logs from all sub-systems and canned "snapshot" command procedures which automatically mail the status of the system to the experts.

During the initial stages of data taking the PAN-DA support team was able to diagnose problems remotely and to store information to perform post-mortem analyses offline without impacting data taking. We used VME Bus monitors, high speed logic analysers, and SCSI bus monitors to find intermittent system problems quickly and systematically.

IV. Next Generation DAQ

The next generation of fixed target experiments at Fermilab, targetted to take data in 1994, require a significant increase in the performance of the data acquisition system. We are collaborating with E781 as well as other potential customers, in specifying the architectural design, system design and implementation of a new data acquisition system.

E781's[4] requirements include a trigger rate of 10 Khz and an event size of 10 Kbyte, for an overall data rate during the spill of 100 Mbytes/sec. The events are buffered while a filtering algorithm requiring at least 10 msec/event of a 1 mip is used to provide a software filter. At present, it is expected that the software filter will require access to less than 10% of each event's data. The filter is designed to reduce the event rate by a factor of 100, resulting in a logging requirement of about 1 Mbyte/sec.

E781 therefore has interesting requirements for large buffer memories to hold the events during online filtering; large amounts of computing power in the data pipeline; and stringent requirements for event building and acquisition rates. E799/P832[5] and P831 have similar requirements for buffering, but much more stringent constraints on the front end dead time.

V. Modelling NEW_DAQ

As the first stage in understanding the requirements and constraints of these new systems, we are building a model of the possible data acquisition architectures using MODSIM[6] - an object oriented system modelling tool reported on elsewhere at this conference.[7]

Parameterizing the system will allow us to investigate the interrelationships between trigger rate, event size, dead time, software filtering time and data throughput. We are also collaborating with other groups in modelling of the data acquisition upgrade of Fermilab collider experiment (CDF)[8]. (see Figure 2).

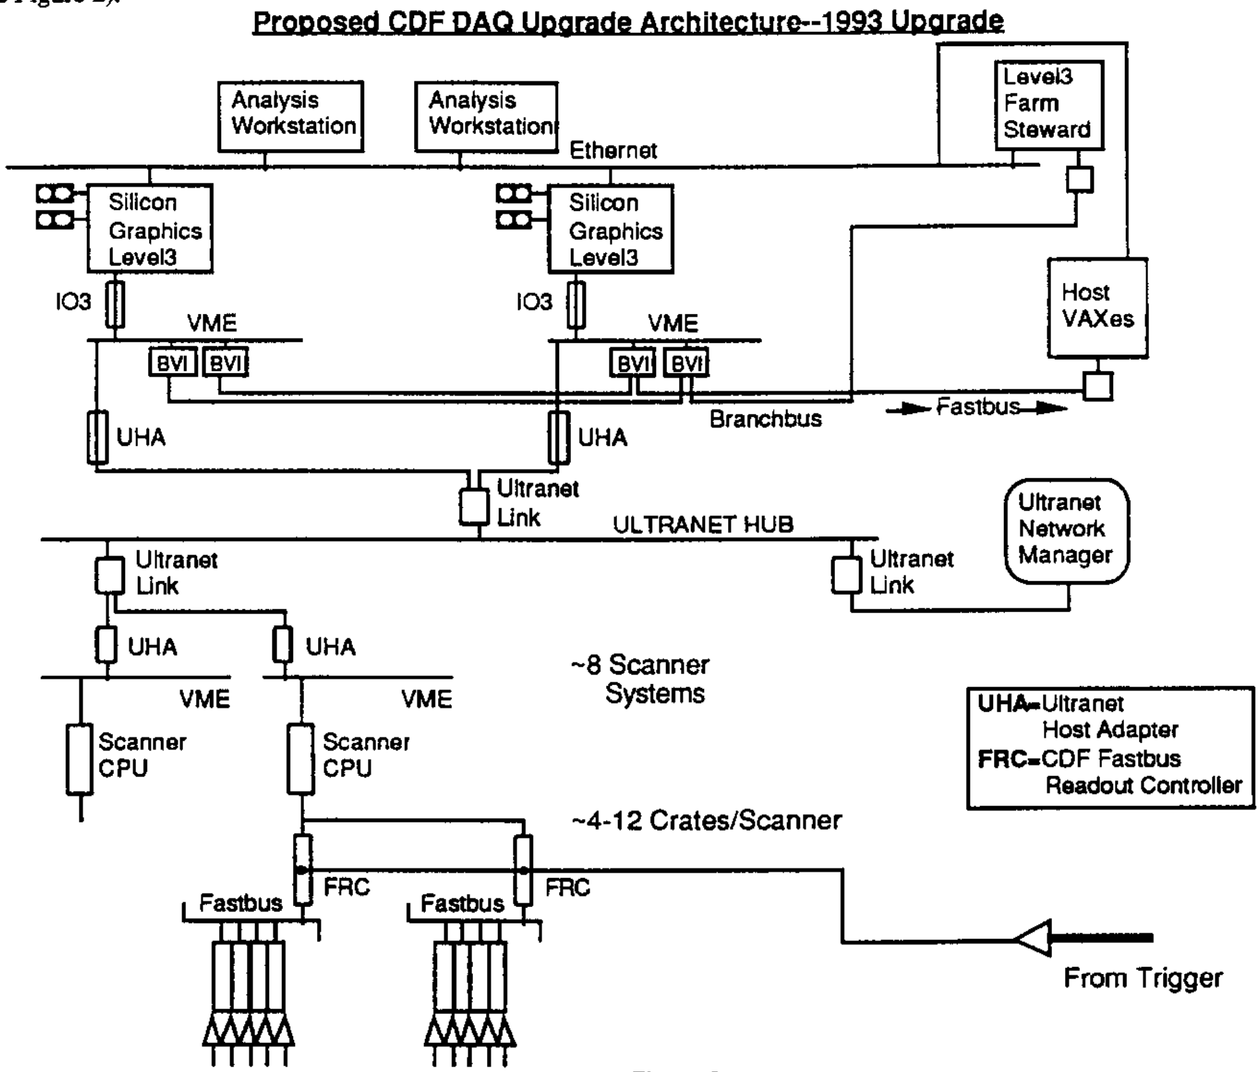

Figure 2

This will include a commercial Ultranet Switch as the event builder sub-system, being fed by eight input streams, reading out the 50 front end FASTBUS Crates. Intelligent routing of the sub-events through the switch will provide completely built events to multiple-cpu Silicon Graphics computers which perform software filtering on the event data. Event rates of up to 100 Hz for event sizes of up to 250 Kbyte are attainable. Again, software filtering will reduce the logging rate to 1 Hz. Several different architectures will be modelled and evaluated for the fixed target experiment requirements.

VI. ARCHITECTURE OF NEW_DAQ

The architecture being considered for this new Data Acquisition system consists of the following components (see Figure 3):

- parallel front end readout crates in FASTBUS, with very low deadtime readout controllers;
- very large sub-event buffers - up to several gigabytes;
- Event building across VSB or some other dedicated high speed backplane or synchronous bus into VME-based memories or cpu boards;
- RISC-based VME and/or VSB accessible processors for software filtering of the data (single board or commercial UNIX workstation computers);
- parallel logging of data from each filter processor, or across VME onto available I/O media;
- distribution of events across ethernet or fddi;
- distributed control and monitoring of the system from a distributed system of Unix hosts.

VII. HARDWARE FOR NEW_DAQ

We are evaluating the new generation of ADC and TDC modules - LeCroy modules, Phillips modules and in-house developed Fermilab modules - to meet the front end digitizing requirements.

A new FASTBUS readout controller board may well be needed. In particular, the requirements of E781 will require a long distance (optic fibre?) connection from the front end crates to the next level buffers. CDF have identified a need for a new readout controller, which is being specified and designed at this time.

Large, high speed memories were developed at Fermilab and have been used in several experiments to date. Though originally built as simple FIFO buffers, parallel control boards have been developed to allow record word counts to be stored in parallel to the data, and read out in front of the main data stream. Interfaces for the VDAS memories to VME, CAMAC and FASTBUS, which support I/O at up to 40 Mbytes/sec, have been developed at Fermilab. We are also evaluating commercial VME dual and triple ported buffer memory boards.

The computing requirements for a software filtering processor may easily be met by a number of off the shelf Unix workstations. Tests are being done to evaluate which can most effectively be used as online software filtering processors. CDF is committed to using Silicon Graphics Workstations in this role. Other options include the Sun SparcEngine, IBM R6000 etc.

We are continuing to evaluate new I/O media, although so far the promise of the double density, double speed 8mm tape drives has not materialised in any useable form.

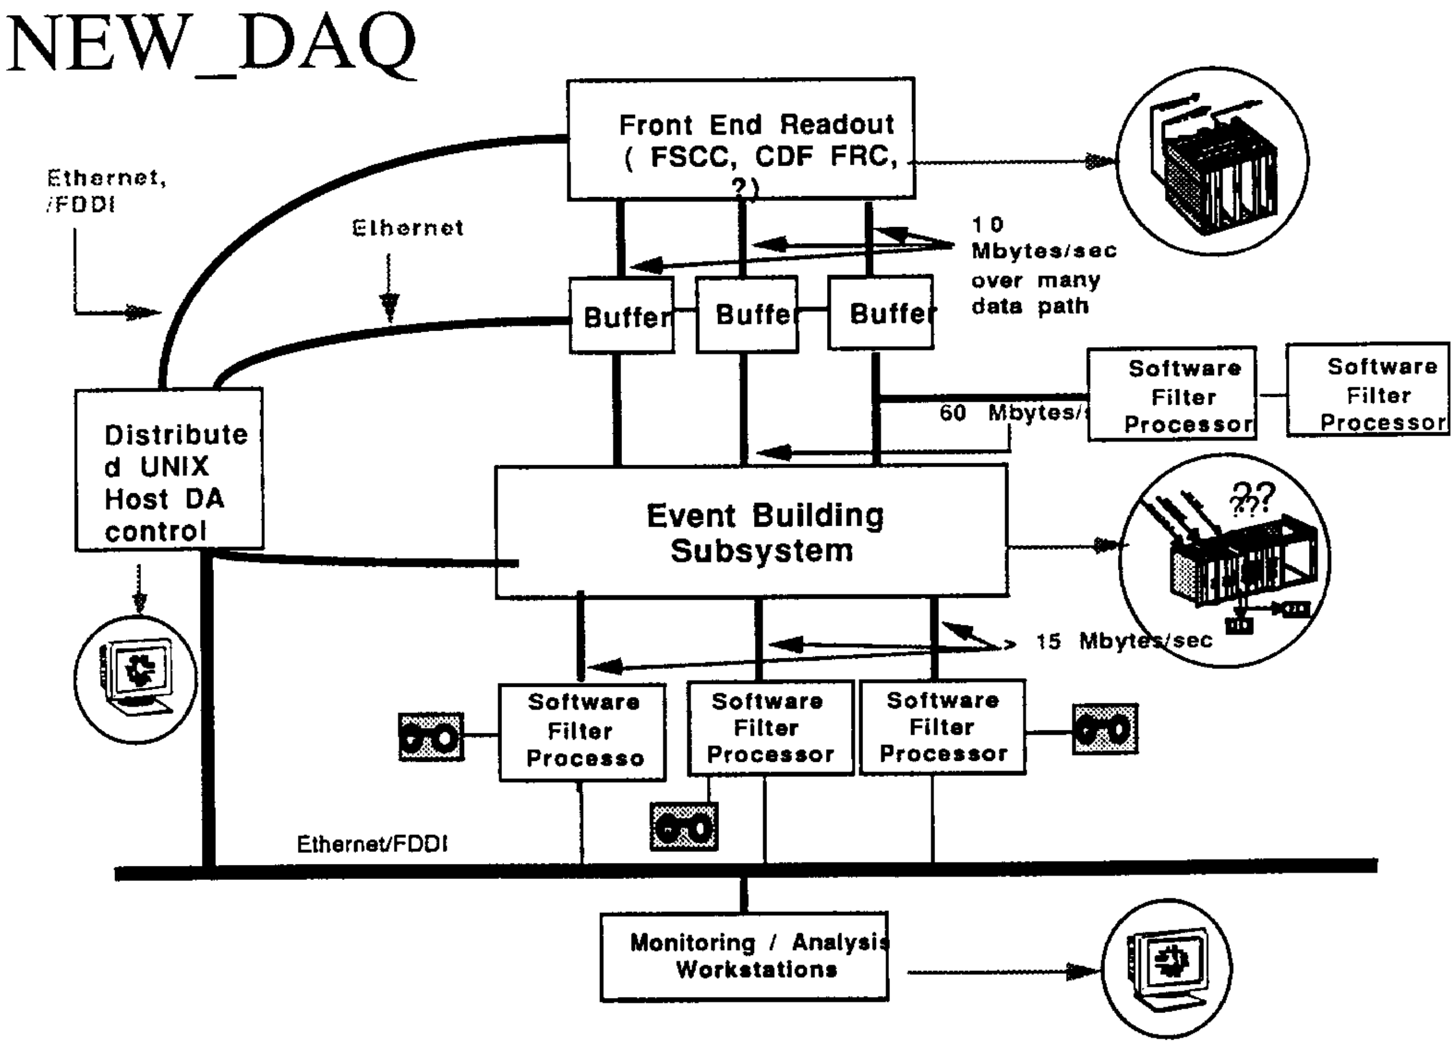

Support for a Unix based host must include consideration of the control and monitoring requirements of a distributed data acquistion system. We need to understand the throughput issues of logging devices from Risc based workstations. We are currently looking for SCSI/Data acquisition bus interfaces and as a start, evaluating the CES CBD parallel CAMAC branch highway driver. We are investigating the availability of SCSI or VME FASTBUS host interfaces.

Our strategy is to not restrict support to any particular Unix platform. All applications we develop in house are released on four platforms. We intend to use this strategy to enable us to add support for new platforms as they are adopted by experiments.

VIII. SOFTWARE FOR NEW_DAQ

We are committed to designing and implementing the system requirements and design of NEW_DAQ using online CASE tools as described in a paper at this conference.[9] We are using CADRE's TEAMWORK for all new software development projects.

We are planning to migrate from the pSOS real time kernel used in the embedded processors for PAN-DA to a kernel such as VxWorks or Lynx which have more fully developed host based software development support and networking applications. We found the lack of good program development tools to be a hindrance in the development of PAN-DA.

We are gaining experience in the use of Motif and X-windows as software support layers for graphical user interface input and information presentation paradigms. We expect to commit to these for use by the host based NEW_DAQ components. We are coding a Unix based message reporting system to gain experience with these tools.

We are developing a prototype project in C++ in order to gain the experience needed to determine the appropriateness of Object Oriented Programming to our applications.[10] Other prototype projects will be used to evaluate the Eiffel OOP Language and to investigate possible commercial distributed data bases for use in the Host Application layers of NEW_DAQ.

IX. PLANS AND SCHEDULE

Our schedule is inevitably tied to the Fermilab experimental program. Components of NEW_DAQ that may be appropriate for the Collider experiments will be needed within eighteen months. Components for the fixed target experiments can be developed over a longer time frame.

We plan to converge on the detailed architecture of our new data acquisition system by the end of 1991, to be followed by the challenge of detailed design and implementation and the excitement of commissioning the system as part of a living physics experiment.

X. REFERENCES

[1]Berg, et al, "The PAN-DA Data Acquisition System", IEEE Transactions on Nuclear Science, Vol.NS-36, No.5, October 1989.

[2]E687, Memorandum of Understanding, Fermilab internal note

[3]T.Nash et al. The ACP Multiprocessor system at Fermilab, XXIII International Conference on High Energy Physics, Berkely, Ca, 1986.

[4]E781 Memorandum of Understanding, Fermilab internal note.

[5]DA requirements for E799 Phase II, internal note, T. Yamanaka, Fermilab

[6]MODSIM - CACI Products Co

[7]Data Acquisition Studies for the SSC, E. Milner, paper in RT91

[8]CDF Data Acquisition Upgrade, October 1990, (pink book).

[9]Design of Digital Sky Survey DA and Online System, D. Petravick et al, RT91

[10]Object Oriented Design and Programming for Experiment Online Applications, G. Oleynik, CHEP 91.