

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.

DOE/ER/03130--44C

DE89 013955

Data Acquisition at D0*†

David Cutts and Jan S. Hoftun

Brown University

Providence, RI 02912, USA

Christopher R. Johnson and Raymond T. Zeller

ZRL

Cranston, RI 02905, USA

ABSTRACT

A large array of MicroVAX computers will be used for high-speed data acquisition and online filtering in the D0 experiment at Fermi National Accelerator Laboratory. This system has been described elsewhere [1-3]. We discuss new developments related to this system, in hardware and particularly in software. The new hardware includes multiported memory boards for faster and more direct input and output of event data. Most of the software development is related to the crucial role of the event filters, requiring an elaborate mechanism for control and monitoring of the FORTRAN based programs running in each of the MicroVAXes.

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.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

1. The D0-experiment

The D0 collaboration at Fermi National Accelerator Laboratory is in the process of building a large new detector to study proton-antiproton collisions at 2 TeV Center of Mass energy. The D0 detector has three major components: a central tracking section with driftchambers and transition radiation detectors, a massive Uranium-Liquid Argon calorimeter and a muon detector with drift chambers and magnetized iron outside the calorimeter. The data acquisition task in such a detector is very complex. The proton-antiproton interaction rate is 100,000 Hz with each interaction supplying on the average 250K bytes of data from the total of 100,000 readout channels.

The interesting events from a physics standpoint are only a small fraction of the total number of collisions taking place. Before complete digitization, high-speed, special purpose electronics (Level-1) will be used to bring the rate of potentially interesting events down to 200-400 per second. These events are then taken over by the Level-2 filter system described here which is targeted to pass about 1-2 per second.

2. High Level Data-aquisition System

When the Level-1 trigger has indicated that an event might be of interest, digitization of the data is started and the Level-2 system is signalled. A "Supervisor" MicroVAX decides which node (or possibly nodes) should receive the data for that event and enables that node before initiating transfer of data from the front-end electronics. The digitizing electronics in 100 VME crates loads the data into VME buffer memories (one per crate) with external

ports. The shipping of this data from the buffer memories to a set of Q-bus based MicroVAX computers is controlled by 'Data Cable Sequencers'. Each Sequencer controls one of 8 parallel data-cables, each dedicated to a part of the detector. Each data cable is connected to a new generation of Multi-Ported Memory boards (ZRL Q22MPM) in each of the Level-2 MicroVAX processors [4]. These MPM's are real MicroVAX memories (using a PMI port) with external ports for the data-cable input and possible coprocessor attachments. The coprocessors can either be attached as daughter boards or will sit on a separate board which also serves as a high speed output channel for sending the accepted events to a host system. The transfer from the VME buffer memories to the Multi-Ported Memory in the MicroVAXes is done at a rate of up to 40 MB/second per cable (aggregate rate of 320 MB/sec).

This system requires no extra event building step: the complete data for one event from the 100 VME crates is immediately available in the private memory of one (or more) of the MicroVAXes.

3. Software Aspects of Level-2 System

The filter part of the Level-2 system will remove over 99% of all the events passed to it. This means that the control and monitoring of the filter's operation is very important. It is also important to have flexibility with respect to the order in which the algorithms are applied as well as to the threshold values used in each of the various algorithms. We base the filter framework on the concept that each algorithm forms a 'TOOL', which is as much as possible independent of all other TOOLS. A single filter (many will be running at the

same time) may then be built up of a sequence of these TOOLS. Both the sequence of the TOOLS defining one filter and the parameter values used for each call to a TOOL may be changed in the running system when a new run is started. Each TOOL may take part in several filters with different sets of thresholds and other parameters for each call. Each Level-1 trigger (up to 32) may activate more than one filter sequence (for a total of 128 possible filters). The TOOL algorithms used will have to be tested thoroughly before being certified for use in the actual system. The filter code written in FORTRAN may be used in offline mode under VAX/VMS before going into the running Level-2 system. No real changes are needed to make it run under VAXELN.

3.1. COMPONENTS OF SUPPORT AND MONITORING SOFTWARE

The software which supports and monitors the system described above will have to be complex. There are a number of individual computers all needing to work together to achieve the highest possible throughput in the data aquisition system while maintaining good knowledge about the data being rejected. Some of the components are part of the commercial VAXELN Dedicated Runtime System used on the farm of MicroVAXes. These are EBUILD for assembling separate .EXE files into down-loadable files and EDEBUG for full symbolic, remote debugging of these programs. Other components were developed especially to meet the needs of the D0 implementation of this system. These programs include an editor which writes a set of control files (L2STATE), direct control of the Supervisor (SUPCON), monitor which displays an updating screen of information from each processor in the system (L2_MONITOR) , and histogram control/display (D0HIST).

3.1.1 The L2STATE program.

The L2STATE program is used to edit files used for node definitions (L2STATE.DAT), for hardware configuration of Supervisor (L2SUPER.DAT), for definitions of filter "TOOLS" in the Level-2 programs (L2TOOL.DAT) and for definitions of types of Level-2 nodes (L2TYPE.DAT). Much of this information is coupled such that for each node described as a Level-2 node in L2STATE.DAT, a number is entered for the type definition pointing to one of the entries in L2TYPE.DAT. And in each type definition in L2TYPE.DAT, a list of TOOLS pointing to L2TOOL.DAT is included. The L2STATE program also generates the routines needed to complete the main Level-2 program thereby ensuring that the correct TOOLS are included. A way of locking access to these files as well as to other files used to keep track of run numbers and other run conditions has been implemented.

3.1.2 The SUPCON program.

SUPCON is a standalone program for controlling the running in the Level-2 system as well as for monitoring the overall dataflow through this system. It uses a set of routines needed by other programs which do direct run control, in particular a program called COORDinate (the heart of the host's ONLINE system) and by other monitoring programs being planned.

3.1.3 The L2-MONITOR program:

L2-MONITOR connects to a server in one of the VAXELN nodes in the system which receives unsolicited monitor data from each of the Level-2 nodes. The rate of this monitor messages may be changed via the Supervisor.

3.1.4 The D0HIST program.

D0HIST manipulates, displays etc. a set of predefined histograms in each Level-2 node. These histograms are intended for monitoring purposes. The program uses HBOOK/HPLLOT for the booking, filling and display of the histograms. An addition to HPLLOT is included where single data points are sent from the nodes for use in a real-time updating histogram display.

3.2. INTERNALS OF A LEVEL-2 NODE

Before receiving an event, the Level-2 node is in a wait mode, such that as soon as the first data-cable is done transferring the data, the filtering process starts. The data will immediately be ready in the standard D0 memory management format, utilizing the ZEBRA program package developed at CERN. This is a FORTRAN package managing data through COMMON BLOCKS with an internal pointer structure. Thus the data is directly available to all the FORTRAN routines which have been included in the Level-2 program (TOOLs). The framework controlling the calling of these various TOOLs is organized to achieve the flexibility in calling the TOOLs in any order and with different parameter sets. For each of the possible 128 filters, the Level-2 nodes will be sent a 'script' which is basically a list of numbers of the TOOLs to call and the parameter set to use for each call. The addresses of the TOOLs are kept in a table and the number given in the script is simply a pointer into this table. Each TOOL author also has to supply a routine which is called at the beginning of a run to load possible new sets of parameters.

The same framework structure which is set up for the Level-2 program

in VAXELN will also be used to make versions of processing programs for calibration and other special purposes.

4. Conclusions

We have an architecture for a data acquisition system which has been functioning well in advance of the experiment itself. The new developments in hardware will allow for a higher overall throughput in the system. This new hardware is currently being prototyped and is expected to go into the system in late 1989.

Recognizing the critical role in the experiment of the filter process, we have undertaken an extended design and implementation of the Level-2 software framework which will allow algorithms to be easily included and ordered in the filter program in a controlled manner. This framework will also facilitate the monitoring of efficiencies and other aspects of the data flow through the Level-2 system.

5. References

- [1] D. Cutts, J. S. Hoftun, R. T. Zeller, T. Trojak, R. Van Berg, C. R. Johnson, "D0 Data Acquisition Design", Proceedings of Fourth Conference on Real-Time Computer Applications in Nuclear and Particle Physics, IEEE Transactions on Nuclear Science, Vol. NS-32, (1985).
- [2] D. Cutts, J. S. Hoftun, C. R. Johnson, R. T. Zeller, "The MicroVAX based Data Acquisition System for D0", Proceedings of Fifth Conference on Real-Time Computer Applications in Nuclear and Particle Physics, IEEE Transactions on Nuclear Science, Vol. NS-34, (1987).
- [3] D. Cutts, J. S. Hoftun, C. R. Johnson, R. T. Zeller "Data Acquisition Hardware for the D0 MicroVAX Farm", Presented at the International Conference on the Impact of Digital Microprocessors on Particle Physics, Trieste, Italy (March, 1988).
- [4] D. Cutts, J. S. Hoftun, C. R. Johnson, R. T. Zeller "A Microprocessor Farm Architecture for High Speed Data Acquisition and Analysis", Presented at the IEEE 1988 Nuclear Science Symposium, Orlando, Florida (November 1988)

* Work supported in part by the U.S. Department of Energy.

† Presented at the International Conference on Computing in High Energy Physics 89, Oxford, England, April 10-14, 1989.

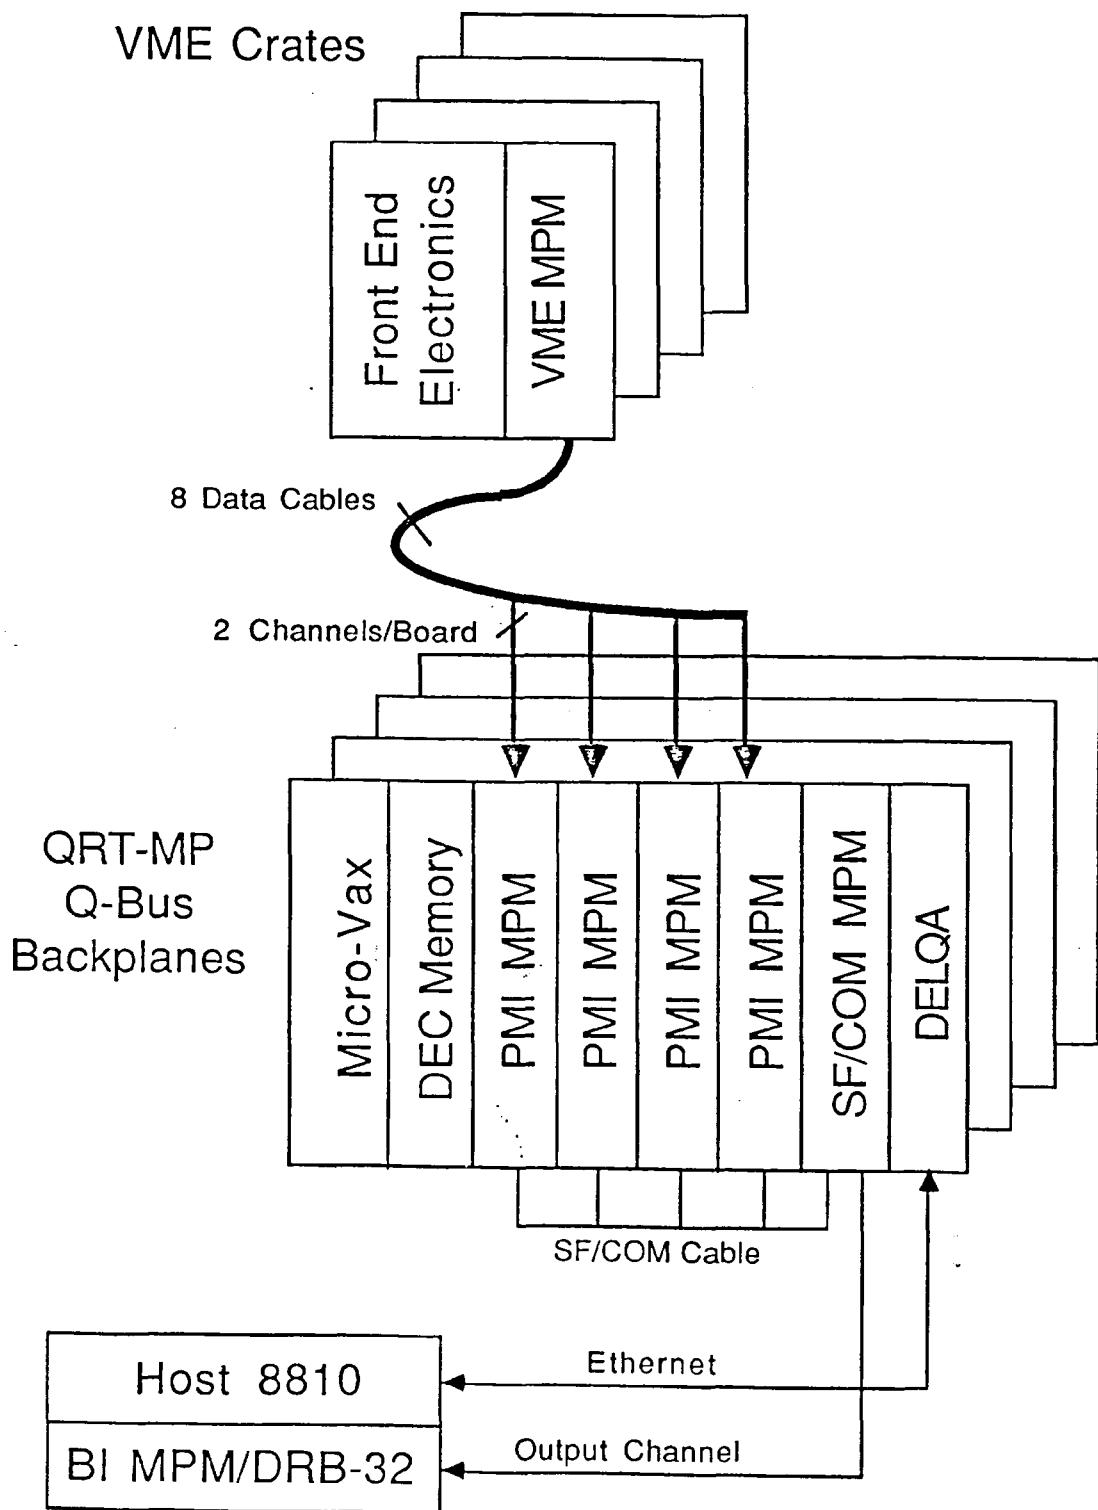


Fig. 1: D0 Data Acquisition Hardware Layout

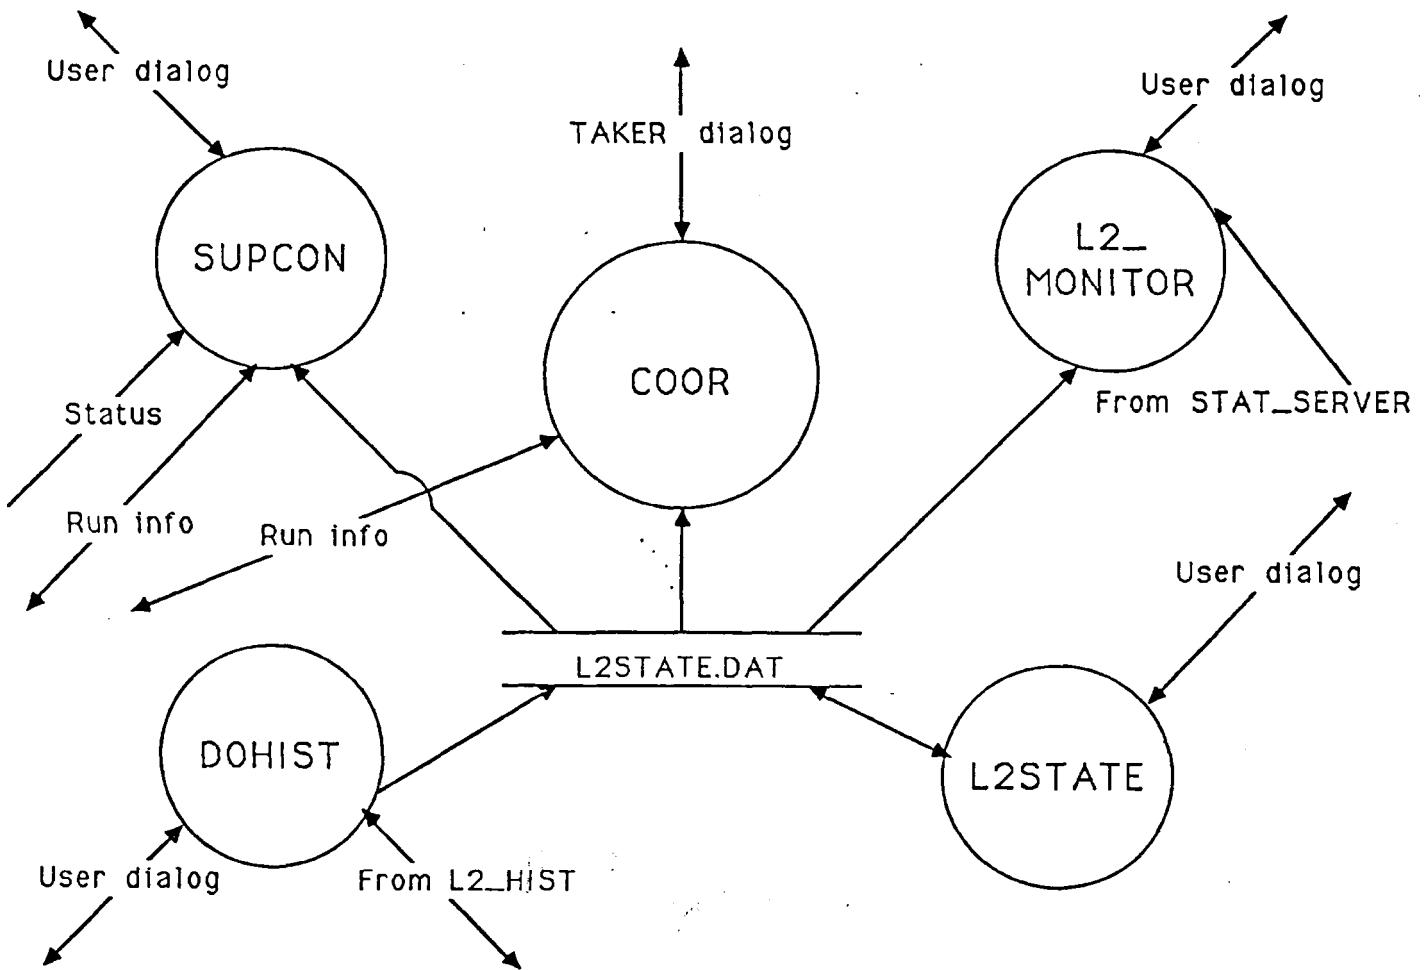


Fig. 2: Support Software for D0 Level-2 System