LEVEL 3 FILTERS AT CDF

J.T. Carroll
Fermilab
P.O. Box 500
Batavia, IL 60510

November 11-14, 1985

Summary

The Level 3 stage in CDF online filtering is currently under development. This system should support a flexible division between online and offline software filters within the constraints of the full data acquisition system. Multimicroprocessor (MMP) structures like the ACP system used by CDF could be improved with multi-rank architectures to meet SSC requirements.

Introduction

The CDF multi-level trigger system must be capable of supporting $\bar{p}p$ interaction rates on the order of 50K Hz. Reducing this to a rate of a few Hz for data logging requires sophisticated custom designed trigger hardware. While it might be possible to make the complete event selection in such hardware, there is an enormous advantage in making part of the trigger decision with software filters coded in a high level language. The software decision is inherently more flexible and allows a larger segment of the collaboration to participate in physics trigger development.

The CDF Level 3 system will use ACP[1] processors executing algorithms written in ANSI FORTRAN-77. For the 1987 run this farm or array of processors is scheduled to be fully operational and provide a minimum processing capacity of 20 VAX 11/780 equivalents. Data flow simulation studies indicate this will require about 30 nodes -- some of the processors are usually idle for input of new Level 2 triggers and output of events accepted by Level 3. Each Level 3 node will process data for a complete event, i.e. a record with data from all components of the detector. Current estimate for the mean event size is 100K bytes and each node will initially have 2M bytes of memory. The CDF Level 3 farm is designed to process events accepted by Level 2 at a mean rate of 100 Hz.

Components of the CDF FASTBUS network required to support Level 3 are shown in Fig 1. When an event is accepted by Level 1 and Level 2 triggers, Event Builder (EVB) modules collect data from front end scanners and reformat this scanner segmented data into detector component orientated YBOS[2] banks. The EVB writes a complete event into memory of an ACP processor in a VME Node Crate. These VME crates are chained together on ACP specified branch buses. A hardware switch connects a bus with processor crates to a master on another branch bus. Events accepted by Level 3 are transferred to VAX host computers for data logging and access by online consumer processes. Both the EVB and consumer VAX use a FASTBUS to Branch Bus (FBBC) slave module for I/O to the Level 3 farm. The processor farm is managed by a MicroVAX II labeled FARM MANAGER in Fig 1. This manager has Q-bus interfaces to both the ACP Branch Bus (QBBC) and FASTBUS (QPI). Farm management functions include initialization of processors and communication with the Buffer Manager (BFM) to specify free nodes and accepted events. The BFM manages data flow from the EVB, through Level 3 to VAX consumers.

Prototype Algorithm

A prototype Fast Calorimetry Filter[3] has been written as a benchmark for Level 3 applications. This algorithm processes data in YBOS detector bank format for the 4π electromagnetic and hadron calorimetry. In normal data acquisition, calorimetry banks received by Level 3 will only have channels above threshold with pedestal corrections already applied by the front end scanners. The prototype filter executes the following steps:

1. Decode each calorimetry bank and build a list with gain corrected energy, transverse energy, rapidity and azimuth angle,

2. Build a list of towers with binning of 15° in azimuth and 0.1 in rapidity,

3. Find clusters in this list of towers and build a list with azimuth, rapidity and energy for each cluster,

4. Calculate jet parameters including directional vectors, transverse energy and global shape parameters.

With Monte Carlo data for CDF central and end wall calorimetry, the average execution time for this filter was 100 ms on a VAX 11/780. Execution time was roughly proportional to event size. This prototype FORTRAN algorithm was coded with minimal subroutine calls and used lookup tables to minimize calculations. Consequently it is somewhat more optimized for execution than offline analysis code which emphasize structured programming.

Filter Strategy

Level 3 algorithms and offline programs can be viewed as a sequence of steps in the same processing chain which prepares data for physics analysis. The separation between functions performed online and offline at CDF will be flexible since Level 3 uses the same data structures and support utilities as standard offline processes. Some of the factors which define this separation between Level 3 and offline functions are as follows:

i) The combination of Level 1 and Level 2 followed by Level 3 software filters must provide a maximum rate for accepted events of a few Hz. The maximum rate for logging 100K byte events to 6250 bpi tapes is about 5 Hz. The long term rate should be about 1 Hz for reasonable offline storage and processing capacity.

ii) The input trigger rate which can be processed by Level 3 is limited by both deadtime and bandwidth considerations. Every trigger accepted by Level 2 must be digitized and buffered by front end scanners. At the design specification of 1 ms for scanner readout, deadtime will be about 14% for a Level 2 accept rate of 100 Hz. The ACP branch bus and VME processor crates have a bandwidth of 20M bytes/sec which also limits Level 3 event input to about 100 Hz. FASTBUS message traffic to manage data flow also cannot handle rates significantly higher than 100 Hz.

iii) Processing capacity available at Level 3 also limits filtering and data reduction at this stage. However, the ACP farm can be expanded to provide at least 400 VAX 11/780 equivalents so installed capacity at Level 3 is a cost effectiveness issue for CDF rather than a DAQ design limit.

iv) Perhaps the most important issue in the division between online and offline processing steps is user confidence in Level 3 hardware and software. As both online DAQ and offline processing systems mature, offline functions can move to Level 3 and improve online event selection.

The advantages of a flexible division between Level 3 and offline processing are more significant than the cost effectiveness of developing new Level 3 code just to reduce the size of the processor farm. Use of highly portable offline routines will allow Level 3 algorithms to be prepared quickly and easily modified with the latest information from offline processing and analysis. Level 3 must have a different support structure since data transfers are initiated by masters outside the farm rather than the algorithm itself. The Level 3 program can be structured to service its algorithms in a manner which makes this difference transparent.

Level 3 can include a variable number of autonomous algorithms which analyze different components of the detector and calculate event attributes required for the selection process. For example, calorimetry analysis could provide a list of jets and a muon algorithm could improve direction measurements using central tracking and/or vertex detector data. Level 3 event selection can then be organized as a series of filters with each filter defined as a set of cuts on event attributes. This strategy allows each filter to apply independent cuts, e.g. cuts on jet parameters may be different for events with and without a central muon. Each filter can also flag events for offline express processing. Event attributes used in this selection process can be logged with each event in YBOS banks created by Level 3.

Implications for SSC

Models for SSC data acquisition typically have about 750K channels and a Level 3 system capable of supporting an input trigger rate of 1000 Hz. With such factors 10 times CDF values, the final rate for logging accepted events is still constrained to a few Hz. Highly portable software providing a flexible division between online and offline processing should be equally advantageous for SSC experiments. In addition SSC may also need a flexible partition between online and offline computer hardware.

Under the current strategy every Level 3 node at CDF executes the same algorithms. The more complex events and higher trigger rates at SSC could be processed more efficiently in a multi-rank architecture. For example, Level 3 at SSC could be organized as follows:

Rank 1: Input rate = 1000 Hz
Accept rate = 100 Hz
Execution time = 1 sec/event - VAX 11/780 time
Calorimetry analysis with energy cuts using full detector segmentation.

Rank 2: Accept rate = 10 Hz
Execution time = 10 sec/event
More global detector filters with improved jet and lepton triggers using some tracking data.

Rank 3: Accept rate = 2 Hz
Execution time = 100 sec/event
Full reconstruction of cones around jets and leptons.

Each rank in this hypothetical example would require 1000 VAX equivalents but there would probably be significantly less memory at Rank 1. Processing at Rank 3 could reduce the standard 1M byte event size by discarding tracking data not associated with accepted interactions.

References

1. T. Nash et al., "The Fermilab Advanced Computer Program Multi-Microprocessor Project", Proceedings, Computing In High Energy Physics, Amsterdam (Netherlands), 1985.

2. D. Quarrie, "YBOS Programmers Reference Manual", CDF Note 156, 1984.

3. G.B. Chadwick, L.J. Moss and D.M. Ritson, "Current Status of Proposed Level III Software", CDF Note 246, 1984.

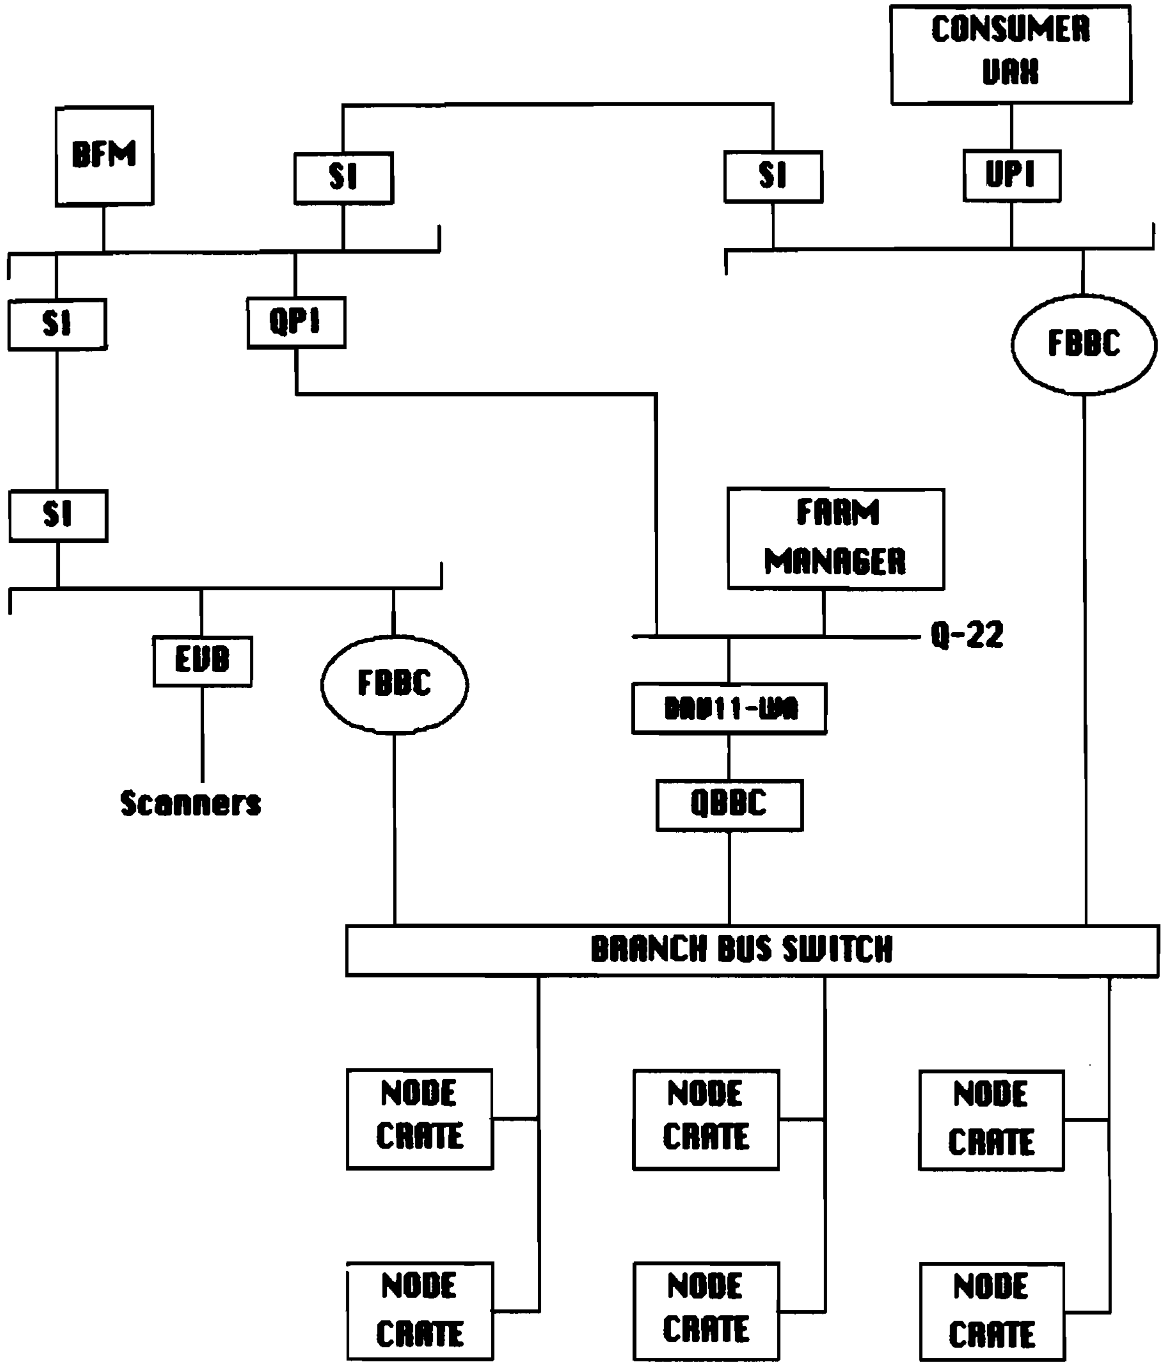

Fig. 1: CDF LEVEL 3 STRUCTURE