TRIGGER PROCESSORS FOR HADRON COLLIDERS*

I. Gaines
Fermi National Accelerator Laboratory
Batavia, IL 60510

Summary

The Collider Detector Facility (CDF) collaboration is designing and constructing a powerful, general purpose detector system for use at the Fermilab 2 TeV center of mass energy antiproton-proton collider. The detector will have approximately 75,000 channels of electronics and must be able to deal with a raw event rate of 50 kHz, corresponding to a luminosity of 10^{30}. The multi-level trigger processing system to be used in this detector is described, with emphasis on the general features of detectors at hadronic colliders which have imposed certain architectural choices on the CDF triggering and data acquisition system.

Introduction

A number of considerations lead one to consider different types of triggering schemes for experiments at hadron colliders than for those at electron-positron colliders. In particular, the high event rates and extreme complexity of the events requires powerful, yet flexible triggers that can make relatively high level physics decisions, are easily programmable so they can be modified and checked, and do not contribute significant amounts of dead time to the data acquisition process.

Unfortunately, these three requirements are somewhat contradictory. Speed must often be traded for flexibility and ease of programming. A well-known solution to this problem is to provide a series of increasingly complex triggers, each successive level of which makes more detailed decisions (in correspondingly larger amounts of time) on fewer events (due to the rejections by earlier levels of the trigger).

This is the solution which has been adopted for the Collider Detector Facility (CDF) at Fermilab. This detector will be used at the Fermilab 2 TeV center of mass energy antiproton-proton collider. The detector design attempts to provide full coverage over the 4π solid angle around the interaction region for particle tracking, fine-grained electromagnetic and hadronic calorimetry, and muon identification using a variety of different detectors. Magnetic analysis is provided for tracks in the central region by a large superconducting solenoid, and for muons in the antiproton direction using iron toroidal magnets. An isometric view of the detector is shown in Fig. 1.

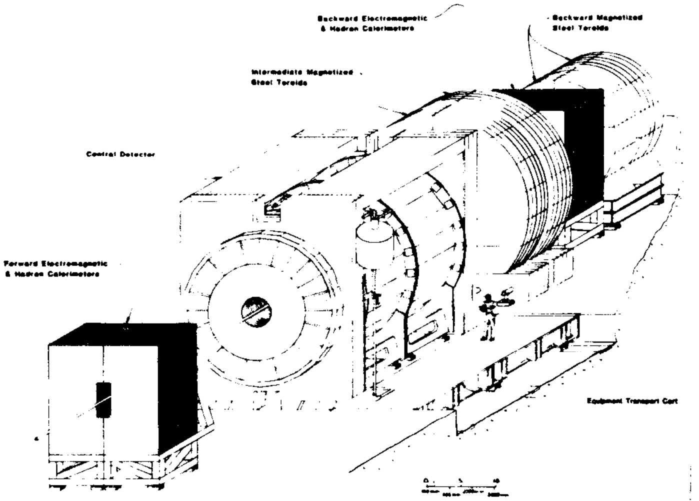

Fig. 1. An isometric view of the CDF detector.

In total, there will be approximately 75,000 individual signal sources including drift chambers, photomultipliers, cathode strip chambers, and cathode pad chambers. Further information regarding the detector can be found in the CDF Design Report.[1]

The data acquisition system is described in Ref. 2. It is a multi-function, distributed intelligence, measurement and control system which provides a variety of services in addition to data gathering. In general terms, the system consists of signal conditioning and digitizing front end electronics located on the detector which are controlled remotely by a FASTBUS based network of processors. The precise configuration is still under design and will continue to evolve as experience is gained at Fermilab and elsewhere.

Other features of the CDF triggering system besides it's multi-level character are also dictated by general considerations. The first of these is the relatively long time (expected to be at least 3.5 microseconds) between beam crossings at a proton-antiproton collider. These mean that the lowest level of the trigger will contribute no dead time as long as it makes its trigger decision faster than this interval, and that there is nothing to be gained by making the first level decisions any faster than that. This allows a certain simplicity of design and means that there is no need for ultra high speed elements in the triggering system.

Finally, the overall cost of the system is another important consideration. With the large number of channels needed in a general purpose detector, the cost per channel must be kept as low as possible. This means that the majority of channels will have no provision for fast read-out, as this would add substantially to the cost per channel. A relatively small number of channels, possibly 10 percent of the total, will be equipped with special fast read-out electronics independent of the standard data acquisition read-out path. The first levels of the trigger will make use only of these fast read-outs. Only after the event has passed the first few levels of triggering will the slow process of reading out the entire event be initiated, and the full event data will then be available at the highest levels of the trigger.

Thus, general considerations of the triggering needs for a general purpose detector for a hadron collider set the basic structure of the trigger system for CDF. There will be a multi-level trigger, with the early stages examining a portion of the event data using a special read-out path, but with no requirement for any decisions faster than a few microseconds. Higher levels of the trigger will examine the entire event, and should allow complex physics related algorithms to be used for final event selection. The details of the CDF triggering system will be described in the remainder of this paper.

---

*Operated by the Universities Research Association under contract with the United States Department of Energy.

Triggering Overview

At the design luminosity of $10^{30}/cm^2/sec$, the inelastic interaction rate is approximately 50 kHz which must be reduced to the tape writing rate by the trigger system. The rate for writing events to magnetic tape is constrained to about 5 Hz by two independent considerations. First, 5 events per second is close to the maximum rate at which a standard 6250 bpi tape drive can be operated. Second, data written at this rate for one month is estimated to require at least one year of available off-line analysis capability. A three level hierarchical trigger strategy has been chosen in which each level produces a rate low enough so that the dead time introduced by the next level is not significant. Within this constraint, the trigger requirements at each level are as loose as possible, leaving more restrictive decisions to higher levels where more information from the detector is available, and a longer decision time per event is allowed.

The first two levels of the trigger system will be used to reduce the rate from 50 kHz to about 500 Hz before digitization. Prompt signals from the detector for these two levels of the trigger structure are provided by the front end electronics in the form of drift chamber hit bits and analog sums of calorimeter towers. The level 1 trigger decision occurs in the time between beam crossings and so is deadtimeless. If a candidate event is flagged by level 1, prompt signals are passed on to level 2 for a more complex and time consuming selection process, incurring dead time. Level 2 is estimated to require of order 20 microseconds to make its decision. Limiting the level 1 rate to 5 kHz then gives an acceptable dead time due to level 2 of 10 percent. An event accepted by level 2 is digitized and stored in buffer memory. Since the digitization process takes approximately 1 msec, the level 2 trigger rate is limited by deadtime considerations to about 500 Hz.

Level 3 of the trigger structure is used to reduce the event rate from 500 Hz to the tape writing rate of 5 Hz. This level's decision criteria should be easily modifiable to accommodate changing physics requirements and increasing knowledge of both the trigger and detector operation. Accordingly, level 3 is configured as a set of independent processors which work on the entire event record using an event analysis and selection program written in a high level language. Those events which pass the filter criteria are sent to a data logger to be written on magnetic tape at an average rate of 5 events/second.

The bandwidth and processing power requirements for the level 3 system are formidable. Assuming 10 percent to 20 percent detector element occupancy and full data compaction, an event is expected to consist of approximately 10,000 32-bit words. This requires a bandwidth of up to $5*10^6$ words per second at the input to level 3. Then, assuming that about 10^5 machine instructions are needed to process an event on the average, level 3 must achieve the equivalent of $5*10^7$ instructions per second in performance.

The CDF data acquisition system is then composed of two cooperating and concurrent subprocesses. The first, or triggering and digitization subprocess, involves the lower system levels, including the front end electronics, scanners, local processors, and level 1 and level 2 triggers. The second, or event selection subprocess, employs the level 3 trigger processors to select a subset of the digitized events for logging to magnetic tape and/or transferral to the host computer. The data flow for an event in this system is shown on Fig. 2, while a schematic diagram of the system components and interconnection is shown in Fig. 3. Further details on the CDF data acquisition system besides the triggering aspects discussed here can be found in Ref. 2, 4, and 5.

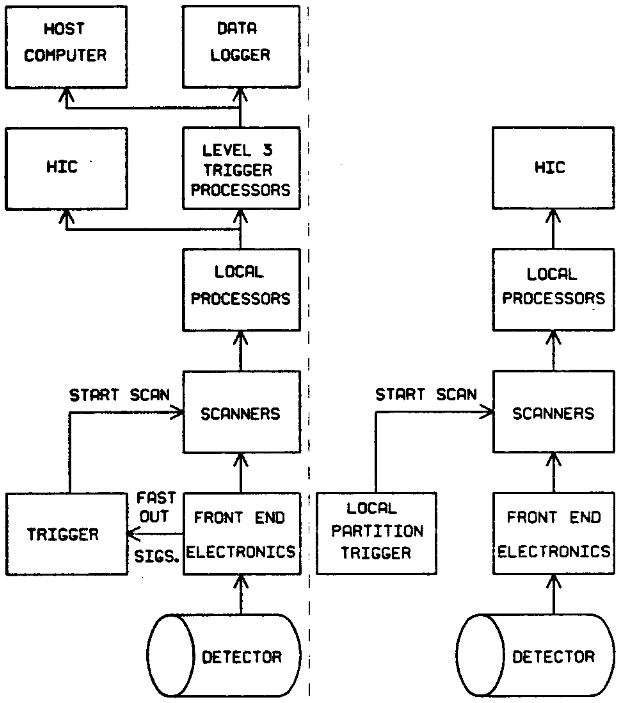

Fig. 2. a) Data flow in the global partition, illustrating the process of normal data acquisition.
b) Data flow in a local partition, illustrating the use of an independent subsystem for running diagnostics and doing calibration. HIC stands for "Human Interface Computer." There are several of these minicomputers distributed throughout the system.

Level 1 and 2 Triggers

These triggers respond to analog signals and drift chamber hit bits delivered directly to the trigger logic from the front end electronics over 5,000-10,000 dedicated cables. These signals include pulse height information from sums of calorimeter modules, timing signals from muon drift chambers, hit bit latches from tracking chambers, and current division pulse heights from tracking chambers. The trigger logic will be located outside the shielding wall. The trigger cables represent the majority of the cables for the experiment that must penetrate the shielding wall, and these cables must either be disconnected or manipulated in some manner when the detector is moved in and out of the interaction region.

The level 1 trigger makes its decision in the time between beam crossings (roughly 3.5 microseconds) so as to generate no dead time. This is expected to provide enough time to allow the level 1 trigger to identify all inelastic events with a transverse energy greater than a predetermined minimum, with more than a given number of calorimeter cells having a transverse momentum deposit greater than a preset value, and to identify events with muon candidates in either the central muon drift chambers or the forward toroids.

Beyond that, the level 1 trigger should introduce as small a bias as possible into the event sample. Up to a luminosity of 10^{29}, the level 1 trigger could in

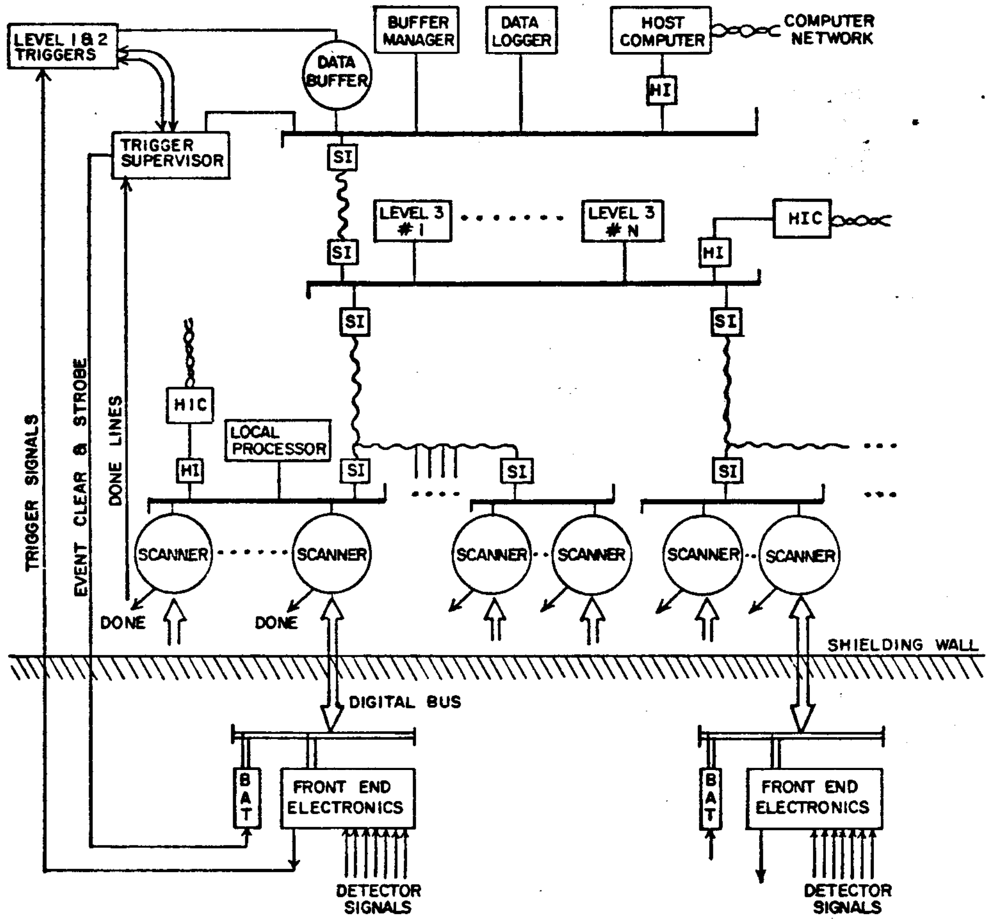

Fig. 3. Data acquisition system block diagram. SI's are FASTBUS Interconnects and HI's are FASTBUS Host Interfaces.

fact accept every inelastic interaction without introducing significant dead times at later levels of the trigger system. At higher luminosities, level 1 will be expected to select roughly 10 percent of the total inelastic interaction rate, or about 5,000 events/second.

A schematic outline of the level 1 logic is shown in Fig. 4. This fairly straightforward logic should not require the development of any special processors, but can be implemented using conventional electronics.

The level 2 trigger makes a more sophisticated decision based on the same data as that available to level 1. It selects events according to the general topology of the event, including energy clusters in the electromagnetic and hadronic calorimeters and muons in both the central and forward muon detectors.

The speed requirements on level 2 are that it not introduce large amounts of dead time when processing as many as 5,000 events per second, and thus the level 2 decision process can average no more than 20 microseconds per event. However, since the trigger decision can be asynchronous and analog information is preserved for a few milliseconds on sample and hold circuits, the level 2 processors can take up to several hundred microseconds for a subset of events provided that the majority of events are rejected in less than 10 microseconds. The level 2 processors must reduce the rate by at least another factor of 10, down to roughly 500 events per second. Events passing the level 2 selection criteria are then digitized by the normal readout process (a

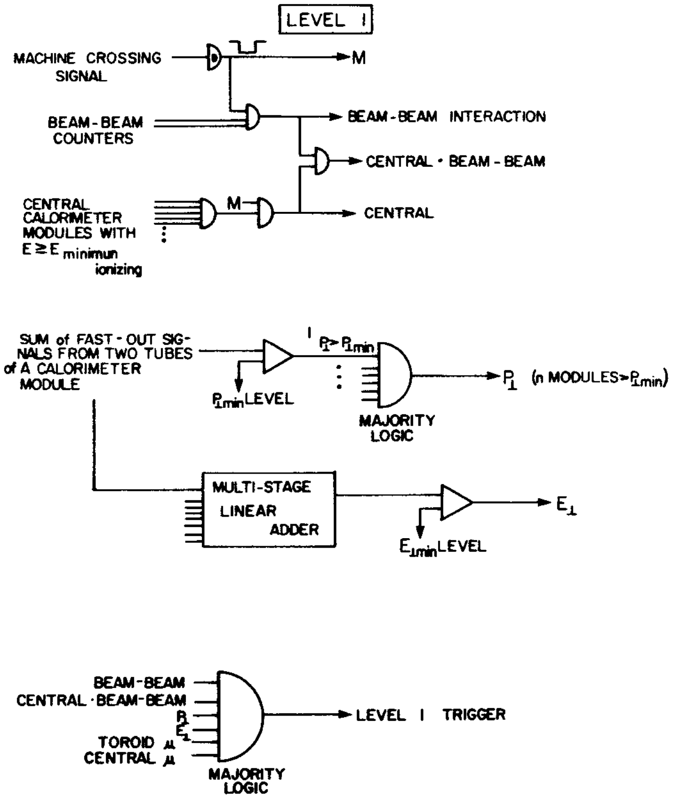

Fig. 4. Schematic of the level 1 trigger logic.

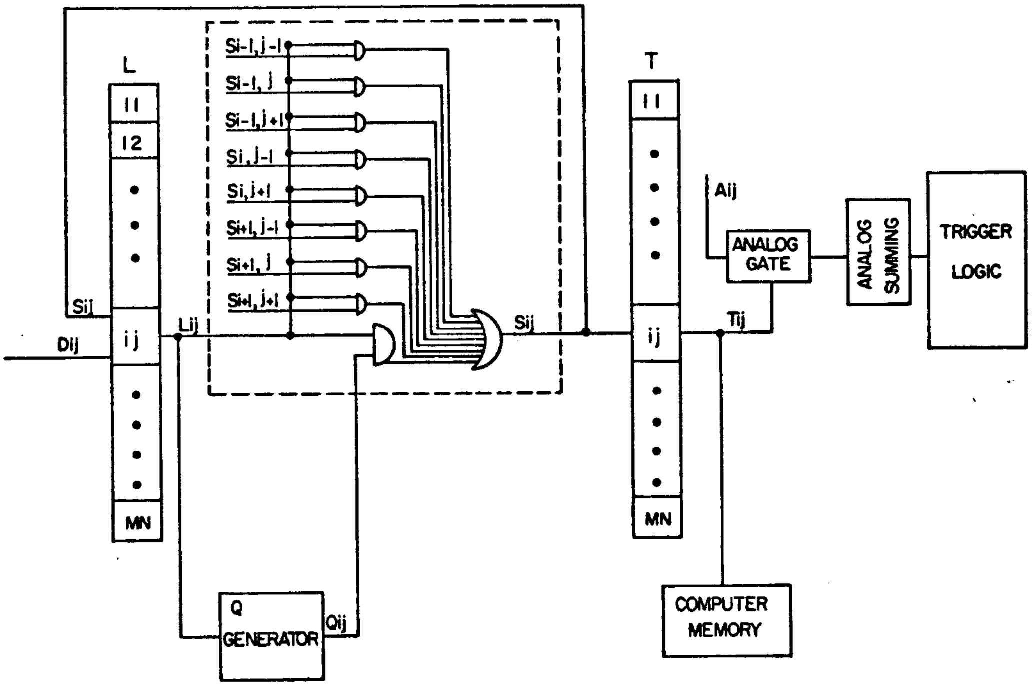

Fig. 5. The cluster finder.

slow process, requiring up to 1 millisecond per event), and the entire event will then be available for consideration by the level 3 triggering system.

The level 2 trigger processors will consist of two major portions: a collection of processing engines and a series of decision engines. The processing engines have the job of doing the actual processing of the input data and producing lists of muon tracks, energy clusters in the calorimeters, and central tracking candidates. These will be primarily hard wired modules, programmable only by reloading memory look-up tables or changing a programmed logic array. They need to work relatively rapidly to allow quick rejection (less than 10 microseconds) for most events, but do not need to be particularly flexible, as the types of calculations needed can be well predicted in advance. New types of triggers (for instance, selecting particle types with Cerenkov counters) will require additions of new detectors and new triggering cables as well as new level 2 processing engines. The initial complement of processing engines should be sufficient for almost all types of triggers which could be formed from the information initially available for level 2.

An example of a processing engine is the cluster finder shown in Fig. 5 and described more fully in Ref. 6. This device will use analog pulse heights and find a group of neighboring calorimeter modules all of which have transverse momentum deposits above a minimum threshold. It will prepare lists of the size, location, total transverse momentum, and electromagnetic or hadronic nature of each cluster.

The decision engines, on the other hand, can be somewhat slower and thus more flexible. They will be programmable devices, possibly using bit-sliced microprocessors, which will run simple programs using the lists of tracks, jets, muons, and electromagnetic showers prepared by the processing engines. A large number of triggering criteria based on the overall event topology are then possible, allowing the trigger requirements to be easily modified as the physics interests of the experiment evolve.

The overall timing of the level 1 and 2 trigger processors is shown in Fig. 6. This diagram shows how the gate and clear process is suspended by a level 1 accept decision, resuming after either a level 2 reject decision, or after the full event readout following a level 2 accept decision.

Level 3 Trigger Processors

The level 3 trigger processors have the task of performing the final event selection, reducing the event rate from 500/second down to about 5/second. The event has already been digitized through the standard read-out path, and so the level 3 processors will have the entire event at full precision to examine.

Despite the very large total processing required at level 3, as described above, there is no particular speed required for any individual event. This is due to the fact that the events are buffered, and thus level 3 processing produces no dead time regardless of how long a single event takes to process, provided that the total amount of level 3 processing can handle the total event rate.

The premium at level 3 is therefore not on processing capability of an individual CPU, but rather on total processing capability per dollar. The most cost-effective way of providing the large total amount of level 3 processing power is with a large number of small CPU's, each of which will process one event for a relatively long time. The overall processing demands are satisifed by many CPU's processing many events in parallel.

However, the individual level 3 CPU's cannot get too small. Aside from the requirement that a single processor be able to handle an entire event, the processors must be programmable in high-level languages. The level 3 trigger selection criteria will likely involve complex physics calculations including extensive pattern recognition and reconstruction of both tracking and calorimetric data, and such programs can be conveniently written only in high level languages.

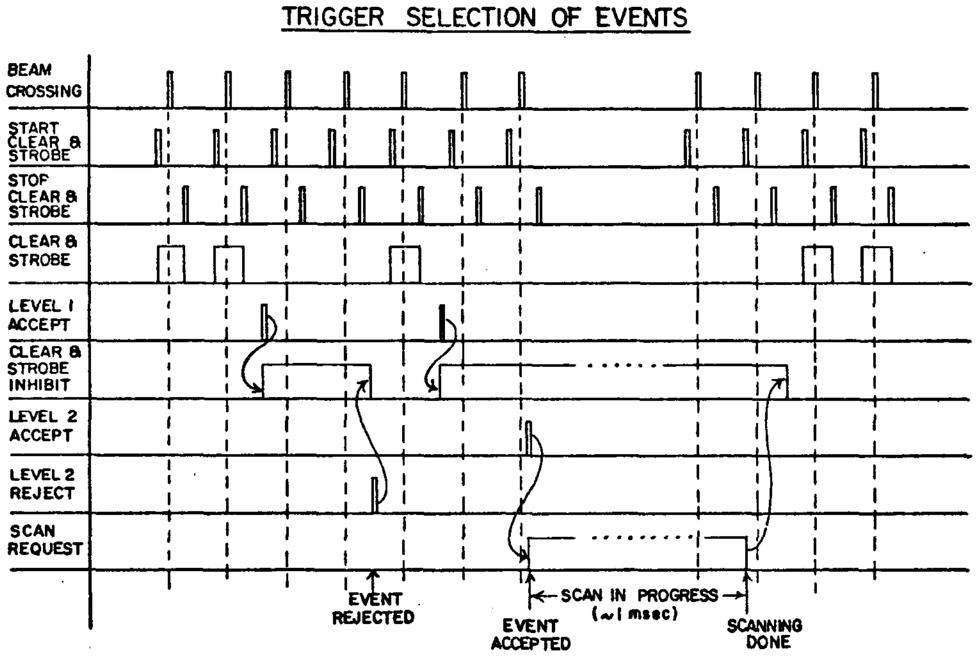

Fig. 6. Timing diagram showing the interaction between the level 1 and level 2 triggers and the trigger supervisor.

Moreover, it will aid program development and debugging if the level 3 processors can execute the instruction set of some larger CPU, which can then be used to develop and test the programs to be used for event selection.

Thus, the preferred implementation for the level 3 processors is a CPU that executes the instruction set of some popular main frame computer together with a large amount of memory, all built on a single FASTBUS card. It is hoped that such devices will be commercially available by the time they are needed in the CDF detector; if not, we will need to develop such processors ourselves.

References

1. Design Report for the Collider Detector Facility (CDF), D. Ayres et al., Fermilab internal document, CDF 111.

2. Design of the Collider Detector Facility Data Acquisition System, A. E. Brenner et al., Fermilab internal document, CDF 108 (1981).

3. FASTBUS Modular High Speed Data Acquisition System for High Energy Physics and Other Applications, U.S. NIM Committee.

4. The Fermilab Collider Detector Facility Data Acquisition System, A. E. Brenner et al., Submitted to IEEE Nuclear Science Symposium, 1981 (also Fermilab publication FN-350).

5. RABBIT System Specifications, T. F. Droege, K. J. Turner, T. K. Ohska, Fermilab internal document, PIN 52.

6. A Fast Cluster Finder, J. Freeman, Submitted to the IEEE Nuclear Science Symposium, 1981.