



Fermi National Accelerator Laboratory

FERMILAB-Conf-81/39-EXP
2100.000

A REVIEW OF PROGRAMMABLE SYSTEMS ASSOCIATED
WITH FERMILAB EXPERIMENTS*

Thomas Nash

May 1981

*Invited talk at the Topical Conf. on the Application of
Microprocessors to High-Energy Physics Experiments, CERN,
Geneva, Switzerland, May 4-6, 1981.



Operated by Universities Research Association Inc. under contract with the Energy Research and Development Administration

A Review of Programmable Systems Associated with Fermilab Experiments

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

In this paper we review the design and application of programmable systems for Fermilab experiments. The high luminosity fixed target environment at Fermilab has been a very fertile ground for the development of sophisticated, powerful triggering systems. A few of these are integrated systems designed to be flexible and to have broad application. Many are dedicated triggers taking advantage of large scale integrated circuits to focus on the specific needs of one experiment. In addition, the data acquisition requirements of large detectors, existing and planned, are being met with programmable systems to process the data. Offline reconstruction of data places a very heavy load on large general purpose computers. This offers a potentially very fruitful area for new developments involving programmable dedicated systems. Some of the present thinking at Fermilab regarding offline reconstruction processors will be described.

Introduction

The use of programmable special purpose digital systems is quietly changing the way we do high energy physics. Experimental triggers, the high speed decisions on whether or not to record an event, can now be made so sophisticated that small cross sections can be selected with confidence in a high luminosity environment. Very large detector systems with high counts of individual detector channels can be managed using microprocessor based systems that process the raw data, then buffer and pass it on to a host online computer to be recorded on tape.

General purpose computers are still used, where they traditionally have been, for data recording and the subsequent data analysis. To date the use of dedicated programmable devices has been generally in areas (triggering and raw data processing) inappropriate for general purpose computers. As I will note in examples later, dedicated systems can be made to handle specific tasks up to tens and even hundreds of thousand times faster than a general purpose computer. It is thus no surprise that there is already in existence at least one special purpose system, a CPU emulator, in an area as sacred and traditional an application of large general purpose computers as the reconstruction of tracks in multiparticle

spectrometers. In fact, planning and design of still more powerful, second generation, reconstruction processors has started.

My charge in this review is to cover the use of programmable systems associated with Fermilab experiments. At first it appeared to me that this subject would be contained in a discussion of two versatile large scale triggering systems that have been developed at Fermilab, the M7 and the ECL-CAMAC system. However, a check of the experiments on the agenda at Fermilab has shown that there is, in addition, an impressive number of often modestly scaled electronic systems with original and creative designs catered to the needs of a diverse array of individual experiments. This creative activity is to large degree encouraged by the Fermilab fixed target environment. Sophisticated triggering is strongly beneficial to many Fermilab experiments, which need to focus on a small cross section in a high luminosity interaction. At the same time it is made feasible by the excellent duty factor of the Fermilab accelerator.

Because of the pressure of work at active experiments there is very little that has been written down by the designers of most of these systems. Gathering information for this review turned out to be much of a journalistic endeavour, involving interviews carried out in electronics rooms all over Fermilab. As a result of this rather unusual method of obtaining information for a technical article, some inaccuracies, misunderstandings and incompleteness may have resulted, for which I apologize in advance.

Triggering Systems

The concept of multi level triggering has been essential to the recent development of sophisticated triggers. The problem of trigger decision making is bounded on the one end by the rate of detected beam interactions which, varying with experiment, can be as high as 5×10^8 /sec. At the other are the data recording limitations of an online computer. The latter, depending on event complexity, rarely exceeds 100-200/sec. The dead time introduced by a decision making process is simply $n_i t_i$ where n_i is the rate at which decisions are required and t_i is the mean dead time to make one decision. A total dead time of 10% is typically allocated to trigger decisions in modern experiments. (As much as another 50% may be required by the rate limitations of the data recording computer). Thus decisions running at the rate of detected beam interactions (10^6 /sec for example) must be made rapidly (100 nsec to stay below 10%). Such low level trigger decisions, even using the capabilities of the fastest electronics, are usually not made very sophisticated. They can, however, be used to select events for more complex, high level, decisions that take, say 10 μ sec and can thus run at $\leq 10^4$ /sec. In this way, multiple steps of event filtering

Table 1: Comparative Table of Programmable Trigger Systems at Fermilab

#	Experiment Physics	Decision ²	Process Time	Rate In	Rate Reduction	σ Out	Process	Experts/ Designers	Cost: 3 Parts and Assembly	Status	First Use for Data
272	Primakoff K^*, ρ	L: Trk multiplicity	60 ns	5×10^5	10^4	7 mb	Bellcord, ECL	Kerns	\$ 5K	run	Summer 78
401	$\Upsilon_p \rightarrow \psi$	L: Scint. hod. $\mu\mu$ or ee, multiplicity	300 ns	$> 10^5$	500	1 μb	DC Register Logic	Sippach	\$50K	run	Fall 79
400	$n\bar{p} \rightarrow \psi, \text{charm}$	H: 2 clean μ trks H: $M_{2\mu}; K_{\pm}^{\pm}$	50 μs 75 μs	300 3500	2 4	0.5 μb 0.5 μb	M7 M7	Droege, Gaines, Harding	\$20K	construction	Winter 81
326	$\pi^+ A \rightarrow \mu K$	L: trks point to tgt H: $d\mu$ trk patterns	20 ns 12 μs	5×10^8 200	2×10^6 20	50 nb 2 nb	ECL gates, RAMS, etc.	Summer, Isalla	\$25K	run	Winter 80
516	$\Upsilon_p \rightarrow (\text{charm})p$	L: hadronic interact. H: single p recoil, M_X	100 ns 10 μs	10^6 2500	400 35	250 μb 7 μb	NIM analog sums ECL-CAMAC Processor	Barsotti, Bracker, Martin, Nash	\$75K	running	Spring 80
594	$\nu_{\mu} e \rightarrow \nu_{\mu} e$	L: ν interaction H: cc veto fat shower veto	50 ns 50 ns 600 ns	5000 1500	3 3	10^{-37}cm^2 $4 \times 10^{-38} \text{cm}^2$ 10^{-38}cm^2	NIM NIM ECL 10000	Mattison, Kerns Droege	\$10K	Spring 81 tested	?
663	$h^+ p \rightarrow \Lambda(\bar{h})$ polarization	L: ϕ_{π} in spectr. H: track patterns	50 ns 80 μs	10^6 400	2000 5	80 μb 20 μb	NIM TTL-MWPC shift reg. readout	Kobrak, Courlay, Melanson	\$10K	testing	Spring 81
612	$\Upsilon_p \rightarrow X\bar{p}$ Low l_t	L: tag and recoil H: TREAD track	50 ns 200 μs	5×10^4 500	100 10	5 mb 500 μb	NIM ECL 10000	White	\$10K	running	Spring 81
537	$h^+ A \rightarrow \mu\mu$	L: Scint. hod. $\mu\mu$ H: ≥ 2 clean DC μ , $M_{\mu\mu}$	100 ns 5 μs	5×10^6 500	10^4 10	50 μb 5 μb	NIM ECL-CAMAC Processor	Areti, Conetti	\$75K	running	Spring 81
623	$\tau^+ p \rightarrow \phi\phi$	L: $2K^+ 2K^-$ H: $M_{\phi} \times 2, P_{\phi} \times 2$	75 ns 75 ns	10^6 10^6	250 30	80 μb 2.5 μb	ECL gates, PROMS, etc.	Fenker, Green	\$10K	testing	Fall 81
605	$pA \rightarrow h_1 P_T$ pairs, singles leptons, hadrons	L: Scint. hod. 2 trks H: clean DC trks, P_L M, particle ID.	25 ns 3 μs	3×10^7 3×10^4	1000 100	10 nb 100 pb	ECL RAMS Modular ECL Trig. Proc.	Glass Sippach, Knapp, Heuing, Benenson	\$ 5K \$75K	construction	Winter 82?

NOTES: 1. All numbers are approximate and intended for general information only.
 2. L: Low level trigger. H: High level trigger.
 3. Substantial design and diagnostic costs are not included in these approximate numbers.

are used to take advantage of the complexity and power of modern large scale integrated circuits and memories without saturating the experimental dead time.

In most of the experiments to be described in this review the triggering is done in two steps. Typically the first step, which we will call the low level trigger, filters the data down to $<10^4$ /sec. This trigger is often used to latch and digitize the data in various conventional CAMAC type modules where it waits until a further decision, the high level trigger, is made whether or not to pass on the data for recording on tape. The low level trigger commonly also gates fast digitizing electronics which provides the data for the high level trigger processor.

Table 1 summarizes relevant facts concerning the triggering systems I will describe. The Fermilab beam that these experiments must work in is typically on for a one second flat top with a repetition period of ten seconds. Within the one second flattop the beam is primarily structured only by the 53 MHz RF frequency of the accelerator. This results in the beam coming in "buckets" of 1 nsec length every 18.6 nsec. This structure means that low level trigger timing need not be refined below ~ 15 nsec, a great convenience. Beyond this RF frequency the duty factor of the flattop is $\sim 80\%$, more or less, depending on how hard the control room is pushing for high intensity. Neutrino experiments use a fast spill of 1 msec length, one or more times per accelerator period.

Low Level Triggers: The First Data Filter

In many cases a straight forward low level trigger made with conventional commercial NIM modules is used before a high level trigger. However, a number of experiments have developed sophisticated, home built, fast trigger systems to meet their low level needs. An early example was the Bellcord system designed by Cordon Kerns¹ for E272 which measured the radiative widths of K^* and ρ in a 1978 run in the Meson Lab.² The Bellcord system (Fig. 1) operating at a rate of 5×10^5 /sec counted the multiplicity of tracks in MWPCs and triggered a gate to digitize pulses from a liquid argon π^0 detector in 60 nsec. The, essentially analog, system consisted of 16 coaxial cables each with 64 inputs from sense wires. A hit on any of the ~ 1000 wires generated a standardized pulse on the coaxial cable where it was summed with pulses from other hits. The summed pulse was evaluated at a level detector to determine if a predetermined multiplicity requirement had been met. The Bellcord trigger was combined with a signal from an ECL triple coincidence matrix of MWPCs that vetoed unless there was a useful interaction. The combined low level system triggered the liquid argon ADC gate at a rate of 20-50/sec. The name comes from the bell cord on city buses which allows any of the (presumably less than 1000)

passengers to request a stop.

Early this spring Experiment 326, a Princeton-Chicago experiment studying $\pi^-A \rightarrow \pi^+\pi^-x$ on various nuclear targets, concluded its run in the high intensity Proton West pion beam. A visitor to this experiment soon notices the almost complete lack of commercial NIM and CAMAC equipment. R. Sumner, with the Princeton engineer Marias Isaila, has designed and built a complete set of state of the art electronics. This is the only running experiment that has an example of a programmable system in each of the three areas being covered here, low and high level triggers as well as a raw data readout system. The latter two systems will be covered in following sections.

Because of the high interaction rate of several times 10^8 , this experiment's low level trigger runs essentially at the 53 MHz rate of the beam RF structure. The experiment (Figure 2) consists of 7 toroidal magnets with drift chambers and scintillation counter arrays in the gaps between the toroids. The scintillators, used for both low and high level triggers, are arranged in octants. The radial size and number of the scintillators in each gap varies to provide appropriate resolution and acceptance. There are a total of 34 scintillators/octant, 272 overall. The low level trigger operates on fast (differential ECL) pulses from these counters and fires a gate to latch the scintillator and digitize the drift chamber pulses a few hundred times/second. This selection of 1 out of $\geq 10^6$ events is accomplished with ECL coincidence matrix arrays that select scintillator hits in the first few gaps that correspond to tracks that point to the target. Three 8x8 hard wired AND gate arrays are used in each octant. Two are used to form 5x7 matrices of the hit patterns between gaps 1 and 2 and between 2 and 3. The third is used to require a hit in the fourth gap. The coincidences are three fold ANDs with the third input connected to a computer loaded latch. Software is available to allow easy online changes of the trigger matrix requirements.

The coincidences are ORed on each board and (Figure 3) the outputs of the matrix boards for each octant are ANDed to fire one of 8 updating 20 nsec one shots. These are connected to the 8 address lines of a ~10 nsec access time 256x1 Fairchild RAM. This memory is loaded with acceptable octant patterns and its output is the low level trigger, running at 200/beam spill, that generates the latch and digitizing gates and fires the high level processor. The length of the 20 nsec one shot is chosen to allow for inter toroid time jitter and the memory access time. The convenience of the Fermilab RF frequency is apparant in this design. We will return later to further discussion of Dick Sumner's elegant and inexpensive electronics.

There are two other experiments setting up to run within a year that will use somewhat similar coincidence array low level triggers. A Fermilab, Florida State, Tufts, Notre Dame, Arizona, VPI, Vanderbilt group, Experiment 623, intends to measure $\pi p \rightarrow \phi \phi$ inclusive (searching for n_c and glue balls) with the Meson Area's Multiparticle Spectrometer. They have just completed a successful test run. From the simplified view point of the trigger the experiment is shown in Figure 4. A basic pre-trigger, made with commercial NIM modules, is used in this case to latch and digitize data and start the low level trigger on any beam interaction. With a beam of 10^7 pions on a 10% target, the low level trigger will have to operate at 1 MHz. The bend plane MWPC wires are ORed together to form 32 channel "hodoscopes" out of each plane.

A triggering system for this experiment has been designed and is being tested by D. Green and H. Fenker of Fermilab. The low level trigger³ for this experiment examines latched data from the 3 MWPC "hodoscopes" measuring the bend plane, selects the 6-22 GeV momentum band appropriate for Cerenkov identification of K^\pm , identifies and counts K^\pm tracks, and allows further decision making if there are at least 2 K^+ and 2 K^- tracks. A simplified drawing of this dedicated, almost hard wired, system is shown in Figure 5. Programming flexibility is allowed by pins that can connect the gates to individual lines on the various busses. Each of the 32 channels from PWC1 gates two 8 bit priority encoders, one each for positive and negative tracks. The inputs to each pair of encoders are 16 coincidences of PWC2 and PWC3 vetoed by the appropriate Cerenkov signal, corresponding to the $6 < P_{K^\pm} < 22$ GeV roads for the particular PWC1 channel. These encoders act both as the final step in this triple coincidence and to provide encoded data for the subsequent high level computations.

This system, and the high level part to be described later, is an example of what two physicists, having a basic understanding of logic and digital circuitry, can accomplish for the complex needs of an experiment when they plug ECL 10000 integrated circuits together instead of commercial NIM modules. Between them, one year was spent on design and assembly. Also preparing to run next year is the Columbia-Fermilab-Stony Brook, Washington, KEK, Kyoto, Saclay, CERN collaboration with Experiment 605, son of the famous epsilon dimuon experiment.

This will be a very big experiment (Figure 6) by any standards. Plans for trigger electronics are large in scope and advanced in design. The low level trigger will generate digitizing gates and start a high level trigger processor (described in a later section). Similar to the two low level triggers discussed above, this trigger uses a triple coincidence array connected to three hodoscope planes to make a preliminary selection of tracks in a very high rate environment. The three banks of scintillation counter hodoscopes (called Y_1, Y_2, Y_3) are located with the wire chambers,

one each at stations 1, 2 and 3. The hodoscopes, and the trigger, are being designed to operate at rates of 50, 25, and 12 MHz over 12, 17, and 13 channels respectively. The 3 dimensional array space has two bands near the diagonal which, according to Monte Carlo studies, will correspond to acceptable trajectories for positive and negative tracks. Individual 4x4x4 cells of the large matrix are built using 12 nsec access time Fujitsu (MB7072) 256x4 RAMs with 4 AND gates (Figure 7). Packaged 6 to a single width CAMAC module the memories can be downloaded with triggering patterns through the CAMAC system. The outputs of the cells are ORed together to provide the low level trigger. Note that only cells in the vicinity of the diagonal bands are required as an absent cell is equivalent to a 0 in memory. It is expected that less than 100 cells will be required to cover the triggering bands. The trigger should safely operate in 25 nsec and will thus handle even the rates of this experiment in a 3×10^{12} protons/spill beam. The design of this system, which is presently being assembled, is due to H. Glass of Stony Brook.

High Level Triggers: Trigger Processor Systems

We turn now to high level triggers, an area where the availability of LSI circuitry has allowed great advances in experimental electronics. Three large integrated triggering systems have been designed for use in Fermilab experiments.

The first processor to be designed and still the most (visually) beautiful piece of electronics at Fermilab is the Magnificent Multi-Muon Mass and Momentum Monitoring Machine, sometimes referred to as the M7.⁴ It was designed by Fermilab's T. Droege, I. Gaines and K. Turner and built in 1977. The M7 (Figure 8) is a stored program device that has selected features of a general purpose computer, slimmed down, souped up and optimized for the high speed reconstruction of spectrometer tracks. The architecture and instruction set was based on a careful analysis of the essential computations involved in track reconstruction. In one major machine cycle of ~145 nsec the M7 can carry out instructions of the form $E_k = A_i \times B \pm C_j \times D$ as well as those where the \pm is replaced by any logical operation available in an ALU. A 4-ported memory (built of 4 identically loaded but independently readable 1024x12 bit memories) provides the operands for such instructions. Two of the operands and the destination are indirectly addressable using index registers. The flexibility of the 64 bit instruction word and the simple register structure has made the ad hoc addition of new instructions a relatively sanguine matter. Intended as a one of a kind machine, the M7 is made of ECL 10000 logic mounted over ground planes and interconnected via wire wrap posts and could be duplicated with little difficulty.

The original intention was to use the M7 to trigger on high mass dimuon events to be produced in neutron induced interactions in Experiment 400 in the Broad Band Photon Beam. A program was written to find and measure dimuon tracks in drift chambers and compute their mass. This program was used to debug the M7 using simulation events generated by a PDP10. Test programs were written in BASIC, an interactive language typical of those useful for debugging this kind of system.

The first on line use of the M7 has, in fact, not been E400 but rather a sister experiment on the same spectrometer, E401, (Figure 9) which was a measurement of ψ photoproduction in e^+e^- and $\mu^+\mu^-$ final states run in Fall 1979. This and subsequent work with the M7 has been carried out by D. Harding. The pretrigger in this experiment involves a conventional NIM coincidence of a beam veto, the TC counter and at least two hits downstream. This latches the scintillation counter arrays and MWPCs. The low level trigger is then carried out using the Sippach DC register logic system. This fired $\sim 300/\text{sec}$ on dimuon or dielectron signatures in the muon hodoscopes and the electromagnetic shower counters, subject to an additional multiplicity requirement.

Instead of drift chambers, in E401 the M7 was connected to the non-bend plane 2 mm MWPCs (P0-P4) whose wires were ganged together in groups of 8 or 16. A program of 175 instructions identifies tracks pointing to the target. Starting with hits in the most downstream chamber P4 it searches for hits in P0 within a window of a straight line to the target. If it finds a hit in P0 it searches for hits in P1,P2,P3 on the P0-P4 line. The trigger requires either 1 perfect 5 chamber track or 2 lower quality tracks. This cuts the trigger rate by $\sim 1/2$ to 160/sec, below the saturation limit of the experiment's data buffer.

The next use of the M7 will be on E400 which is expected to be running next fall or winter. It will be connected to drift chambers and proportional tubes in and behind the muon wall. A program similar to the original one will be used to compute dimuon masses. In addition there is expected to be a K trigger (correlation of high momentum tracks with void Cerenkov counter cells). Possibly also a downstream (charm) vertex will be reconstructed based on a set of 1/4 mm wire chambers to be added for this run. The goal is to obtain a factor of 4 trigger reduction with the M7 in this experiment.

In order to filter the data further and allow running at still higher luminosity, there is a plan to expand the event buffer and run the M7 loaded with portions of the full offline track reconstruction program in the 10 sec inter spill period. The reconstructed information would be used to select events to be written on tape and to get a big jump on the offline work. It is believed that the M7 could hold the complete track reconstruction package and process an event in 10-15 msec. Thus ~ 500

events could be processed and selected between spills. Because of limits on physicist time available to translate programs from Fortran to M7 code, it is not expected that the full program will actually be done online. Two of 13 subroutines of the main program have already been converted, an impressive start which presages reconstruction processors, a subject I will return to in a later section.

In 1978 a second major trigger processing system, modular in concept, was designed at Fermilab by E. Barsotti, S. Bracker (Toronto) and myself with M. Haldeman and the Beam Systems Group of Research Services⁵. J. Martin (Toronto) was heavily involved in setting up the first application of the system for E516 and in designing the recoil detector to which it was connected. This system is referred to as the ECL-CAMAC Trigger Processor System⁶ and was built for the new Tagged Photon Spectrometer. It has been used by E516, a charm photoproduction experiment being carried out by a group from Fermilab, UC Santa Barbara, Toronto, Colorado, Oklahoma and Carleton/NRC. It was conceived as a large extension in power and flexibility of the modular NIM fast logic philosophy, familiar to physicists, rather than as a streamlined, high speed, instruction programmed computer like the M7. The system makes heavy use of large arrays of fast access memories as table look ups in place of combinational logic. Programming this processor is done by cabling modules together at the front, as in NIM fast logic but with multi twisted pairs, as well as by loading control and memory information into the modules via CAMAC. The modules operate in an unlocked, non handshake mode. Far more powerful than fast logic, these modules operate on multi-bit words, finding tracks, calibrating pulse heights and computing complicated functions. Single module operations typically take 50 nsec.

Although the system operates in a high speed unlocked mode, system algorithms may include hardware processing of nested loops, conditional branching and subroutines, operations normally associated with serial computer-type processing. Unlocked operation means that each module only waits for data specifically required at its input, minimizing process time. Operation is controlled by a system of Ready lines: a module starts when all required input Readys are set and when it has completed its operation it raises its output Ready. Except in esoteric loop situations, intermodule timing need not be a concern of the user. Timing between two asynchronous loops is handled automatically by the Stack module functioning as a dual ported random or serial read and write buffer memory.

To permit system testing, nearly every intermodule data path can be monitored or controlled by CAMAC read and write commands. The system is packaged using CAMAC hardware and a modified CAMAC crate with an ECL backplane and a CAMAC TTL to ECL translator in slot 23 next to a conventional Type A Crate Controller. As seen from the host computer the system responds and transmits according to CAMAC standards.

Three assembly techniques were used. Low volume modules were assembled on an ECL wire-wrap CAMAC Kludge-board designed for the project. Modules required in larger quantities and designed in the early stages were constructed using three layer printed circuit boards with the middle layer a ground plane. This technique is limited somewhat in chip density because of layout requirements introduced by ECL signal path requirements. Later modules were produced using the "Multi-wire" technique, a commercial process that uses a computerized wire laying machine to set small insulated wires in an adhesive insulator over the two sides of a PC board containing a power grid and ground plane. A well defined transmission line impedance is obtained with Multiwire and layout expenses are reduced since only a key punched wire list is required.

The ECL-CAMAC system has been extremely reliable. E516 has been running steadily, with only a 10 day shutdown, since November. During this entire 5 months it has not been necessary to replace a single module out of the 80 modules in 6 packed CAMAC crates. A total of almost 200 modules has been built to date about half for E537, the second major experiment to use the system. There have been more than 30 different module types designed (including about 6 for E537).

The heart of the system is the Memory Look Up (MLU) module, a high speed table look up device which contains (up to) 16 35 nsec Fairchild 10470 4K ECL RAMs. Over 40 of these units are now in use. With these memories it is possible to precalculate functions of arbitrary complexity, store results in tabular form, and, when needed, access the appropriate location in memory to obtain a real time answer in 50 nsec or less. The MLU is an extremely general device in a flexible and easy to use package. It accepts up to 16 bits of input which is used to address the internal memory. The number in the addressed word is the MLU output. Both input and output fields may include several variables allowing multi-dimensional functions. A front panel switch allows a choice of 5 output (and thereby input) word length modes. With 12 input bits there may be up to 16 output bits, with 13 in, 8 out, and so forth. If output bits are not needed appropriate RAM chips may be left out, or 1K RAMs can be used to reduce the MLU time from 50 nsec to 35 nsec.

The physics user of the MLU does not need to be concerned with the internal memory organization. To generate an MLU load, a Fortran subroutine is written relating the input variables in vector form to the output variable vector. A "compiler" on the online PDP11 calls this subroutine for every possible input value and generates the correct bit pattern for the memory which is then stored on disk. From the disk file it can be downloaded through CAMAC with a simple procedure in about one sec per MLU load. The "compiler" accesses a "configuration file" which contains all relevant information about the processor configuration such as which variable is on which bits at the input.

These functions can be very complex, involving whole reconstruction algorithms and many sub-subroutines. Several MLU loads involving 8192 different addresses for E516 take as long as 45 minutes to be computed by the PDP11, or 0.3 sec/address. The MLU operating at 50 nsec is thus almost 10 million times faster in making this computation. Examples of MLU applications in use include: Fast pulse height calibration, high speed multiplies and divides, trigonometric and logarithmic functions, missing mass calculation, dE/dx pattern recognition and particle identification, track projections, dimuon mass calculation, process control, conditional branching, and as a single module, programmable, timing independent replacement of conventional NIM logic. Data passing between MLUs is often packed and encoded with complicated functions in order to optimize the precision of the final calculation over the complete parameter range despite the relatively limited number of bits available. This packing is made part of the MLU function.

The Stack module contains 32 words of 16 bit wide 12 nsec memory. Internal arbitration circuitry allows simultaneous, asynchronous read and write operations at up to 25 MHz. Both read and write can now be done either sequentially or under random access. Several modules may be linked together to expand the storage in either the word or bit direction. The Stack has two principle applications. One is to receive and hold raw data from the fast ADCs and TDCs and the MWPC and drift chamber latches. The second function is to act as a FIFO (or random accessed) buffer between asynchronous loops.

Another important module is the DO Loop Indexer which is used primarily in conjunction with two stacks in track finding situations. The module acts as a controller that cycles through all values of two indices (I,K) much like a Fortran double index DO loop except that the upper limits (and the lower limit of K) may be increased by external flags while the loop is in operation. This allows the indexing of Stacks while they are still filling with data. In typical applications this module adds only 5 nsec to the loop time as it precomputes all possible I,K pairs while the rest of the loop is operating and, at the last moment releases the pair appropriate to the index limit flag status.

A track finder subsystem (Figure 10) was designed specifically for E516 to find straight line track segments in three evenly spaced wire chambers (I,J,K). The system has turned out to be of more general use than originally thought. By adding MLUs, as has been done in E537, the basic sub system can also be used in situations with arbitrarily spaced chambers. It could also be used to correlate multiple hits in X,U,V chamber sets which involves an identical algorithm to that of the evenly spaced chambers, as well as for other track projection and matching operations.

The track finder system uses 3 Stacks (two contain I and K chamber hits and the third stores the output track parameters), a DO Loop Indexer, a special Track Finder Module (TFM). The middle chamber J hits are stored in a hit array located in the TFM (or in a new Hit Array module designed for E537). A pair of hits selected from the I,K stacks by the DO Loop Indexer are used by the TFM to compute a projection to the middle chamber. This projection is used as an address to access the middle chamber hit array which contains a bit on at the projection address if a hit has been found. Simultaneously, track parameters (track slope and vertex for E516) are computed and, if a bit was found, output to the Track Stack. The MWPC centroid processor can load extra bits around a hit in the hit array and generate a window wide enough to account for projection error. Alignment errors can be similarly compensated by shifts in the hit array bit pattern. A single track finder loop takes ~130 nsec. The time to find n hits increases as n^2 (rather than n^3) because of the two dimensional search algorithm that is used. Thus, an event as complicated as one with 10 tracks can be unscrambled accurately in only 13 μ sec.

Other general purpose modules in the original collection include several test modules, level converters, a General Logic Module (simple ANDs, ORs), a Fan Out Module, and a Quad 4-bit Scaler unit used to count event associated quantities, such as number of tracks, which can be used for diagnostic monitoring or trigger decisions. Added recently, for E537, are an Arithmetic Logic Unit, a Router and a Comparator Module, which were designed and built at the University of Athens.

In E516, a charm photoproduction experiment, this extensive system has seen its first use. The experiment uses a conventional NIM low level trigger which is a coincidence of a signal from the photon tagging system and an indication from electromagnetic and hadronic calorimeter energy sums that there has been a hadronic interaction. This intentionally loose trigger runs at ~ 2 1/2 times the total hadronic cross section vetoing most of a large e^+e^- pair background. The trigger processor is connected to a sophisticated recoil system and the tagging system. The recoil detector has 3 cylindrical MWPCs concentric around a 1.5m liquid H_2 target. Outside these are 4 cylindrical annulus layers of scintillator, two plastic and two liquid, in 15 azimuthal sectors covering 15/16 of 2π . The scintillators measure the dE/dx of tracks from the target allowing a determination of energy, E , and particle type. The MWPC coordinate read into the trigger processor comes from an induced cathode signal and measures the distance along the beam line and thereby the production angle, θ . The pulse heights are digitized in 1.4 μ sec (for 8 bits) in a fast ADC/TDC system designed by Cordon Kerns.

The processor carries out a complete recoil detector reconstruction in ~ 10 μ sec (which in embellished form takes ~ .2 sec on a Cyber). It reconstructs all tracks, associates each track with an azimuthal sector by comparing the projection into the inner scintillator with end to end timing information from the scintillator, calibrates the phototube pulse heights (correcting for pedestals, attenuation, and scintillator saturation), determines particle type (π or p) and E from the pattern of dE/dx in the 4 layers, selects the most upstream vertex and counts the number of tracks at that vertex, identifies and counts the number of neutral tracks, calibrates and determines the photon energy (k) from tagging system information, computes the missing mass ($M_x(k, \theta, E)$) recoiling from the most upstream proton, and finally if, there was exactly one proton track at the most upstream vertex and M_x is in a predefined range (2.5-11 GeV in E516) and other scaler counts are acceptable, it triggers. Figure 11 shows 3 histograms of M_x as computed by an independent offline reconstruction program for data taken in a Spring 1980 test run under 3 different online trigger processor M_x range cuts. The trigger defined above has resulted in a factor 30-40 reduction in trigger rate. The processor has successfully and accurately carried out the trigger it was programmed to do. Whether the physics resulting from this choice of trigger is an equal success remains to be seen. Figure 12 is a simplified drawing of this trigger. Note that there are 4 asynchronous loops: the neutral pattern loop, the track finder loop, the big track loop (known familiarly as the grand loop), and the sector finder loop (embedded in the grand loop). Another paper presented to this conference by J. Martin (Toronto) provides details on experience with this processor in E516.

As has been noted, a second major experiment is just starting to run with an ECL-CAMAC system processor. This is E537 (Figure 13) a Fermilab, Athens, McGill, Michigan beam collaboration measuring di muon production in a \bar{p} enhanced negative beam in the High Intensity Area of P West. Sergio Conetti and Hari Areti of McGill have been principally responsible for this installation. The conventional NIM low level trigger requires a beam track, ≥ 2 quadrants of a triple coincidence of hodoscopes in a Fe-Concrete muon wall, \geq two tracks in a hodoscope just in front of the wall, and a counter with a hole just downstream of the Cu absorber. This trigger is now running at several hundred events/second. The high level trigger makes heavy use of specialized versions of the ECL-CAMAC track finder system containing MLUs and is designed to select events with at least two clearly identified muons with mass above a predetermined value. The trigger has been extensively tested using tapes of raw data downloaded into a new test module which transmits the data at real time speed into the processor. Thus even before having live experience the experimenters were confident of obtaining a factor 5 reduction of data rate for a mass cut set at > 2 GeV

with the parts of the system presently to be used and a full factor of 10 when their complete system is up (later this spring or in fall).

The triple coincidence of muon counters are latched and used by the trigger as roads defining acceptable areas of the drift chambers for track finding. A system designed by Cordon Kerns uses the road latch to strobe out appropriate hodoscopes and DC wire latch signals through a priority encoder to the track hit Stacks. In the first step the processor finds pairs of hits from the bend plane views of D4 and D6 and computes their slope, α , and bend point intersection, a measure of the projected production angle, θ_x . This is stored in a Stack where it is available for the next asynchronous loop. In this loop three parallel questions are asked: Does the track candidate extrapolate to the hit muon counters? Is the momentum (P , determined in an MLU assuming a point target from α and θ_x) above a predetermined value (now set at 4 GeV)? Does this candidate extrapolate to a bit in the hodoscope CPX? If all answers are yes, P , θ_x and a road index are put in a (pair of) stacks for use by the final mass computing loop. The procedure to this point is repeated for all hit roads, and, if necessary, for other DC combinations. To compute the mass correctly the non-bend plane production angle, θ_y , is also required which is determined from hits in the hodoscope CPY. In the final loop the mass of all muon track pairs from different Roads is computed in a series of 4 MLUs: one computes $(P_1 P_2)^{1/2}$, the second the bend plane opening angle which is combined with the y projected opening angle in a third MLU, and the fourth MLU computes the mass and triggers if in an acceptable range. The y plane angle is presently computed by taking hit differences in CPY. When there are more than two hits ambiguities will reduce the effectiveness of the trigger. A more sophisticated θ_y computation has been built, but not yet implemented, which sorts out the ambiguity using quadrant information from the muon counters. When connected (later this spring or in fall) this will add a factor of 2 to the trigger reduction which is expected to then be a full factor of 10. Further details about the trigger work in E537 is contained in a paper presented to this conference by H. Areti.

The use of the modular ECL-CAMAC system is beginning to move further afield. A small system is presently being prepared at Berkeley for a nuclear physics experiment. A large collaboration including several bubble chamber groups has recently proposed (P698) a large facility including a high repetition bubble chamber, a streamer chamber and a spectrometer. They require a trigger (on high P_T muons from bottom decay) for which an algorithm like that of E537 appears appropriate.

Not as far in the future, E605, the large high P_T experiment mentioned earlier, will use many aspects of a large modular parallel pipelined trigger system similar to the ECL-CAMAC pipelined trigger in philosophy but different in that it will be clocked. This system has been designed over

several years by W. Sippach and B. Knapp and is described in in the addendum to Fermilab Proposal 627 where the complete on line reconstruction of tracks in a high multiplicity, new multi particle spectrometer is discussed. This proposal has not been approved, but the processor system will be used in a Brookhaven experiment and in a E605. In E605 the processor will be connected to the wire chambers at the 3 stations shown in Figure 6. It will find straight line tracks to the target in all views (bend and nonbend plane). The trigger will fire depending on how many tracks are reconstructed in these views and may also include a P_T and mass calculation. The experiment will have a large event buffer, the "Mega Memory", at the end of a high speed, nonstandard, "Transport Bus" allowing up to 1000 events to be written on tape per spill. Trigger processor track finders will be hung on this bus and will select events to go into the Mega Memory.

Sippach and Knapp have designed their processor system with the most general reconstruction problems in mind. They have developed a consistent set of principles for inter module pipelined communication. Each module will contain several holding registers and data between them will be transferred from register to register under control of a central clock, expected to have a 25 ns period. Three bits of control information accompany the data in the processing path: a Valid bit, a block transfer Complete bit, and a Hold bit indicating that the present communication must be postponed. The latter two bits allow for a more complex, structured large scale system than the single Ready bit of the ECL-CAMAC Processor. The Hold, starting at a module not ready to receive data, propagates backward up the pipeline, one module per clock cycle, just as the Valid goes forward. This important feature is accomplished with an extra holding register in each module. It allows the possibility of module operations taking longer than a cycle (not presently used) and, most important, it allows conceptual ease and optimal synchronization in complicated structures with converging pipelines. Conflicts never exceed the worst case single module interference.

The clocking and all the extra registers have another important advantage for complex systems: one can simulate the entire processor offline with a Fortran program that describes the precise state of every register for each clock cycle. Each module type corresponds to a separate subroutine. Given a module list and cable connections (a Configuration File in the ECL-CAMAC Processor language) the simulation program can provide an accurate representation of the system at any moment in time. Although for simple systems a simulator like this can be effective even for an asynchronous processor, the clocking becomes very important for complex structures.

The system avoids standards like CAMAC and the Fast Bus, but it does include a control/communication bus link to the online computer. This can be used to access any register in the system from the outside, allowing diagnostic exercises and memory loading at rates ~ 1 MHz.

The modules, built with ECL 10000 circuits, are kept conceptually simple by the requirement that operations between their input and output registers be complete in one 25 nsec cycle. The modules are interconnected like the ECL-CAMAC system on flat ribbon cable but here more than a dozen modules are permitted on one cable bus. At least 12 module types have been conceived to date and we list these below. Many of these will be used by E605. Not surprisingly there are strong parallels between these modules and those of the ECL-CAMAC Processor System described earlier and these are noted in parenthesis.

- I Binary index generator (DO Loop).
- L List, in 3 versions including one with 2 ports indexed by "pages" (Stack).
- N Normalizer, computes $ax+b$ from table look up (MLU).
- + Adder, allows ALU operations (ALU,Quad Scalers).
- M Map (Hit Array).
- ≥ Comparator compares input with predetermined limits.
- S Switch merges two lines onto one (Router).
- P Page counter allows indexing of data by "page".
- Multiplier.

A simple example of how some of these moduls are used and the symbol notation that has been developed is the 3 chamber straight line track finder shown in Figure 14. This is very similar to the track finder described earlier (Figure 10). More complex subsystems built from these building blocks will be used in E605 to find lines in more than 3 planes, to find bend angles, and to match views.

Dedicated High Level Trigger Systems

We turn now to a group of triggering systems somewhat smaller in scope than the large processors just described but no less imaginative in meeting the specific needs of their experiments. We start with the high level system Dick Sumner has designed for E326, the dimuon experiment in the high intensity pion beam (Figure 2) that was discussed in the low level trigger section.

The E326 trigger is intended to select opposite sign dimuon events in a beam with a halo of 10^7 muons/sec. As we have seen the low level trigger selects > 200 events from more than 10^8 interactions/sec. It was originally expected that the low level rate would be several thousand/sec.

Since the rate is an order of magnitude better than this, the high level trigger does not have to work as hard as it might. Looking at the 272 latched scintillator signals, 34/octant, the trigger compares the latch bit patterns with a set of up to 1024 patterns deemed acceptable by a Monte Carlo program and stored in a 32 x 1024 ECL RAM array. In practice only 291 acceptable patterns were needed. The patterns in each octant were encoded in 21 bits, 7 scintillator planes x 3 bits each. Sequentially stepping through the 291 patterns at a 25 MHz rate, each stored pattern was compared with the 8 octant latched patterns simultaneously. Cycling through the 291 patterns thus required $\sim 12 \mu\text{sec}$. Each stored pattern carried sign information with it, and a count was kept of the number of octants with patterns corresponding to +, -, or ambiguous signed tracks. A trigger was generated if exactly 2 tracks of opposite or ambiguous sign were formed or if there were > 2 tracks. The final data recording trigger rate was about 10 per spill.

Because of the anticipated high low level trigger rate, the memory was designed to readout serially with a 10 ns period. Since the Fujitsu 256 x 4 12 ns RAMs were not available at the time, Sumner linked 4 20 ns 10410 256x1 RAMs through a 10141 4 bit parallel loaded 10 ns shift register to create a 100 MHz 1024 bit serial memory. 32 such units made up the complete track pattern memory which could be loaded conveniently from the online computer. In order to compare stored patterns with latched track patterns at the 100 MHz design frequency an 8x1 multiplexer was used as shown in Figure 15. One of the up to 8 lines from the latches of one scintillator plane was selected onto the output of the multiplexer chip by the 3 bit pattern for that plane. Identical arrangements for each of the 7 planes (and the 8 octants) was used in parallel and the 7 outputs for each octant ANDed to signal a pattern match.

Experiment 663, a UC San Diego, UC Davis, MSU, Carleton experiment in the Meson Area to measure $h^+ \Lambda(\bar{\Lambda})$ inclusive polarization ($h=K^+, \pi^+, p(\bar{p})$), uses a small magnetic spectrometer including 8 layers of MWPCs (4X and 4Y views) and a scintillator wall. The experiment, at last word, is in the throes of beam testing preparatory to its run this month. The conventional low level trigger uses NIM logic to require a track in the beam telescope and a hit in the scintillator wall away from the beam. This is expected to run at 400/spill with a 10^6 negative particles/spill beam which is 4-5 times higher than the online computer can handle with reasonable dead time. To reduce this rate a simple but clever high level trigger has been designed and built by Hans Kobrak, S. Gourlay and H. Melanson around the existing shift register MWPC readout of the experiment. The data shifting out from each plane at 10 MHz is now passed through a card which counts the hits within a defined window. The window can either be defined by down loading from the online computer or (in the future) by having a hit in an

upstream chamber determine a downstream plane's window. The count pattern in each of the 8 scalers are formed into an address to a 64K RAM array which contains on bits at the addresses of acceptable patterns. The memory is made up of 16 Intel 2147 55 nsec 4K RAMs. The efficiency of potentially acceptable patterns was determined by Monte Carlo. Convenient interactive software allows physicists to try out different lists of patterns with simple commands to the Eclipse which downloads the patterns through CAMAC. The goal is to reduce the data recording rate below 100/sec so that the dead time will be less than 30%.

These experimenters demonstrate how the availability of LSI chips can make possible powerful, though simple, triggers even within the constraints of limited personnel and financial resources of a small group in a small experiment.

Another small group that has built a trigger from ICs instead of NIM modules is the Rockefeller University E612 experiment in the Tagged Photon Beam. This experiment uses the TREAD detector, a 15 atmosphere H₂ detector similar in concept to the TPC at PEP, as a target to study low |t| recoiling protons in photoproduction. TREAD contains 8 concentric rings of signal wires surrounded by two sets of 16 azimuthal scintillation counters within the pressure vessel. The low level trigger requires a signal from the photon tagging system and any of the scintillors with appropriate antis. Then about 500 times per spill a trigger designed by S. White determines whether or not to digitize and read out the analog information from the wires and scintillators. This jumper programmable trigger is built with ECL 10000 circuits. It requires that 4 of 8 sense wires be on and timed in 20 μ sec windows corresponding to a track with production angle 45°-90° striking the scintillation counter that fired the low level trigger. The present trigger rate is ~ 10/spill. The system also provides control information for the analog multiplexers that route sense wire pulses to sample and hold circuits.

The high level trigger of E623, the Multiparticle Spectrometer measurement of $\pi p \rightarrow \phi\phi$, uses the encoded K^+ bend plane momenta and angles obtained during the low level phase of this trigger (described earlier, Figures 4 and 5) to compute the bend plane mass and P_T of the left most and right most K^+ pairs (Figure 16). It sets bits on 4 lines corresponding to M_L , M_R , P_T^L , and P_T^R being within programmable ranges. These lines are used by simple NIM logic to generate a data recording trigger. The mass and P_T are computed primarily in table look up memories, similar in concept to the MLUs of the ECL-CAMAC system but here implemented with 10149 PROMs. Some programming is also done via DIP switches and comparator circuits. The trigger takes 150 nsec (75 nsec for each of the low and high level phases) and was designed to be fast enough to fire the spark chambers which are still in use in the MPS. PROM changes would allow the system to be used for other pair triggers such as $\mu^+\mu^-$, $\bar{p}p$ or $\Lambda\bar{\Lambda}$.

A neutrino experiment such as E594, a Fermilab-MIT-MSU-NIU group which is preparing to measure $\nu_{\mu}e$ elastic scattering, is the last place one would expect to find a special triggering system. Instantaneous rates in such an experiment are obviously low. Wrong! With 5ν interactions/1 msec spill, the instantaneous rate is 5 KHz, similar to that of the other high level triggers we have discussed. When fired, the flash chambers of this experiment cannot be refired during the 1 msec spill. Since the cross section of interest is much smaller than even the neutrino cross section ($\sim 10^{-4} \sigma_{CC}$) this would introduce a substantial unnecessary dead time. Thus we find that when ν experiments focus on small parts of the total cross section, their triggering needs are surprisingly similar to those in hadron or photon beams.

Between the planes of sand and shot of this experiment (Figure 17) are flash chambers. Every sixteenth gap contains a proportional tube hodoscope of 1" granularity. Four of these 1" tubes are tied together to form one cell for readout. The total mass of the detector now 70% complete will ultimately be 340 tons. The low level trigger fires on any ν interaction using the 35 planes of proportional tubes, an upstream charged particle veto, and conventional NIM logic. A NIM charged current veto requiring no μ out the back of the calorimeter reduces the trigger rate by 1/3 but leaves the dead time still high at 50%.

To allow a further reduction, a high level processing system has been built and tested. This identifies the width and length of showers in the calorimeter as measured by the proportional tubes. The system, which has had design input from T. Mattison (MIT) and T. Droege and C. Kerns (Fermilab), uses priority encoders to determine the addresses of the first and last block of tubes in a given plane with signals over a minimum ionizing level threshold. Subtracting these addresses gives the width of the shower and a similar technique longitudinally can give the length. The system uses ECL logic to carry out this analysis in 600 nsec, fast enough to make a decision to fire the flash tubes reliably.

A typical electron shower is narrow and clean and can be selected by vetoing on "fat (hadronic) showers" with at least two planes with shower widths wider than 4 cells (16" or 40 cm). Such a veto successfully reduces the trigger rate, and dead time, by a factor of three bringing it close to the design goal of 1/2 trigger/fast spill. The problem is the presence at the few percent level of events which are clearly electron showers but have a track separated from the central region by a gap. The shower width criterion of the fat shower veto causes such events to be lost. The experiment has decided that, at least for its first running, it prefers to tolerate a larger dead time and fewer events to the, presumably, energy dependent systematics caused by the fat shower veto. After some data is accumulated and the biases of this trigger are understood it may be used

for further high statistics data accumulation. Despite the difficulties, this trigger work is very significant as it represents the first step of neutrino physics into advanced triggering and demonstrates how wide spread triggering efforts are at Fermilab.

Detector Monitoring and Data Acquisition Systems

Commercially made microprocessor controlled high voltage and ADC data acquisition systems have been in use at Fermilab for the past two years. The Lecroy Research Systems HV4032 is a 32 channel phototube high voltage unit that can be used with other units to service large systems of tubes. Via a daisy-chain of these units and a CAMAC module the voltages can be set and read back by the online computer. An F8 microprocessor continuously checks and maintains the voltages at the set values. The Lecroy System 2280 is a high density ADC data acquisition system that can handle up to 1008 ADC channels in a CAMAC crate. A microprocessor controller allows automatic pedestal subtraction and data "sparsification". Data is only read out in channels a user selectable number of channels away from a channel with data above a variable threshold. The program for the processors of both Lecroy systems are prepared and loaded at the factory to meet the needs of experiments. Perhaps for this reason and the NIH (Not Invented Here) resistance syndrome, which apparently works in both directions, experimenters have not taken advantage of the full capabilities of these powerful commercial systems.

Experiment 326, the Princeton-Chicago di μ experiment, uses a microprocessor controlled multi hit/wire drift chamber system⁷ designed by R. Sumner (Princeton), H. Sanders and R. Lenski (Chicago). Each 8 bit Signetics 8x300 micro controller scans 500 wires for hits subtracting the stop time and handling multiple hits. Only hits are read out. The E326 system has ~ 3500 wires costing about \$50/wire. Programs are stored in PROMs which experimenters were able to change to fit their needs because of the availability of a cross-assembler on their online computer.

Although gross features of the design of the Colliding Detector Facility for the Tevatron's $\bar{p}p$ colliding beam experiments are beginning to settle down, much of the design is still far from final. One thing that is clear is that programmable controllers and microprocessors will be essential features of the extensive data acquisition systems. The total projected cost of the facility, something between \$20 and \$60 million dollars depending on your accounting methods and how you take inflation into account, gives an indication of its scope. The detector will cover as much of the solid angle as possible with fine grained calorimeters and drift chambers to measure the secondaries from the $\sqrt{s}=2000$ GeV interactions.

Fermilab's Tom Droege and his group have been developing designs for a data digitizing and acquisition system for the CDF. They are designing for a system Droege estimates will involve (within a factor of 2) 20,000 TDC and 50,000 ADC channels. Each channel is connected to a sample and hold circuit. These are later strobed into a common ADC and read out. On the drift chamber themselves are the first two of four stages of amplification (MECL10115). To eliminate ground loops, power to the amplifiers will be sent on the multi twisted pair cable that will carry the signals about 10 m to a crate containing the digitizing electronics. These crates will be located at convenient locations around the detector. After further amplification at the crate the signals are processed by a simple Time to Voltage Converter (TVC) and this voltage held for later digitization. The chip count is ~ 2.5 per wire, so that quantity costs as low as \$5/channel are expected. In a similar manner signals for integrated charge digitization (ADCs) charge up a capacitor during a gate and the capacitor voltage is held. Each voltage hold circuits is capable of <1 mv/sec drift and is connected to two independent (for speed and redundancy) analog busses in the non-standard crates.

Each of the two busses in the crate (Figure 18) has its own controller (called the EWE II) which contains an ADC that can be switched onto a bus line. It also contains a D-A for threshold comparison and to monitor power supply voltages, temperatures, etc. The EWE II controllers are connected via a differential ECL twisted pair daisy chain to controllers in other crates at the detector and to a master controller located far away in a fast bus crate in the control room. The master controller, named the TDS, Table Driven Scanner, handles up to 1000 individual channels through the EWE IIs. It contains a large tabular memory containing addresses, labels, and instructions to the EWE IIs on what action is to be taken at the given address. The memory is stepped through sequentially and the instruction transmitted over the twisted pair cables to the EWE IIs. Digitized data coming from the remote crates passes through these cables to a dual ported output buffer in the TDS. A complete scan through the 1000 modules controlled by the TDS to find and digitize all hits should take 1 msec for a typical event. A processor will be located in each of these Fast Bus crates to allow testing and monitoring of different parts of the detector independent of the central online computer. Through the processors it will be possible, for example, to load the TDSs with instructions to a EWE II to connect a specific analog bus line from a single phototube directly to the twisted pair cables coming upstairs so that they can be viewed on a scope.

Although, as noted earlier, this does not yet represent a final and approved design, it does indicate the direction that will necessarily have to be taken to cope with the complexity of the CDF detector.

The Future: Reconstruction Processors

The inefficiency of general computers on reconstruction problems is most apparent for the data from high multiplicity, high luminosity spectrometer experiments where each of 20 or more million events will consume as much as a second of Cyber 175 time. The large main frames are inefficient for these problems because of their emphasis on performance in high precision, long word length, floating point calculations rather than short integer computations more relevant to the indexing and simple arithmetic required during reconstruction. Furthermore, limitations in memory available, though of long word length, often forces programmers into sacrificing speed for space.

For a long time to come conventional computers, flexible and easy to use, will be in heavy demand for final physics analysis and histogramming. However, experiments that can easily burn as much as a year of a large computer's time point to the need for dedicated reconstruction processors. A first giant step in this direction was taken several years ago by Paul Kunz at SLAC with the development of the 168E System using relatively low cost emulators of the IBM 370-168 main frame. The plans we described earlier to use the M7 at Fermilab as an inter spill data filter, of course, imply that the M7 could be used to process data just as well from tapes. This would give its programmers more time to prepare complete track reconstruction programs which it is estimated will run 3 times faster than on a Cyber 175. The use of the M7 in this way, perhaps slowed by the scale of program conversion, represents still another step toward taking the reconstruction load off the big CPUs.

E. Barsotti of Fermilab, S. Bracker of Toronto and I have on occasion discussed ideas for more advanced and powerful reconstruction processors. Bracker has spent time this winter studying the basic design considerations of reconstruction processors at some length and has summarized these in a series of notes ⁸ from which I draw heavily in the discussion that follows. One can imagine at least three rather different approaches, including the emulator approach of the 168E System. Another direction is an extension of the modular structure of the ECL-CAMAC triggering system or the Sippach-Knapp trigger processor for E605, both described earlier. That it would be possible to carry out a full online reconstruction in a high multiplicity experiment has been convincingly argued by Knapp in several proposals to Fermilab and other accelerators.⁹ Although offline reconstruction could in all likelihood be carried out with such systems with reasonable hardware cost effectiveness, it would not be with great flexibility or ease of programming. A third approach involves the use of a multi processor network and appears most promising at this point.

There are three important characteristics of reconstruction algorithms that encourage multiprocessing: events can be reconstructed independently of each other, the problem is easily (and usually) broken into a number of

independent substantial subalgorithms, and there are always a set of basic instruction sequences (such as finding straight line tracks) that are repeated so frequently that they account for a significant fraction of the computing effort.

Motivated by these characteristics, one could link a number of microprocessors together each with extra hardware specially configured to deal with a particular part of reconstruction. This hardware would include extensive memory for table look up to increase the speed of many computations. It would also include track finder subsystems such as that in use in the ECL-CAMAC triggering system. These special hardware procedures would be incorporated in the processor much like the sophisticated "instruction" set of the M7 trigger processor. If these processor units are linked serially (Figure 19) with each one passing complete results on to the next stage, bottle necks will clearly develop. Even if buffer memories are used between steps there will be statistical bottle necks due to differences in the average time taken by different processes which will cause significant inefficiency in hardware utilization even for small differences in the process times.

Since we are not considering real time processing, our emphasis is on designing good system utilization as opposed to elapsed time optimization. This approach maximizes event throughput for fixed hardware cost. From this point of view an approach better than serial processing is to network the processor units with several processors loaded with each procedural step. Data packets passing between the processors would be under the control of a dispatching processor which would keep track of which processor was ready for a new data packet requiring a specific procedure. In more familiar computing terms, one can think of the dispatcher as a main reconstruction program calling subroutines, the procedures in the processor units, which represent different aspects of the reconstruction and which are likely to have been prepared by different members of the experimental group. There is also a similarity in concept to array and data flow processing systems available for general purpose computers. The difference is that here the individual steps in the array are whole small procedures of the scale of track matching or peak finding.

The dispatcher would also have some capability to change the procedure loads of the individual processors and to balance and optimize processor utilization. Each processor should have the capability to efficiently handle more than one procedure without reload. Preliminary indications are that with 10 procedures and 20 or 30 processors, each capable of executing two procedures without reload, a system would run with little time lost to idling. The plan is to use simulations this summer to study networking arrangements and processor design in detail.

An example of a network configuration under consideration is shown in Figure 20. Data is passed on the ring to and from the dispatcher in packets that can be 2000 words in length. The bus could be designed to handle $\sim 10^6$ words/second. Each packet leaving the dispatcher would contain destination and procedure information that would tell the processor when to grab a packet from the ring and what to do with it. The dispatching function, NEXT STEP (processor, procedure) = F(status, processor availability), can be contained in a random access memory lookup table. The ring structure is far simpler conceptually than a complete crossbar interconnect system that would allow the processors to talk simultaneously to each other. It is, however, obviously subject to buss saturation limitations. This is mitigated by the fact that individual procedures will take a long time on the scale of data packet transmission time on the bus. To some degree additional rings, sub-subroutines with dispatchers on the main ring, can be used to control the bus traffic particularly in tight loop situations. Note also that the dispatching processor itself, if made sophisticated, is a procedure that could be handled by several processors on an auxiliary ring.

If simulation studies indicate the bus saturation problem intractable in linear bussing systems, one can turn to a "telephone company" data-switching network also controlled from a single dispatcher. Here the dispatching processor is the most likely area for saturation problems which can probably be handled with use of tables, as noted above, and high speed search hardware.

What makes these approaches appear so promising is the availability of powerful microprocessors such as the Motorola MC68000 or the new Intel APX-432 32 bit units, one of which we would probably use. The APX-432 architecture is specifically designed for network applications and its instruction set is matched for very efficient use of the high level structured language ADA. The new structured languages are more expressive than Fortran and apparently not difficult to learn. We can expect some resistance by physicists to leaving their mother tongue, Fortran. But not far in the future, I suspect, an increasing amount of reconstruction crunching will be done in processor networks like the ones we are studying, and the language of the programs will carry one of the strange names, ADA or Pascal or Praxis or C.

Conclusion

The breadth and potential of the work going on at Fermilab in advanced electronics for experiments is clear from this review and requires no further emphasis. So I hope I will be forgiven if I conclude with a brief sermon. The main reason for all the successful activity in electronics is,

of course, the powerful benefit that can be brought to experiments with modern integrated circuits. There is another factor which shouldn't be forgotten. This area of work is an unusually rewarding, and entertaining, escape from the sometimes extreme pressures and challenges elsewhere in the high energy physics environment. There is nothing intrinsically wrong with this, but there are dangers some of which can already be sensed. There has been surprisingly little communication between workers in this area even within one laboratory. The result is a substantial overlap of effort. The relatively low cost of this activity, apparent from Table 1, means that the usual checks and balances of a large experiment tend to be less rigorously applied. This leaves much freedom for creativity and much responsibility. The responsibility is to see that the electronics not become an end in itself but a part of the common effort to do good physics.

Acknowledgement

I would like to thank the following people for their help and information they provided: S. Bracker, C. Brown, S. Conetti, T. Droege, H. Fenker, I. Gaines, D. Green, D. Harding, T. Jensen, A. Jonckheere, C. Kerns, B. Knapp, H. Kobrak, J. Martin, W. Sippach, H. Sticker, R. Sumner, F. Taylor, J. Walker, and S. White. I would also like to thank Marti Bennett for typing this manuscript in a short period of time.

References and Footnotes

1. Cordon R. Kerns, IEEE Trans. Nucl. Sci 26 , 239.
2. D. Berg et al , Proc. XX Intern. Conf. , Madison, 1980, P537.
3. Strictly speaking this is not a low level trigger as we have defined it because the system proceeds to conclusion of the "high level" trigger whether or not the "low level" trigger is satisfied. It is described here because of its conceptual relationship with the other triggers in this section.

A "fast abort" signal is generated in parallel with the low level trigger operations if a count of PWC2 PWC3 multiplicities does not fall within a prescribed range.
4. T.F. Droege, I. Gaines, K.J. Turner, IEEE Trans. Nucl. Sci. 25, 698.
5. Also making major contributions to the original design effort were R. Hance, B. Haynes, T. Soszynski and K. Treptow. Recent module designs (for E537) have been done by J. Stoffel and P. Kostarakis (Athens).
6. E. Barsotti, J.A. Appel, S. Bracker, M. Haldeman, R. Hance, B. Haynes, J. Maenpaa, T. Nash, T. Soszynski, K. Treptow, IEEE Trans. Nucl. Sci. 26 , 686 and ECL/CAMAC Trigger Processor System Documentation, Fermilab TM821 (1978)

7. H. Frisch, N. Giokaris, J. Green, C. Grosso-Pilcher, R. Lenski, H. Sanders, R. Sumner, IEEE Trans. Nucl. Sci. 27, 150.
R. Lenski, H. Sanders, R. Sumner loc. cit., 655.
8. S. Bracker, Some General Thoughts About Multi-processing and other private notes, 1981.
9. B. Knapp et al, Fermilab Proposal 627 and Addendum.

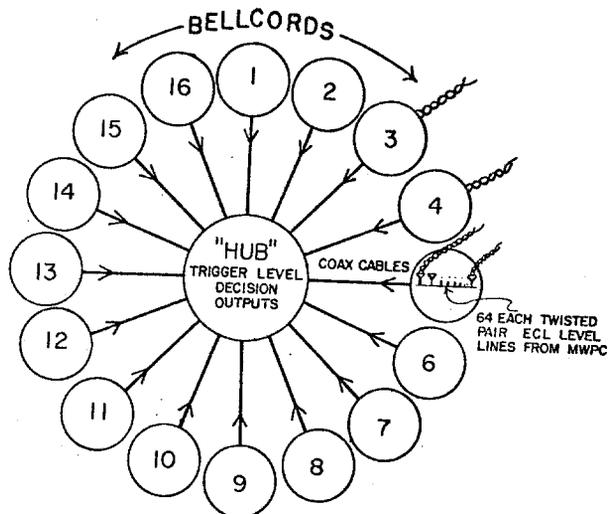


Figure 1: Bellcord system for E272.

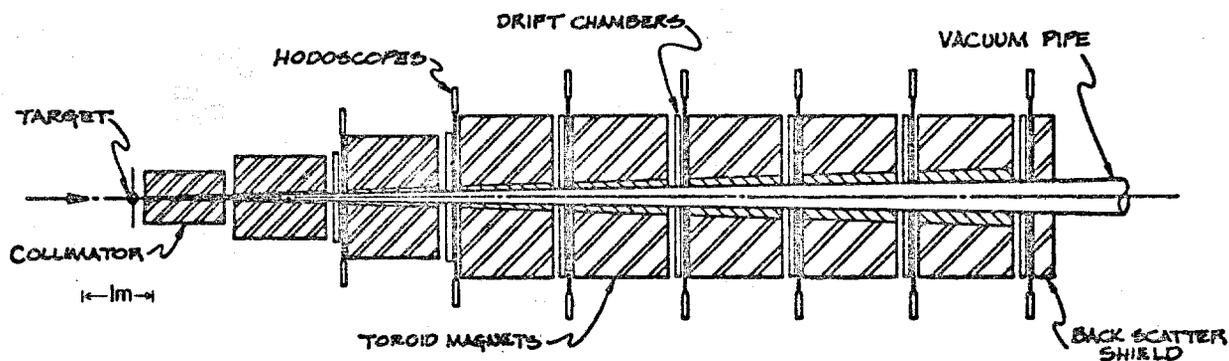


Figure 2: E326 dimuon magnetic toroid spectrometer in High Intensity Area.

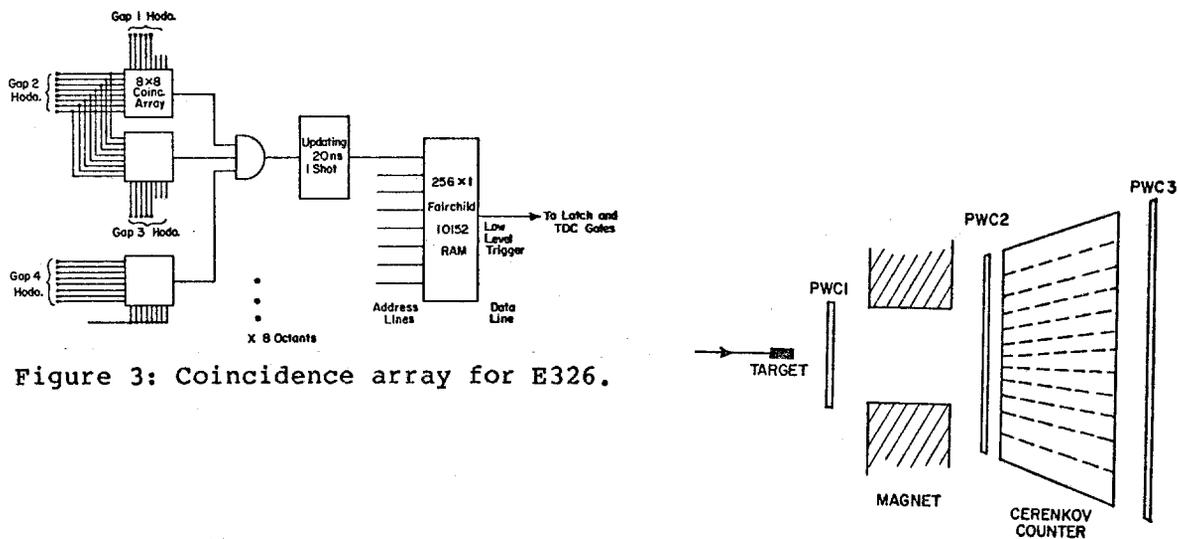


Figure 3: Coincidence array for E326.

Figure 4: E623, trigger view.

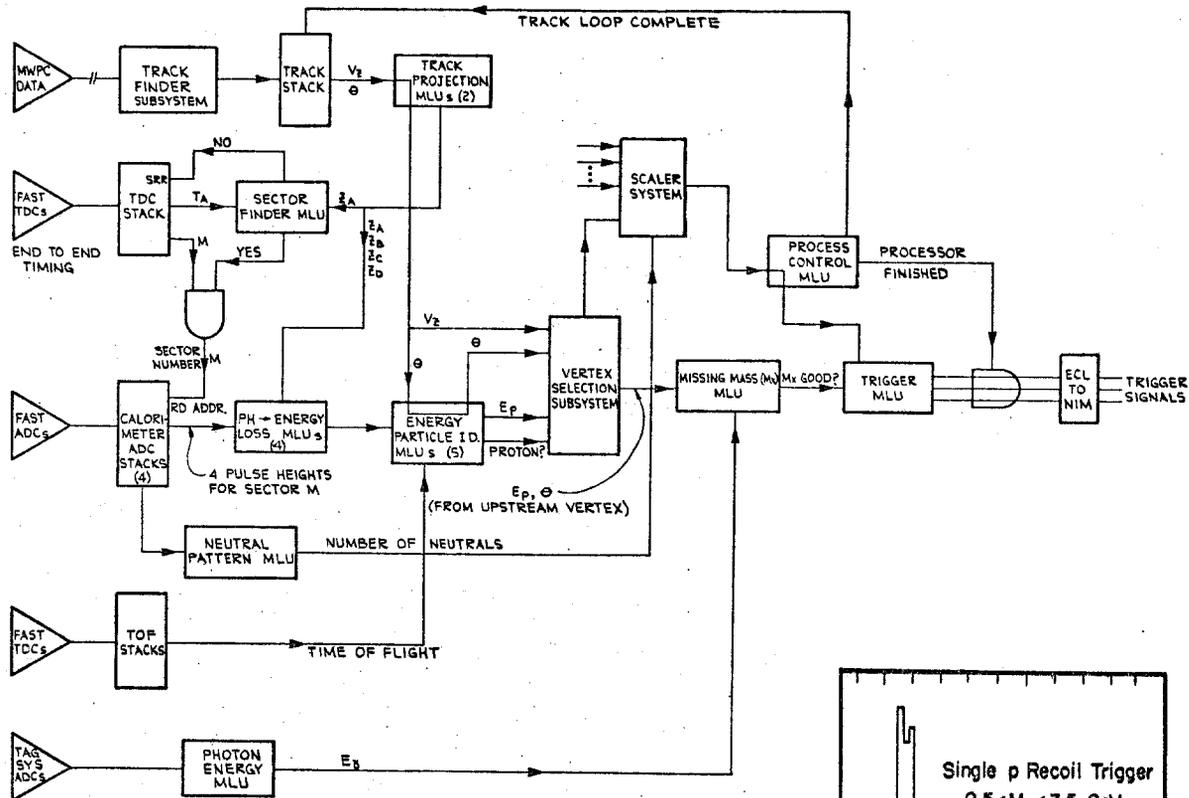


Figure 12: E516 recoil trigger, simplified schematic.

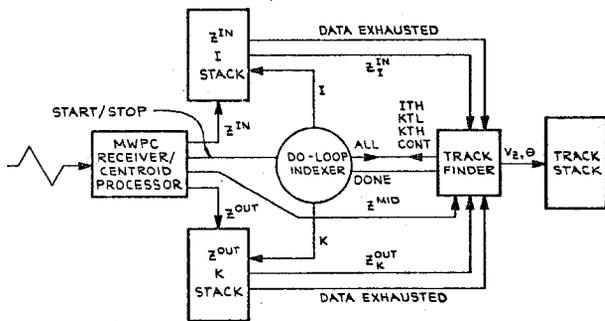


Figure 10: Track finder subsystem.

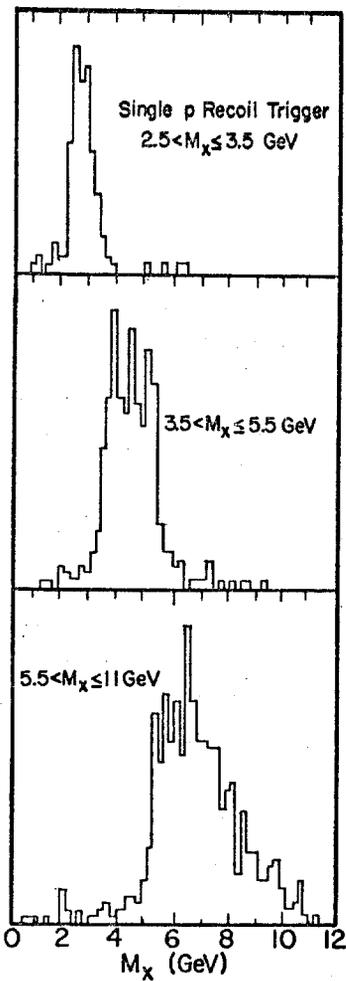


Figure 11.

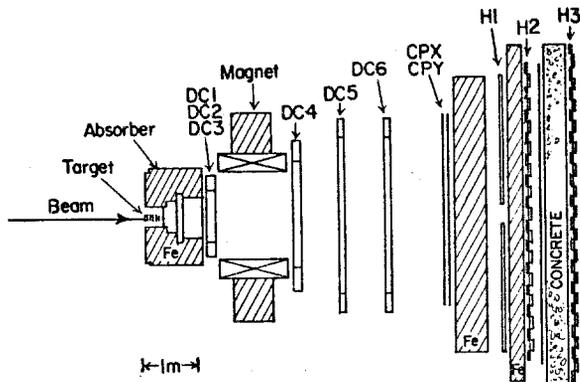


Figure 13: E537 π , \bar{p} spectrometer.

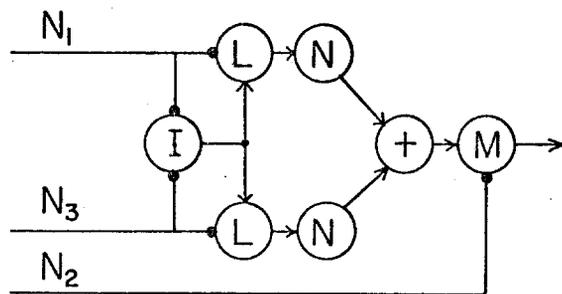


Figure 14: Three plane line finder.

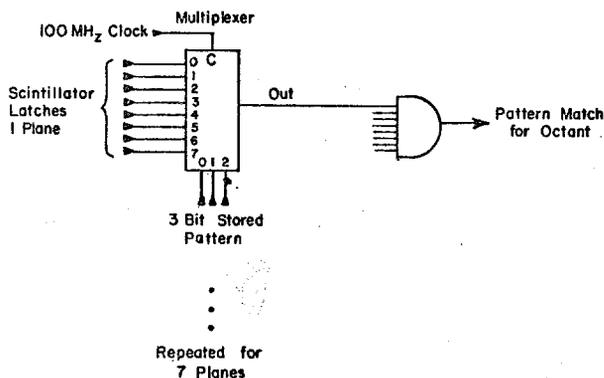


Figure 15: Pattern matching in E326.

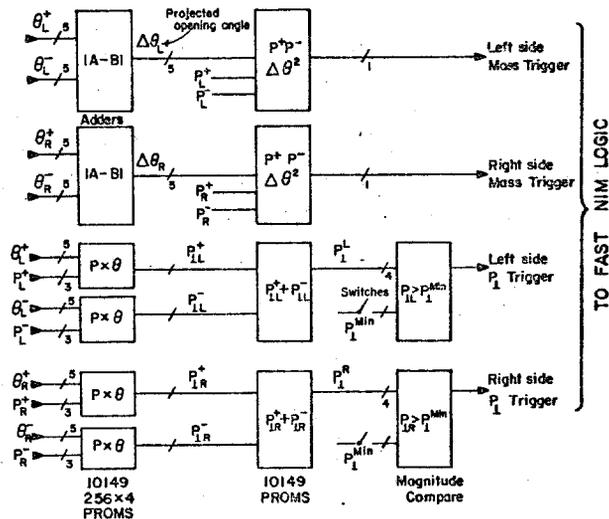


Figure 16: E623 P_T and M computation.

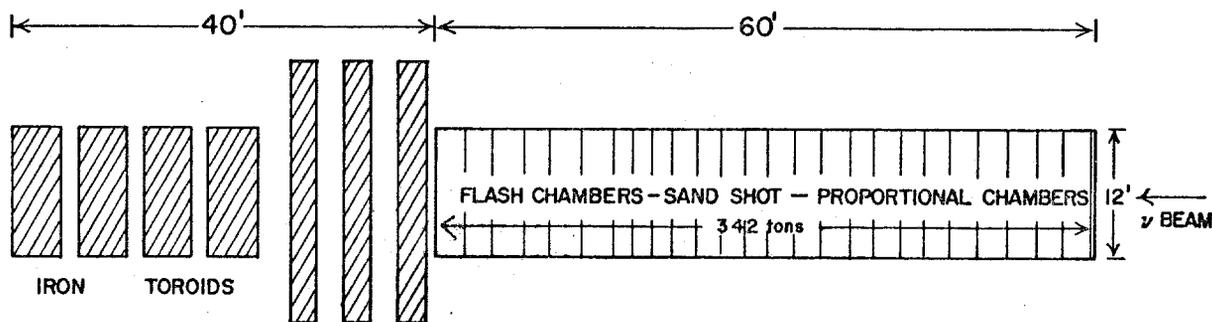


Figure 17: E594 elastic neutrino-electron scattering detector.

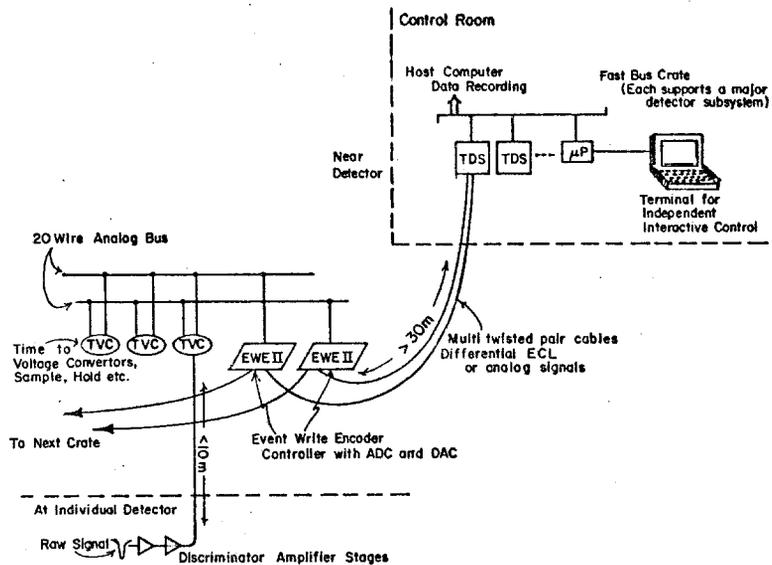


Figure 18: Plans for the Colliding Detector Facility data acquisition system.

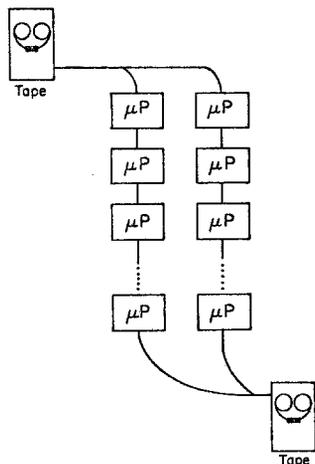


Figure 19: Serial multi processing system.

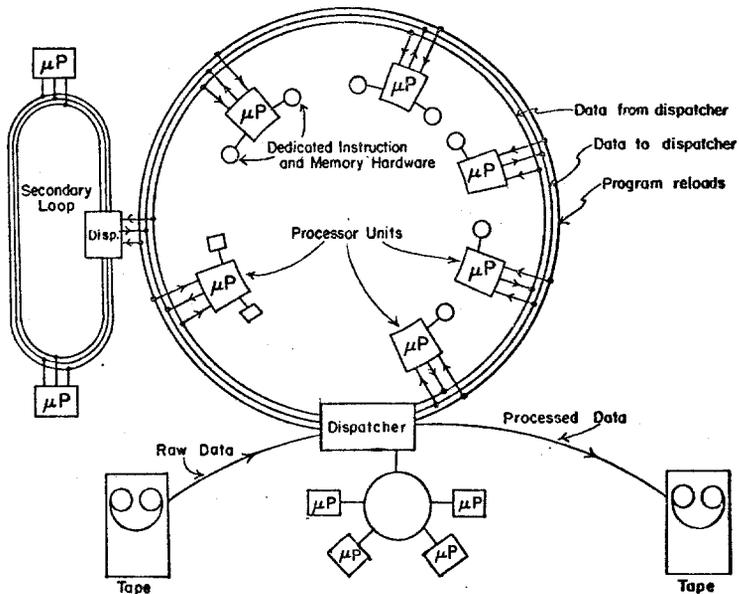


Figure 20: Reconstruction processor network.