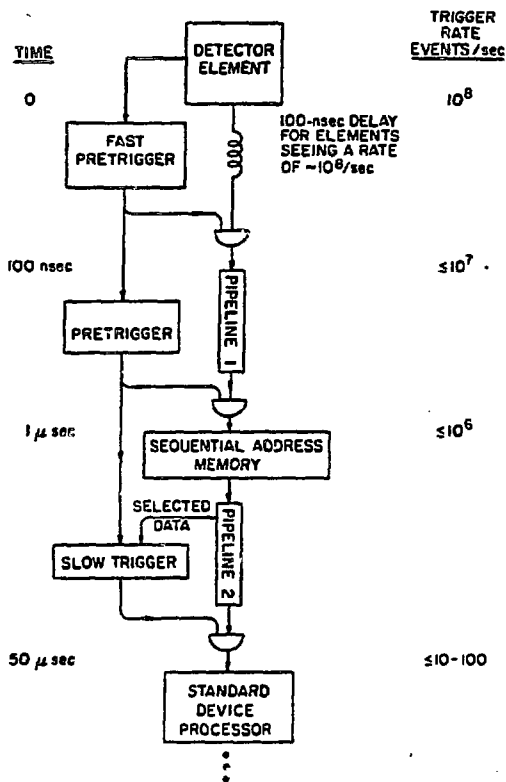


REPORT OF THE SUBGROUP ON FAST PROCESSING *

201 2000 B.G. Gibbard, Cornell University R.J. Plano, Rutgers University -5605000
 0373000 L.E. Kirsch, Brandeis University M.S.Z. Rabin, U. of Massachusetts -4003000
 6054000 G. Moneti, Syracuse University E. Willen, BNL -0936000

I. INTRODUCTION

We have examined the flow of data and the simultaneous processing needed to reduce the 10^7 to 10^8 triggers per second expected at ISABELLE to a number of events on the order of 10 to 100 per second which would be written on magnetic tape. We assumed that within 100 ns of the event, a fast pretrigger would have reduced the data rate to at most 10^7 per second. At that point, data from all sense elements in the experiment are fed into a 1- μ s-long pipeline (Fig. 1).



NOTICE
 This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

Fig. 1. Trigger and pipeline data flow.

*Supported in part by the United States Department of Energy.

Within the first 1 μ s (while the data are in the first pipeline) we assume that another level of triggering reduces the trigger rate to at most 10^6 per second. The data are then fed into a second pipeline which is 50 μ s long (Fig. 1). During the 50 μ s the data are in the second pipeline, a more sophisticated level of triggering (slow trigger) reduces the trigger rate to a level that can be handled by standard data processing techniques (microprocessors or larger machines), i.e., 10^2 to 10^3 per second.

First we describe the pipelines and the buffer between them, a sequential address memory, and then we present several alternative schemes for the pretrigger and slow trigger. We begin by assuming that a fast pretrigger is provided which will bring the rate in any high rate device down to 10^7 per second.

II. PIPELINE 1

Pipeline 1 must store for 1 μ s data coming in at the rate of up to 10^7 hits/s. We propose two devices, one for digital data (i.e., PWC wire hits) and another for analog data (i.e., calorimeter pulse heights).

For digital data, a shift register scheme similar to the one described by Platner¹ would be used. This device, which now costs \sim \\$1 to 2 per channel, samples data at the rate of 250 MHz and holds it for 1 μ s.

For analog data, a charge-coupled device (CCD) would be used.² Currently available CCDs (e.g., the Fairchild 321 A) can run reliably at rates up to 14 MHz (70 ns/cell). It is not unreasonable to assume, given the current rate of advances in technology, that in several years CCDs will be available with clock speeds of 50 MHz (20 ns/cell). Such CCDs would be perfectly suited to the task at hand. A 1- μ s storage time at 20 nsec/cell requires a CCD with 50 cells. It would have 10 events in the pipeline at any given time (on average). The cost of the Fairchild 321 A is now \sim \\$10 per channel, but it should come down as CCDs find wider usage.

The linearity of output voltage of CCDs with respect to input voltage is on the order of \sim 1%, but the nonlinearities can easily be corrected for by calibrating each chip.³ This would have to be monitored during experimental running. Many of the other problems mentioned by Zeller³ in 1977 (i.e., dark current and glitches) have been improved by Fairchild during the past year (e.g., dark current on the 321 A is 1.5 mV/msec rather than the 15 mV/sec reported in 1977).⁴

For those data lines with analog information where timing to better than 20 ns is necessary, an extra line of digital shift register would be used, with a 1 bit in the position signifying the time of arrival of the signal.

III. PIPELINE 2 AND TRANSITION BETWEEN PIPELINES 1 AND 2

At the end of 1 μ s, as events are coming out of pipeline 1, the pretrigger has reduced the valid data rate from \sim 1 event

every 100 ns to ~1 event every μs , i.e., on average 9 out of every 10 events are not passed to pipeline 2. However, because pipeline 2 is 50 μs long, ~50 events will be residing in pipeline 2 at any given time.

We propose that pipeline 2 be entirely CCDs of a type similar to the Fairchild 321 A, with about 500 cells and a clock rate of 10 MHz. Digital information from the shift register lines of pipeline 1 can be transmitted directly to the CCDs of pipeline 2. For those digital data lines where it is important to preserve the timing information, fast time-to-amplitude conversion could be performed on the output of pipeline 1.

The clock rate of the CCDs used in pipeline 2 will be 1/5 the clock rate of the CCDs in pipeline 1. This will cause a problem if acceptable triggers from the fast pretrigger are closer than every fifth event. Since the average space between acceptable triggers is 10 events, the probability of this problem occurring is $P(10,1) + P(10,2) + P(10,3) + P(10,4) + P(10,5) = 4 \times 10^{-4} + 0.002 + 0.008 + 0.019 + 0.038 = 0.07$, where $P(\bar{X}, n) = \bar{X}^n e^{-\bar{X}} / n!$ is the probability of the nth event after a good event also being a good event. This problem can be resolved in two ways: (1) Gate off the input to pipeline 2 after each good event has been passed to it long enough to avoid the pileup problem; this would cause loss of ~7% of the good events. (2) Use a serial analog memory, similar to Reticon SAM-64 (Fig. 2) but having a sample rate of 50 MHz, which allows analog information to be read into succeeding memory cells and read out of succeeding memory cells but has independent input and output address registers that are incremented with each operation, so that several analog events can be read in rapidly and read out at slower speed.

It is conceivable that at some time the slow trigger may wish to access some of the data in pipeline 2 (e.g., if it has determined transverse coordinates of a track by looking at which wires were




Fig. 2. Serial analog memory.

of a PWC or drift-chamber hit and would now like to get analog information on charge division to determine longitudinal coordinates for that track only.) If this is to be the case, pipeline 2 may be broken somewhere down its length and, with a buffer-amplifier, share some of its analog data with the slow trigger.

IV. A PRETRIGGER USING SHIFT REGISTERS

Here we propose a method of determining p_1 , φ , and θ for all tracks in an event within $1 \mu\text{s}$ and passing this information to the next stage in the processing. Thus, the slow trigger could calculate kinematic quantities such as effective masses, rapidity densities, sphericity, etc. To make the example definite, we choose a specific geometry, i.e., a cylindrical detector in a solenoidal magnetic field. With obvious modifications, the procedure can be adapted to different configurations. The basic idea is taken from Breidenback's fast track-finding scheme used in the SPEAR MK-II detector.⁵

We assume a cylindrical detector with external radius $R = 1.5 \text{ m}$, in a 10-kG axial magnetic field. Three cylindrical proportional wire gaps would be available for triggering, with radii R_1 , R_2 , and $R_3 = R$, i.e., $R_1 = 0.3 \text{ m}$, $R_2 = 0.8 \text{ m}$, and $R_3 = 1.5 \text{ m}$ (Fig. 3).




Fig. 3. Cylindrical detector.

The axial wire would provide the azimuth coordinate φ , and cathode strip readout would provide the axial coordinate z along the beam.

Let us divide (logically) each cylinder into 8 octants and each octant into about 100 azimuthal sectors. Each sector would cover a $\delta\varphi = 2\pi/800 = \pi/400$ and would be made of a suitable number of sense wires ORed together. The hits from each octant of each chamber would be parallel loaded in a 100-element fast (200-MHz) shift register. This would be done by a fast strobe after having allowed for the drift time in the PWC.

A coincidence among the three shift register outputs would indicate the presence of a straight track originating in the axis of the cylinder. The tune of the coincidence would effectively

provide the azimuth of the track. Fan out the shift register outputs to delays of 0 to $n\delta\phi$ (these delays would be variable-length shift registers) looking for coincidences among the three chambers: $n_1\delta\phi$, $n_2\delta\phi$, $n_3\delta\phi$ (Fig. 4). As a constraint, only

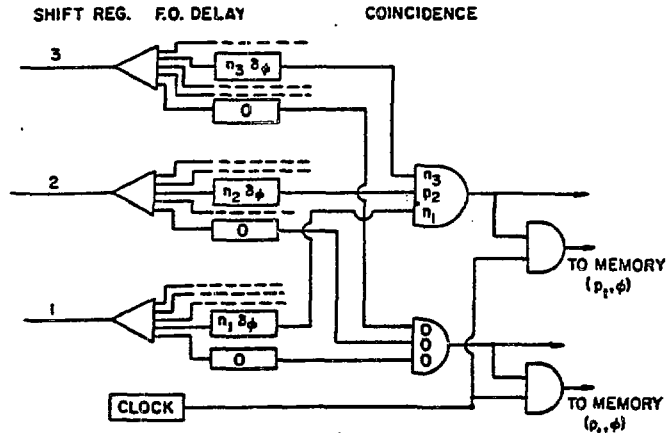


Fig. 4. Cylindrical detector logic.

coincidences corresponding to tracks originating from the beam line would be allowed. A given n_1 , n_2 , n_3 coincidence, combined with a clock measurement, will give unambiguous values of p_1 , ϕ , and charge sign of track. The actual values may be obtained from a lookup table in memory.

Resolution

The projected radius of curvature is

$$\rho = \frac{R^2}{8s}, \quad s = \frac{R}{2} (\varphi_3 - \varphi_2),$$

hence

$$\rho = \frac{R}{4(\varphi_3 - \varphi_2)}, \quad \frac{\delta p_1}{p_1} = \frac{\delta \rho}{\rho} = \frac{4\sqrt{2}\rho}{R} \delta \varphi;$$

with $\delta \rho = \pi/400$, $R = 1.5$ m, and $B = 10$ kG:

$$\frac{\delta p_1}{p_1} = \frac{\rho}{34 \text{ m}} = \frac{p_1}{10 \text{ GeV}/c},$$

i.e., undelayed coincidence means $p_1 > 10 \text{ GeV}/c$.

Coincidences with $ \varphi_3 - \varphi_2 = \delta\varphi$ mean $5.0 < p_1 < 10.0$ GeV/c		
	2 $\delta\varphi$	3.3
	3 $\delta\varphi$	2.5
	4 $\delta\varphi$	2.0
		5.0
		3.3
		2.5

and so on. Make your choice of how many coincidences per octant you want.

Boundary Between Octants

The end of a shift register should overlap the beginning of the next one for a number of elements corresponding to the maximum delay used in the coincidences.

Possible duplication of tracks can be sorted out on the spot by suitable (anti)correlation logic between octants or later by comparing contents of memory.

Detector Efficiency

To avoid inefficiencies, two sense planes should be provided at each R_i . Signals of the two planes may be ORed into shift registers, or two shift registers should be used and their outputs ORed before the triple coincidences.

Axial View

In the axial view (Fig. 5) the intersection region, given its length, does not supply an adequate constraint for a two-point track. The tracks are then required to have three hits on a straight line. The straight-line assumption is good insofar as $p > R$, i.e., in 10 kG with $R = 1.5$ m and $p_1 > 0.45$ GeV/c. To accommodate lower p_1 , there should be three sense cylinders within a smaller radius.




Fig. 5. Axial view of intersection region.

Suppose we have cathode strip readout allowing location of the hit (without center-of-charge calculation) to within ~ 1.5 cm in z . This would mean an angular resolution of 1.4° (25 mrad)

for a $\theta = 90^\circ$ track and of 0.7° (12 mrad) for a 45° track. It seems an adequate resolution for our purpose.

Let us now divide the cylinders into "forward" and "backward" halves and assign a parallel-load shift register to each half cylinder with a number of elements such that each element gets the signal from a $\Delta z = 1.5\text{-cm}$ (i.e., three strips). Since, however, the shift registers have clock rates proportional to the number of elements in them, they all take the same time to shift their entire length (Fig. 6).




Fig. 6. Shift register assignment.

Requiring a triple coincidence between the three outputs of the shift registers is equivalent to requiring a straight track through the origin. By taking coincidences with delayed shift register outputs, one can allow for tracks that do not go through the origin. The scheme is identical to that of the azimuthal view, including taking into account the boundary of the two half cylinders.

There is a one-to-one correspondence between coincidence gates and θ and z (production) of the track. Again, the actual values can be given by a lookup table (memory).

To increase the ability of the system to deal with high multiplicity events, it may be necessary to break the cathode strips into azimuthal sectors with a corresponding increase in the number of fast track-segment recognizing processors. The azimuthal segmentation of the axial view would also facilitate the association of track segments in the two views, which is discussed next.

Association of Track Segments in the Azimuthal and Axial Views

Azimuthal segmentation will provide, in about $0.6 \mu\text{s}$, two separate sets of two-parameter track projections: a set of $\{p_i, \varphi_i\}$ and a set of $\{\theta_i, z_{pi}\}$. The next task is to correlate them and generate a set of four-parameter tracks $\{p_{it}, \varphi_i, \theta_i, z_{pi}\}$.

In a PVC-cathode strip readout scheme, a possible handle is provided by the fact that cathode and anode pulses for the same hit are strictly proportional, but there is a wide variation, due to ionization fluctuation, in pulse heights between different hits (the FWHM of the Landau curve is about 100%). To take

advantage of this information, one could perform a rough 3-bit digitization of every group of anode wires or group of cathode strip signals. This requires seven comparators and seven AND gates. (A simpler and cheaper 2-bit digitization might do the job if the number of tracks per sector is small.) Each group of sense elements would load three parallel shift registers. The ORs of the three would be used in the same way as the single shift register in the preceding sections. However, for a successful projected track element, the composite 9-bit word (3 bits from each of the 3 sense planes) would be saved in a $N \times 9$ -bit word register (N being the maximum number of projected track segments expected in a sector).

Besides the 9 bits containing the pulse-height pattern (PHP) of 3 bits, these words should contain additional address bits (IDW) identifying the track segment, $\{p_{1i}, \varphi_i\}$ or $\{\theta_j, z_{pj}\}$, to which they belong. Each PHP of the axial view would then be compared with those of the azimuthal view, and the best matches (i, j) would be considered as full track candidates with parameters $\{p_{1i}, \varphi_i, \theta_j, z_{pj}\}$.

Storage and Retrieval of the Track Parameters

Each triple coincidence among shift registers should have access to a register containing an identifying word (IDW) < 14 bits. The content of this register and the content of the PHP word must be transferred to the N -track register within a shift register clock period (5 ns in the current example) in order for the coincidence to be ready for another clock pulse (suppose there are two track with the same p_{\perp} in adjacent φ -bins). The transfer operation can, however, be pipelined.

After the track projections have been correlated, two IDWs from the two views merge into one track identification word (TIDW). This can be used to retrieve the actual physical parameters $\{p_{\perp}, \varphi_i, \theta_j, z_{pj}\}$.

In the example under consideration, one should allow the following number of values for each parameter.

Parameter		No. of Elements	No. of Bits
p_{\perp}	8 intervals \times 2 signs of charge =	16	4
φ	100 intervals \times 8 octants =	800	10
θ	100 intervals \times 2 halves =	200	8
z_p	30 intervals	30	5
P	Total	1046	27

Note that it is possible (indeed natural) to encode the address for each parameter independently so that all we need is four (logically) separate read-only memories for a total of about 1046 words in order to translate the IDW into a physical parameter word (PPW) containing $(p_{\perp}, \varphi, \theta, z_p)$. Instead of storing $\{p_{\perp}, \varphi, \theta, z_p\}$ one could equivalently store $\{p_{\perp}, (\sin \varphi, \cos \varphi), \sin \theta,$

z_p]. Translating the TIDWs into parameter words, done serially for all the tracks found, may take 20 ns per track.

Conclusion

With the scheme described it is possible, in about 1 μ s, to recognize most of the tracks in a detector and to make their physical parameters (\bar{p} , z_p) available for further processing. The scheme should be adaptable to geometries substantially different from the one used as an example here. The cost seems to be relatively modest, and, what is more important, using more detectors to introduce either constraints or redundancy will increase the cost of the electronics about linearly rather than according to a power law.

V. AN ALTERNATIVE SHIFT REGISTER SCHEME

An alternative shift register scheme that takes more time than the one just outlined, but is more general, is described below. Eight trigger planes are assumed, with charge-division readout on alternate planes. This scheme uses both shift registers to prepare the data and standard RAMs to determine p , θ , and φ . It is seen as a generalization of the SLAC-LBL MK-II scheme⁶ and the one described above.

Suppose we have eight planes of wire measuring the coordinates in a bending plane. The bend is accomplished by a discrete field region(s) between detector planes or by a distributed field which, though not necessarily uniform, should not vary too drastically. Every second plane (indicated by CD on Fig. 7) will have charge division done on it, and a fast (≤ 10 -nsec) crude (3-bit) ADC result for each end of these wires must be available.

The results from each wire, 6 bits for the current-divided wires and 1 bit for the others, are parallel loaded into a shift register chain containing all the elements of a particular plane in series. The data, as shown in Fig. 8, are then shifted out of these registers at rates controlled by a clock program RAM. The current-division planes must additionally pass through a pulse-height to position conversion RAM. Thus, on each master clock, 16 bits of data, 3 bits from each CD plane and 1 bit from each non-CD plane, are presented to a set of parallel processors, each of which scans for a particular momentum range. Each processor, of which there might be as many as 40 to produce adequate momentum resolution, would consist of one $64,000 \times 16$ -track parameter coincidence RAM (Fig. 7) and eight sets of plane preparation logic (Fig. 8), one for each of the eight planes. The plane preparation logic introduces a delay in the shift register chain for each particular plane which offsets the expected curvature of tracks of a particular momentum. In addition, it determines the width of the window viewed at one time along the shift register chain for each plane (effectively the road width). By separately controlling the shifting rates of each plane and by controlling for each

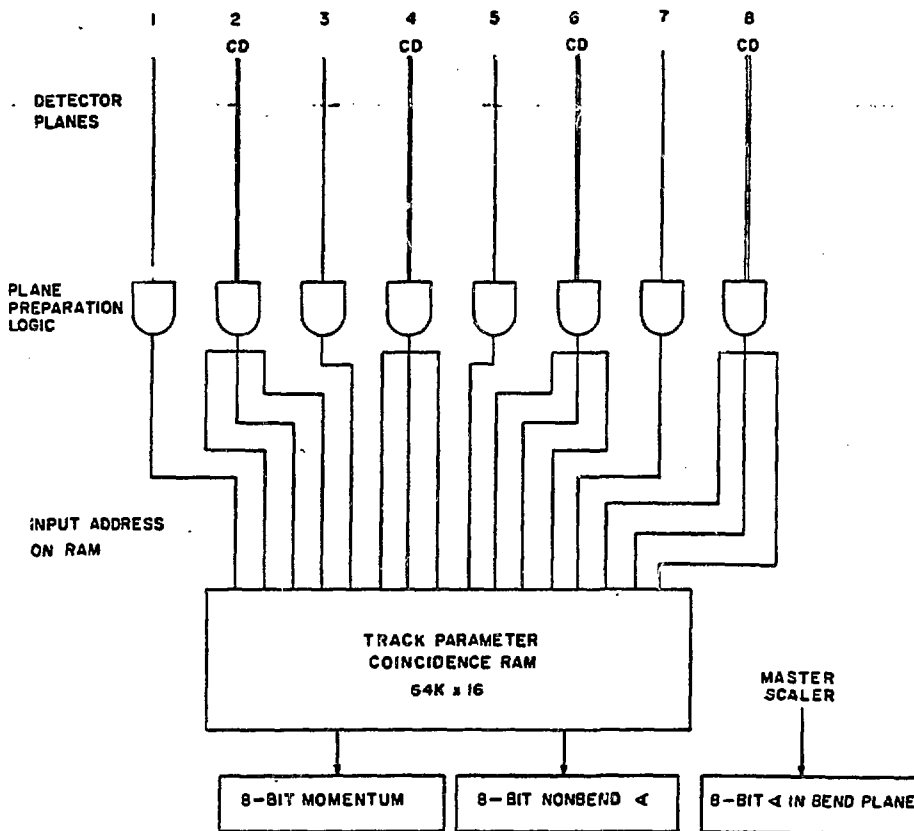


Fig. 7. Overall processor schematic.

processor, independently and dynamically, the delay and road width of each plane, it should be possible to guarantee that, as long as the detector geometry and magnetic fields are not pathological, all tracks will be detected by some processor.

The detection of a track by a processor is recognized by the observation that the current set of 16 bits constitutes an address for which the content of the RAM is not zero. The RAM is preloaded with the momentum and nonbending plane angle best characterized by the particular pattern of bits (some representing wires hit and others representing positions in the current-division direction) constituting a valid track. By the proper combination of plane shifting and individual processor delays for these planes, it is possible to establish the master clock count as a simple function

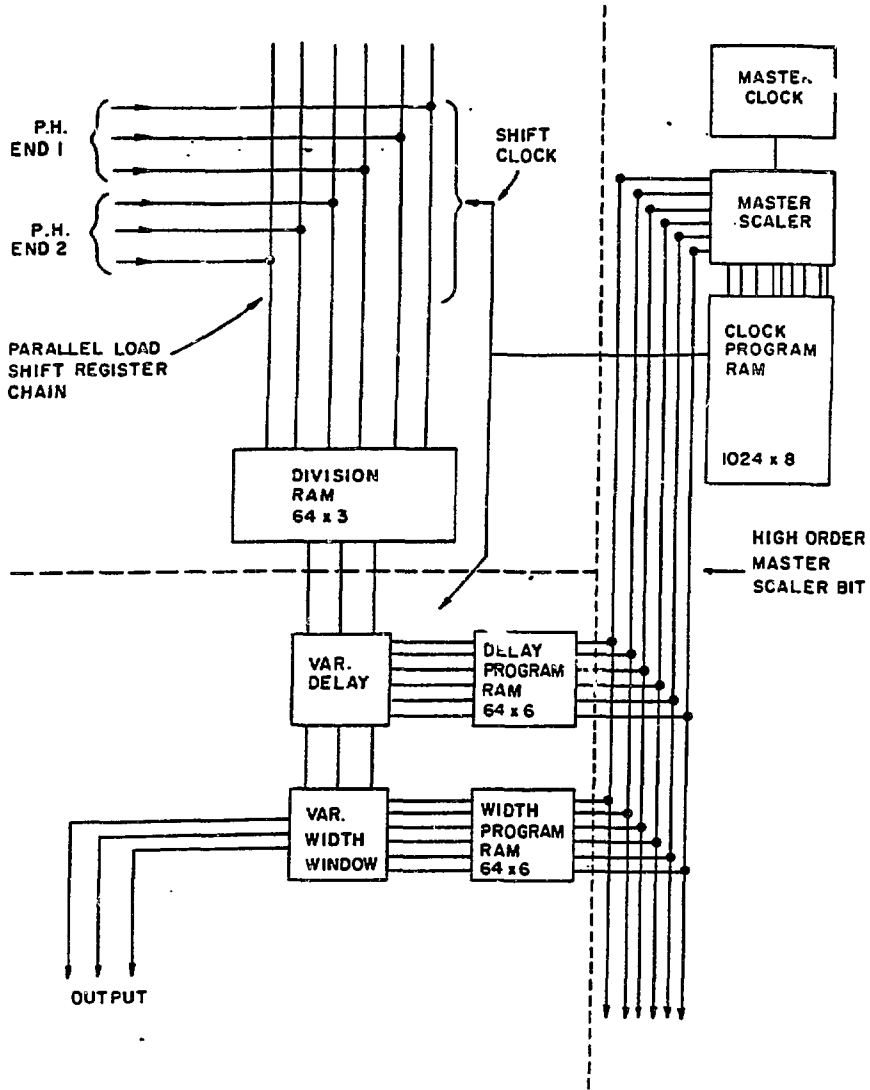


Fig. 8. Typical current-division plane preparation logic.

of the production angle in the bending plane. Admittedly, the effort required to determine the contents to be loaded into the parameter and control RAM would be substantial. The amount of very fast memory required is also substantial, but technological advances in this area are expected in the next few years.

Although the accuracy obtained by this system depends completely on the particular geometry and fields used, it seems

possible for a detector of somewhat limited though large aperture to obtain accuracies in angles of 10 to 20 mrad and momentum accuracies of 6% for momenta < 20 GeV and $\Delta p/p^2 = 0.003$ for higher momenta.

If one considers planes containing 1000 to 2000 wires and master clock rates of 100 MHz, the track parameters can be completely determined in 10 to 30 μ s. Microprocessors operating on these parameters could begin higher-level processing as soon as the first parameters were available. By this means, a very high-level kinematic decision could be made within 50 μ s.

VI. TRIGGERING UTILIZING SPECIAL RAMS

This approach to triggering is based on special RAMs similar to that devised by Platner.⁶ Approaches to the problem of a fast trigger processor are either to devise an algorithm unique to the trigger which can be implemented with high efficiency in hardware, or to build a general-purpose device at some sacrifice in performance. The advantages of the latter are its ease of modification in the event of a changed trigger and the concurrent reduction in money and manpower. Here we consider the performance that might be achieved with such a device.

Platner's RAM is one that has no address encoding. An $n \times m$ RAM thus has $n + m$ input lines, each pair (x, y) corresponding to a bit. If any bit in memory corresponding to a combination (x, y) of 1's input contains a 1, a 1 bit is output!

This system, which we call BITRAM, provides a one-bit answer to a three-dimensional correlation function with a granularity of 128 elements in each dimension (2×10^8 bits). It is completely general in terms of the variables, which may be wire numbers, hodoscope elements, or digitized data from calorimeters, and it is easily modified for different experimental configurations. The generalization of the BITRAM is a device which provides the address(es) of the correlation(s) rather than just the information that one or n correlations exist (Fig. 9). We call this system a WORDRAM. Decisions would be available from this device in 5 μ s depending on the track multiplicity and degree of parallelism of the processor.

As one example of its use, we consider a set of concentric PWCs with axial wires and charge-division readout, placed in an axial magnetic field. As indicated in Fig. 9, these chambers are connected to the BITRAM and 10-event buffer which, in turn, feeds the WORDRAM on average every tenth event. The indices (addresses) produced by this device correspond to a set $\{p_{\perp}, \varphi\}$. A stream of analog information from the charge-division readout would be stored in a CCD shift register for 10 μ s while the WORDRAM determines the set of wires that are correlated. These indices are used to select the analog data for the particular wires used in a correlation, and hence the longitudinal position and angle are determined. The composite addresses now point to a $\{p, \varphi, \theta\}$ table which allows calculation of the effective mass of all pairs.

The time required at each stage is shown on the left of Fig. 9, as is the reduction in data rate.

The problem which is difficult to solve is that, in general, several correlations exist in the WORDRAM and therefore they must

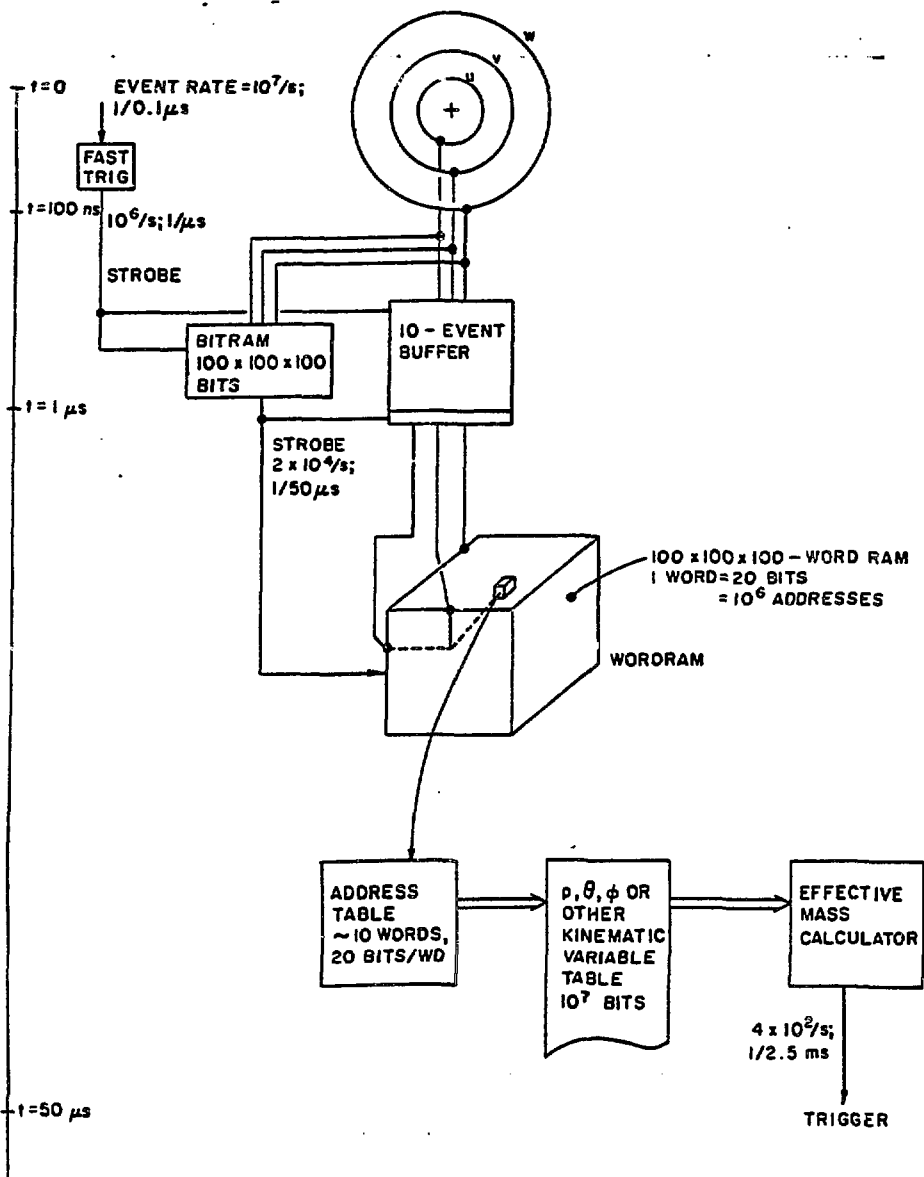


Fig. 9. Concentric PWCs in an axial magnetic field. Assumption: Interaction point and three measurements in PWCs in B-field give unique p , θ , ϕ .

be serially multiplexed onto the address lines. This differs from the situation in the BITRAM, which produces only a one-bit answer. The question thus becomes how to scan the sparse matrix efficiently to find the valid correlations, subject to the constraint that the integrated circuit package have some limited number of output lines ≤ 50 . This requires an internal (on chip) encoding of the address and internal serialization of multiple correlations. A schematic design of such a chip is shown in Fig. 10. The WORDRAM would be

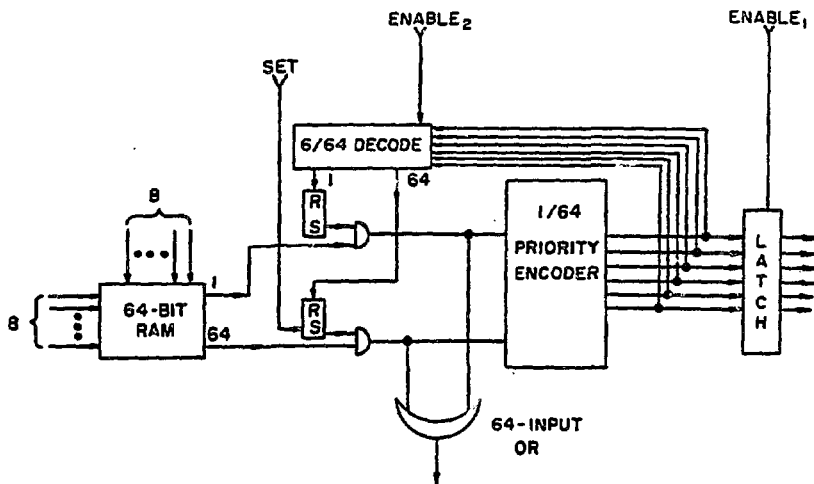


Fig. 10. Schematic design of a self-serializing RAM. A total of 26 pins are required, exclusive of power requirements.

constructed of 256 chips/plane, and the scanning of an entire plane would require $256 + n$ clock times where n is the number of correlations (~ 10). This could occur in $\sim 2 \mu\text{s}$ if all planes are scanned in parallel.

Obviously there are a variety of serialization schemes, but, no matter which approach is taken, a binary division search of 16,000 bits takes 14 probes. Hence, for 10 tracks, 140 probes are required. Thus the scheme is limited to time domains in the range of 10 to 50 μs when the time for the kinematics processor is included.

With all the schemes outlined above it is assumed that track information is passed on to another device which then calculates quantities of interest, such as an effective mass or jet parameters, in the remaining time. This should allow the trigger rate to be reduced to a manageable level within 50 μs .

REFERENCES

1. E.D. Platner, in ISABELLE 1977 Summer Workshop, p. 53, BNL 50721.
2. M.E. Zeller, Ibid., p. 140.
3. F. Kirsten and E. Yazgar, Characterization of Charge Coupled Analog Memories for Nuclear Data Acquisition, LBL Note; Nucl. Instrum. Methods, in press.
4. M.E. Zeller, Private communication.
5. H. Brafmann et al., Fast Track Finding Trigger Processor for the SLAC-LBL MK-II Detector, SLAC-PUB 2033, Oct. 1977.
6. E.D. Platner, Programmable Combinational Logic Trigger System for High Energy Particle Physics Experiments, BNL 22020, 1976.