FERMILAB-PROPOSAL-0879

June 12, 1995

TRACKING ARRANGEMENT for a B COLLIDER TEST SETUP

J. Butler *, D. Christian
D. Husby, M. Johnson, L. Spiegel
Fermilab

M. Procario, J. Russ
Carnegie Mellon University

W. Kononenko, W. Selove *, K. Sterner
University of Pennsylvania

M. Artuso, M. Goldberg, S. Stone
Syracuse University

Abstract

An advanced trigger for selecting beauty events at the Tevatron collider is being designed. This trigger aims to detect the presence of detached secondary vertices. High speed tracking is required, to do this at high luminosity. If a secondary-vertex trigger is to be made the lowest level trigger ("Level-1"), then very high speed tracking is required in the vertex detector. This vertex detector tracking must either give momentum information, on each track, or must be combined with other elements to give momentum information, at primary trigger level.

In this proposal we describe a testbeam setup which would provide tests of several ideas aimed at multi-MHz collider operation with a secondary-vertex trigger as a primary trigger. The initial objective is the testing of ideas on very high speed tracking, and the study of the rate of false secondary-vertex events. The testbeam setup will be usable for subsequent testing of additional components needed for a high luminosity secondary-vertex-trigger B-detector system, including high speed fiber-optics data transport, a prototype high speed data-switch/trigger system, pixel detectors to assist in very high speed tracking, and particle ID elements.

* Co-Spokesmen

1 Introduction

The two crucial characteristics of a secondary-vertex trigger, as for any trigger, are rejection and efficiency. We have designed a trigger system for producing a secondary-vertex trigger, with the ability to work at high luminosity. Our approach is aimed at achieving high rejection and high efficiency, for beauty events, by locating vertices with high precision. The capability of operating at very high luminosity is achieved by the use of massively parallel organization of the tracking and triggering system.

We have carried out extensive Monte Carlo calculations, studying the predicted performance in detail. These calculations have given promising results. Such calculations however can not by themselves provide a solid foundation for committing to construction of a large detector. The test setup described here will enable us to test the actual performance of the new system we have designed for tracking and triggering, with regard to precision of the vertex finding, adaptation to very highly parallel processing in the tracking/triggering system, and susceptibility to confusion from false vertices. The setup will also provide the capability for testing, subsequently, a set of new hardware elements needed in the final system.

The initial phase of the tests will be carried out at low intensity. In later phases, we will operate at higher intensity. At a still later time, and based on the performance in the fixed-target tests, we expect to propose a test run in the collider, under conditions closely approximating those which would obtain in an actual high luminosity experiment using the secondary-vertex trigger.

2 Phase I: Tracking Tests at Low Intensity

In Phase I of the proposed program a small vertex detector system will be operated, in or ahead of a dipole magnetic field. A test beam of order 200 GeV/c, mostly pions, will be used. Figure 1 shows the geometry, when the detector is located in the field.

Operating at an interaction rate of about 1 kHz, this setup will provide:

1. tests of a tracking arrangement designed to allow massively parallel real-time tracking at high luminosity, with momentum information, of virtually all tracks seen in a vertex detector,

2. measurements of vertexing precision in an arrangement designed to give unusually high vertexing precision in a forward-region collider geometry,

3. experimental information, from 10^6 to 10^7 events recorded per day. This will allow

    - comparison with Monte Carlo results on vertex precision,

- observation of whatever effects might give false secondary vertices at the level of 1 in 10^5 to 10^6 events.

More detail is given in Section 4.

Additional tests which we plan to carry out, in Phase I, include 1) tests with a different geometry, in which the vertex detector is moved upstream out of the magnetic field, and 2) tests of particle ID systems appropriate to a B detector. A specific example of the latter will be a study of RICH modules, to be carried out by the Syracuse group. That group is currently constructing a RICH system for the CLEO III experiment, which will use a LiF radiator and a gaseous methane-TEA photon detector. Some of these modules can be used, with various radiators and various photon detectors, to explore the performance available for different angular regions (different momentum regions) in a hadron-collider B detector.

3 Phase II: Tests at High Interaction Rates

We anticipate that before the 1996-97 Fixed Target run is over, useful measurements at high interaction rates can be made, in this setup, with system components needed for high-luminosity operation. We expect that a prototype high-speed fiber-optic data transport system will be ready for testing [1]. We also expect that prototype components of a high-speed data-switch/trigger system will be available for initial testing [2]. Finally, for very high speed real-time tracking, pixel planes would be of great advantage; if such planes are available with suitable readout electronics [3], we would test them, also, under high-rate conditions.

4 The Test Vertex Detector

The vertex detector shown in Figure 1 has a set of features which in a later complete detector would make possible high speed on-line tracking of collider events at luminosity 10^{32} or higher. (The readout electronics used in Phase I of this program will not have that speed capability.)

For fully on-line tracking of all tracks in a vertex detector, at an average rate of 1 interaction per 132 nsec bunch crossing period, a massively parallel computing system is needed. Moreover, for efficient detection of tracks with significant impact parameters, the momentum of each track must be measured on line, so that tracks of low p_t, which have large false impact parameters, will be ignored in the initial vertexing search. A system for data transfer and trigger processing has been designed by Selove and Sterner, and is described in references [4] and [5].

A full secondary-vertex trigger system has been designed by D. Husby and W. Selove, and a prototype module is being built at Fermilab [6].

Two particularly important matters which will be addressed in the test program are the following:

1. In an ultimate detector system, which will use pixels, the signal/noise ratio will be much higher than is obtainable with long strips. For a number of questions, high S/N operation will give results quite different from those for more usual S/N values for long strips. These questions include the matter of efficiency for angled tracks, and the matter of confusion caused by relatively weak signals, from delta rays or from radiation background. To explore these questions, we will include some special planes, with short strips, of order 1.0 cm long, to better simulate the performance of pixels.

2. We will be using a dipole field at the vertex detector, to test the effectiveness of measuring the momentum of individual tracks purely from the vertex detector hits. We note that for a forward-geometry vertex detector a dipole field is much better suited than a solenoidal field, for obtaining long track length in the vertex detector.

The test detector of Figure 1 has the following features, patterned after the design of an eventual high-rate system, and designed to make a very fast tracking analysis possible.

1. Each station consists of a triplet of planes. We expect that ultimately one or more of these planes will be pixel planes; in the test setup they will be microstrip planes. The 3 members of a station give, successively, precise values of y, x, and y. (Track-matching ambiguities will be resolved either by stereo strip planes or by one or more pixel planes.)

2. The first stage of the hit processing will consist of a mini-tracking operation, finding the triple coincidence of hits on a track in the 3 planes. In the test beam setup the analysis will be done off-line; we go on to describe how the ultimate on-line system is designed to operate.

   Ultimately each station of the detector will be subdivided into mini-areas. The mini-areas are of such a size that the occupancy is very low, of order 0.5 tracks per interaction. This low occupancy makes it possible for separate hit-processors to operate on the hits from each mini-area of each station, all in parallel. The triplet of hits, in a mini-area of an individual station, is converted by the hit processor into a 3-dimensional space point.

   The subdivision into mini-areas, each one a "phi-slice", is illustrated in Figure 2. Figure 3 shows the typical location of a hit, taken to be defined 2-dimensionally in a pixel plane, in one plane of one phi-slice of a station. Finally Figure 4 shows the typical location of the triplet of hits in a station. With the low occupancy,

and with the very good signal-to-noise ratio for pixel planes, the mini-tracking operation promises to be feasible with very little complication from amplifier noise or from background radiation.

3. The 3-D space points from the hit processors will be directed, by data switches, into a set of phi-slice tracking processors. We have carried out initial programming for the tracking processor, and obtain a program time of less than 10 microseconds, average, to find the track parameters from the set of hits on a track, This is for a processor running at 100 MHz. We expect this time to be appreciably shortened as the program is refined. For a 10 microsecond average tracking time, and about 200 tracks per microsecond, the number of tracking processors required to just keep up is 2,000.

Prototype hardware is being built at Fermilab by D. Husby *et al.*, in the Electronic Systems Engineering Department. A layout of the trigger system is shown in Figure 5. Further details are contained in Ref. [6].

We now mention a few additional features of this design, which speed up the tracking and which suppress noise hits and poor quality tracks.

1. The mini-tracking eliminates isolated single hits from noise and from radiation background.

2. The triplet of hits, in a mini-area, gives a track slope as well as a hit position. This slope is used in several ways:

   - The sign of the slope makes it possible to steer hits separately to different track processor modules when one track is a "forward" track and the other is a "rearward" track, even if the phi values are in the same phi interval. This separation reduces further the track occupancy in the individual track processors.
   - When there are multiple tracks in a phi-slice, they can generally be sorted, within the track processor, into sets of different slopes. To a good approximation, each set corresponds to a single track.
   - The slope, used together with the transverse coordinate, immediately gives an approximate z-vertex location along the beam axis, typically accurate to a few hundred microns for individual tracks.
   - With a dipole field, the slope in the bend plane changes from station to station in a smooth way, giving a very rapid estimate of the track momentum. If the slope change is not smooth, the track is of poor quality, and should either be discarded or else analyzed more carefully. With a 10 kGauss field, and the station spacing shown in Figure 1, track momentum can be measured in this test setup with useful accuracy for momentum up to 20 or 30 GeV.

5 Equipment and Funds

In the initial test system, all silicon planes will be strip planes. The readout electronics will be closely duplicated from that of E781. Readout will be done with SVX-2 chips, followed by copies of the E781 digitizer modules. The SVX readout system is also very well adapted to later testing of a fiber optics data transmission system: the development of such a fiber optics system is being pursued by M. Johnson, and his plan is to first apply it to reading out the output from SVX chips.

Next we list the funds needed, for new items to be purchased or built. Afterward we give some details on the individual items.

1. Silicon strip planes ..... $45 K
2. SVX readout chips ..... 75 K
3. Digitizer modules ..... 18 K
4. Sequencer module ..... 2 K
5. Beam telescope and trigger counters ..... 4 K
6. 4 Fastbus/Camac crates, power supplies, etc. (Pool items)
7. Vertex detector enclosure ..... 15 K
8. Other mechanical construction ..... 15 K
9. Contingency ..... 15 K

TOTAL $189 K

SILICON STRIP PLANES: Planes of various sizes are needed, for the geometry shown in Figure 1. These are all strip planes, 20 micron pitch. Approximately 10 are 5 cm $\times$ 5 cm in size. From E781 experience, we estimate the total cost of such a plane, together with a PC board to hold the SVX chips, to be $3 K. 10 additional planes, mostly 1×1 cm^2 or 2×2 cm^2, are also needed, and we estimate that cost to be about half as much. The total channel count is approximately 30,000.

SVX READOUT CHIPS: This cost corresponds to a quote from UTMC, for the smallest production run they are willing to make. That size run, 5 wafers, will produce about 10 times as many SVX chips as we need for this test setup; many would then be available for other applications at Fermilab.

DIGITIZER MODULES: The digitizer modules used in the E781 detector each handle about 6,000 channels (48 SVX chips). We require 6 such modules. The cost for parts and fabrication is approximately $3 K per module.

VERTEX DETECTOR ENCLOSURE and OTHER MECHANICAL CONSTRUCTION: These costs are estimates based on past experience at Fermilab.

6 Data Acquisition System

The DA system will consist of the following major components:

1. 6 Digitizer Modules (Fastbus), copies of those in the E781 system.
2. 2 Fastbus Memory Modules, Lecroy 1892, 4 MB each.
3. Control computer, to be supplied by Computing Division.
4. An exabyte tape drive.

The Digitizer Modules will be built under the supervision of W. Kononenko, working with the advice of J. Russ. The other components are available at Fermilab. The details of the DA program, and testing of the system, will be under the supervision of K. Sterner.

7 Testbeam Location

Pre-proposal discussions have been held with Roger Dixon and Bruce Baller of Research Division. A suitable beam location was identified as M-TEST. A beam of 250 GeV/c pions is available there, with an intensity of 10^6 particles per spill for 10^{12} protons/spill on target.

8 Personnel

The major work of arranging the procurement and fabrication of the vertex detector system will be carried out by W. Kononenko, K. Sterner, and L. Spiegel, together with W. Selove. J. Butler will be responsible for arranging for the DA system. Development of fiber optics, by M. Johnson *et al.*, and of pixel planes, by D. Christian et al., are proceeding and one or both of these efforts are expected to provide testable components before the end of the 1996-1997 Fixed Target run.

References

[1] Under development by M. Johnson *et al.*.

[2] Under development by E. Barsotti, D. Husby *et al.*.

[3] Under development by D. Christian *et al.*.

[4] W. Selove, "Impact Parameter Trigger and Vertex Detector for Forward Collider Detector", in Proc. of the Workshop on B Physics at Hadron Accelerators (Snowmass), 1993, p. 617 ff, Fermilab Report CONF-93/267, P. McBride and C.S. Mishra, editors.

[5] W. Selove and K. Sterner, "Data Transfer and Trigger Processor System for a High-rate Heavy-Quark Collider Detector, Univ. of Pa. Report UPR 227E, Aug. 1, 1994 (unpublished).

[6] D. Husby, W. Selove, *et al.*, "Design of a Secondary Vertex Trigger System", presented at LeCroy Conference on Electronics for High Energy Physics, May 10-11, 1995.

List of Figures

Fig. 1 Testbeam layout

Fig. 2 Quadrant triplet, showing phi-slice subdivision

Fig. 3 A track hit, in a single-plane phi slice

Fig. 4 A triplet hit, in a station

Fig. 5 Trigger system (Husby)

DESIGN OF A SECONDARY-VERTEX TRIGGER SYSTEM*†

D. Husby, Fermilab
and
P. Chew, K. Sterner, W. Selove
University of Pennsylvania

ABSTRACT

For the selection of beauty and charm events with high efficiency at the Tevatron, a secondary-vertex trigger system is under design. It would operate on forward-geometry events. The system would use on-line tracking of all tracks in the vertex detector, to identify events with clearly detached secondary vertices.

The Tevatron currently produces about 500 times as many B-Bbar pairs per second as the expected rate from the SLAC and KEK B-factories. With the completion of the Main Injector, the rate will increase by another factor of 10 or so. These production rates offer a potentially extremely important capability for highest sensitivity for B physics. A critical component in exploiting this capability is a suitable trigger system. The trigger must be able to reject minimum bias events by a factor of order 1,000:1, while retaining B events with high efficiency. Clearly a secondary-vertex trigger, if highly efficient, would provide the greatest possible sensitivity.

We have designed a secondary-vertex trigger for a forward-geometry hadron collider experiment at the Tevatron. The vertex detector geometry is patterned after that originally proposed by Schlein [1], and operated by him in a test at the S-p-pbar-S [2]. Figure 1 shows an overall side view of the vertex detector and its 31 detector stations. Planes are mounted inside the beam pipe; they are retracted during a beam store and then brought down to about 4 millimeters from the beam.

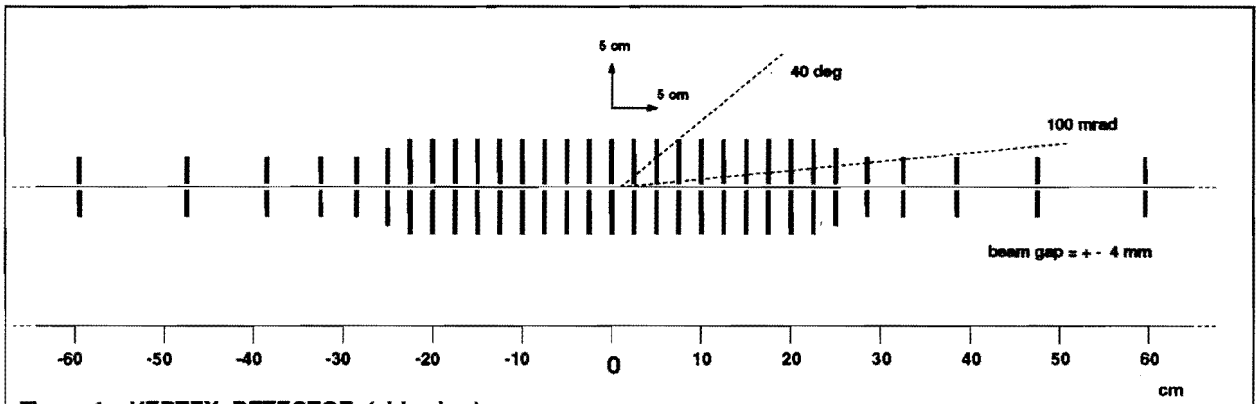

Figure 1. VERTEX DETECTOR (side view)

The primary job of the trigger processor is to analyze every event at interaction rates of 5 MHz or more, and quickly discard minimum bias events while at the same time recognizing and retaining B events with high efficiency. To this end, we have made several major modifications to Schlein's system. The modifications are designed to provide for very fast on-line tracking with momentum information on each track. The momentum information is needed at trigger level in order to do a vertex search using only those tracks which can be extrapolated accurately to the vertex region. Tracks of low transverse momentum will give large smearing of the vertex region, and can cause confusion by producing false secondary vertices. To prevent this smearing, and yet do a vertex search at high speed, we wish to determine the pt of each track, approximately, so that tracks of very low pt can be left out of the vertex search. Momentum information of sufficient accuracy can be obtained, for the forward geometry, purely from the

vertex detector if a suitable magnetic field is present. We note that the momentum information also makes possible a very rapid estimate of the effective mass of the daughter particles at a decay vertex, adding to the ability of the trigger to separate reconstructible B events from other events such as charm events.

In order to obtain fast on-line tracking with momentum information, the trigger system has several important features. Among the most important are the application of a dipole magnetic field at the vertex detector, the arrangement of triplets of planes at each station, the division of station area into mini-areas to facilitate a very high degree of parallelism in processing the individual hits on a track to obtain 3-d space points, and the sorting of hits into groups in azimuthal slices, to facilitate very rapid tracking.

Each of the detector stations consists of a triplet of planes. Ideally each one is a pixel plane to facilitate very fast tracking algorithms. The area of each plane is subdivided into 32 phi slices arranged as 8 "mini areas" in each quadrant as indicated in Figure 2. In this way the average number of tracks passing through one phi slice in one station is brought down to less than one track per interaction for minimum bias events. Thus the tracking problem in each phi slice of each station is a "mini-tracking" problem, and each such phi slice delivers individual station hits which can all be processed in parallel by hit-processors. For the design shown in Figure 2 the data rate on each fiber from each phi slice is approximately 3 hits per microsecond, at luminosity 10^{32}. The hits are fed on 992 separate fiber optic triplets to the trigger electronics.

Figure 2.
Division of Pixel Plane in one quadrant of a station.

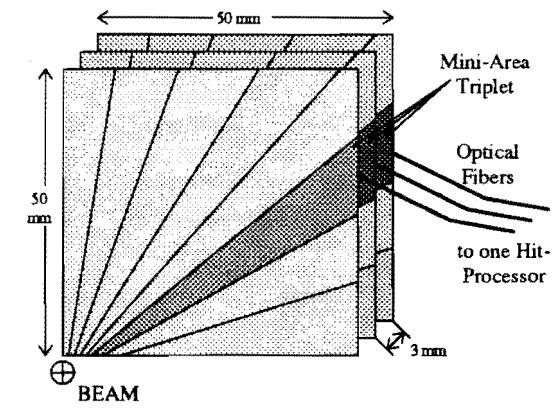

The trigger electronics consists of three levels of processors. At the first level, each fiber triplet feeds a single hit processor where raw hit data is processed independently and then passed to the next level where it is combined with other hits within the same phi slice to form tracks. Similarly, tracks are processed independently (within a phi slice) and combined at the third level to form vertices. At each level, data is combined and summarized, reducing the bandwidth from about 100 Gigabytes per second at the front end to about 100 Megabytes/sec at the output.

Trigger Electronics Data flow through three levels of processors

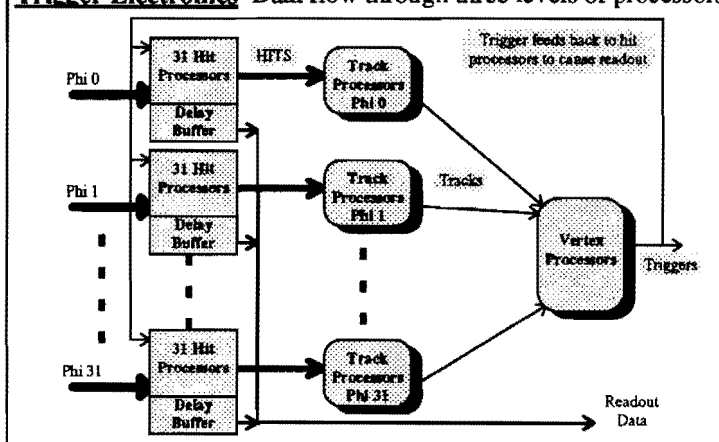

Raw Data from Detector 100 GByte/Sec	992 Hit Processors 330ns/Hit	Hit records 3 GHz 24 GByte/Sec	2048 Track Processors 10 uS/Track	Track records 200 MHz 3.2 GByte/Sec	128 Vertex Processors 16 uS/Vertex	Triggers 10 kHz Readout 100 MB/sec

Initial versions of programs have been written for the hit processor and for the track processor. At luminosity 10^{32}, there will be approximately one interaction per bunch crossing, with a group of 33 bunch crossings occurring at a repetition period of 132 nsec.

The hit processor has to analyze 3 hits per microsecond, on average, during this 33-bunch burst. Initial programming of this operation, by D. Crosetto, for the special purpose processor he has designed [3] gives a processing time of 15 to 20 clock cycles per mini-track [4]. For a commercial processor, with a single operation per clock cycle, it is likely to take somewhat longer.

The track processors have to analyze some 200 tracks per microsecond. An initial program has been written for the tracking operation, and gives a processing time, for a commercial microprocessor, appreciably smaller than 10 microseconds per track. Thus to keep up with the data rate, some 2000 track processors are needed.

Hit Processing

At the predicted hit rate of 3 MHz per hit processor, there's not much time to do detailed processing. For each hit, the processor must do four tasks: extract coordinates, find segments, assemble hit records, and place the raw data into a delay buffer for later readout.

Extract coordinates: When a track passes through a triple layer of pixels, it generally hits a cluster of pixels-- a center pixel and a few pixels on either side as shown in the diagram at right. Accurate hit coordinates are extracted by finding the center of mass of a cluster.

Find segments: it's possible for several tracks to pass through a triplet during one interaction (although in most cases, there will be less than three). A simple segment finder cycles through the list of hits and assigns them to "mini-track" segments. Each segment is represented by a hit record.

Assemble hit records: A hit record is a 64-bit word that contains enough information to describe a track segment's 3-d position and direction, as well as route it to the proper track processor. As shown at right, it contains several fields that must be merged into a single 64-bit word. The first 5 fields are used by the routing circuitry to deliver the hit record to its final destination in a track processor.

Manage delay buffer: Each hit processor must hold raw data in a FIFO delay buffer for several microseconds. When an 'accept' trigger is generated, the processor will transfer data for the accepted event to the readout port. Most of this is done as a DMA operation and requires only a few CPU cycles per trigger.

Hit Processor

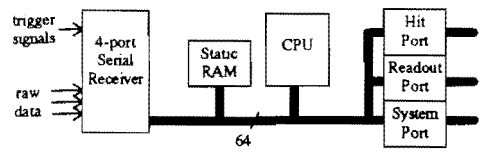

Hit segments in a triplet

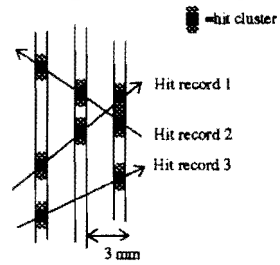

Fields in a hit record

Size	Function
8 bits	Time Stamp
5	Station Number
3	Serial Number
2	Flags
1	Direction
12	Y Slope
15	X coordinate
15	Y coordinate
3	Cluster Size

Track Processing

Hit records are passed through a switching network which sorts them by Timestamp and Direction, (and possibly by Slope), and places them in a buffer in a track-processor's memory. The result of this sorting is that all of the hits for a track are brought together in a single processor's track-buffer. The buffer may contain more than one track, but it will contain complete tracks (as long as they don't cross phi boundaries) which all belong to the same interaction.

There 64 track-processors per phi slice (2048 total), allowing and average of 10 microseconds for processing each track. A load-leveling circuit dispatches new events to the least-busy processor in order to minimize latency.

The format of the track buffer is designed to speed up the track finding algorithm by pre-sorting hits by station number. The buffer has 248 locations arranged as space for 8 hits from each of the 31 stations. The station-number field and serial-number field are used as an index into the buffer when the hit is written. The flags are used by the software when parsing the buffer.

The track processor computes the track momentum and pt, and also a tentative location of the source point if the track originates on the beam line. For tracks with pt below an adjustable cut value, a flag is set in the track output record which is subsequently used by the vertex processor to assign low weight, or zero weight, to this track in the vertex search.

Track Processor

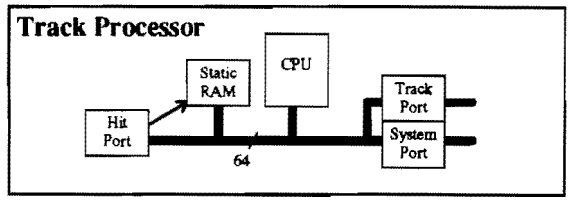

Track buffer

Addr.	Station	Serial#	Flags	
0	0	0	1	(valid)
1	0	1	1	(valid)
2	0	2	3	(end of list)
3	0	x	x	(un-used)
4	0	x	x	
5	0	x	x	
6	0	x	x	
7	0	x	x	
8	1	0	1	
9	1	1	3	
A	1	x	x	
B	1	x	x	
...				
F0	30	0	3	
F1	30	x	x	
F2	30	x	x	
.				
F7	30	x	x	

Track Record

Size	Function
8 bit	Time Stamp
5	Phi Number
3	Serial Number
2	Flags
1	Direction
16	X-Slope
16	Y-Slope
14	X-intercept
14	Y-intercept
14	Momentum
8	py
3	py flags
4	Track quality factor
4	Vertex quality
16	spare
128	total

Vertex Processing

Track records from all 32 phi slices are brought together in the Vertex level. They are sorted by timestamp into vertex buffers. There are 128 vertex processors, allowing an average of 16 microseconds to process each event. When a clearly identifiable secondary vertex is found, and after a rough effective-mass calculation identifies the event as a promising B candidate event, a trigger signal is sent to the rest of the system to cause data for that event to be saved.

Switching Network

When data passes from one processing layer to the next, it must be merged, sorted, and load-leveled. Within each phi slice, for example, there are 31 hit processors emitting hit records that are used by 64 track processors. Data from all 31 hit processors must be combined together and sorted by time stamp. All of the records with the same time stamp must be delivered to a single destination processor. Destination processors must be chosen so that no single processor is swamped with too much work. In addition, all of this must be done at a sustained rate of 10 nanoseconds per hit record.

While this is an engineering challenge, it certainly isn't impossible. At one extreme, such a network could be implemented as a single, 64-bit wide bus with 31 masters and 64 slaves running at 100 MHz. At the other extreme, it could be implemented as 1,984 slow point-to-point links.

The implementation that we have chosen seems to be an optimal middle ground. We divide the bandwidth into four 32-bit buses running at 50 MHz. This keeps all of the parameters within reasonable limits: The bus speed, bus width, and fanout are well within the capabilities of TTL-level drivers. The pipeline cycle speed, 40 ns, is reasonable for low-cost programmable gate arrays. The physical size, 8-16 processors per PC board, is reasonable for a typical 9U-VME-sized module.

As shown in the diagram, processors are grouped into clusters. Hit processors are grouped into four clusters of 8 CPUs. An eight-way combiner merges 64-bit hit records onto a single 32-bit bus. Buses from four combiners feed data into a 4x4 switch that routes words to one of four output buses based on the lower 2 bits of their time stamps. Each output bus feeds a cluster of 16 track processors via a load-leveling buffer manager.

The switching elements are implemented in Field Programmable Gate Arrays (FPGAs). Commodity chips like the AT&T ORCA can implement wide data paths, complex logic, and multiple FIFO buffers on a single chip. The combiner shown at right, for example, is implemented in a single ATT 2C12 chip. It includes eight FIFOs and circuitry for inserting data fields into a 64-bit record.

The sorting switch is implemented using three FPGAs. Each chip is a 12-bit slice of the 32-bit data path. The switch contains a 16-word FIFO at each intersection of four input buses with four output buses. On every 40 ns cycle, a word is transferred from each input to one of four FIFOs in its column. The destination FIFO is selected by the lowest two bits of the word's time stamp. Outputs of four FIFOs in each row are merged together onto a single bus and sent to the output pins.

The FIFOs are necessary to smooth out the data flow. Since it is possible for all four input buses to be sending data to the same output bus, some of the data will have to wait. The FIFOs allow the input buses to continue sending data at the full rate as long as the average rate to all outputs is the same. The FIFOs are 16 words deep, so the switch can accept a sustained burst of 20 hits on all four inputs directed at the same output before a FIFO overflows. Most of the burstiness is eliminated by using the lower two bits of the time stamp to select an output. A time stamp represents 132 nanoseconds. A given output will receive only the data from every fourth time stamp.

Implementation of switching network

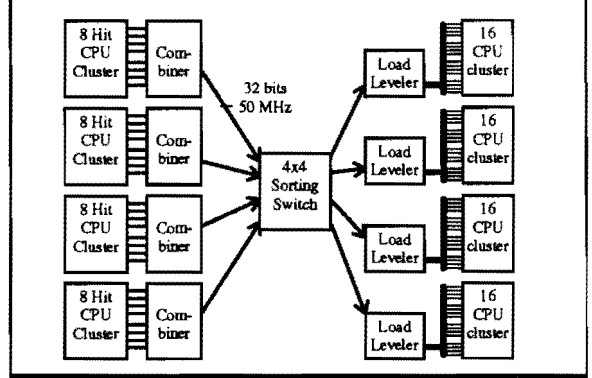

8-way Combiner Chip

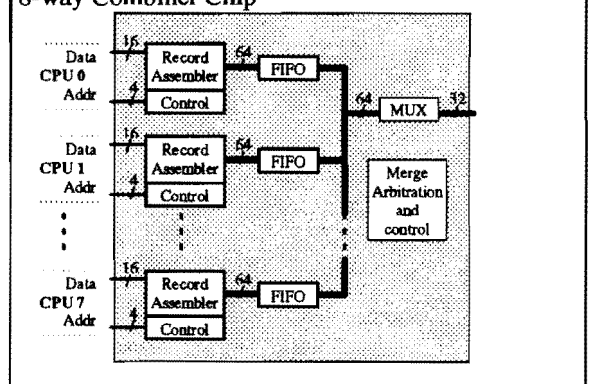

4x4 Sorting Switch

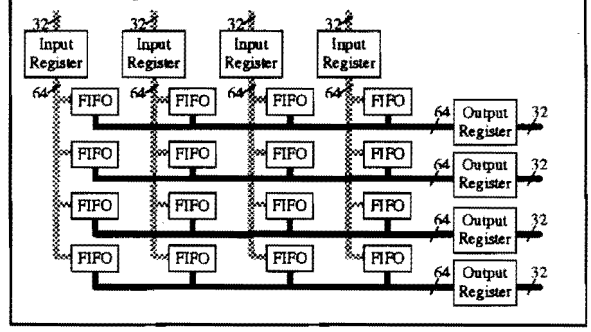

The load leveler performs a similar smoothing function. It keeps track of which processors are busy, and assigns new events to the least-busy processors. An event (interaction) is defined as all hit records with the same time stamp.

The load leveler keeps a translation table that translates time stamps (and other sorting fields) to buffer addresses. When a hit record with a previously unseen time stamp is received, it causes a cache miss which causes the buffer manager to select a free buffer and write its address into the translation table. Subsequent hits with the same time stamp will all be directed to the same buffer. When a processor has finished processing an event, it releases the buffer so that the manager can use it for another event.

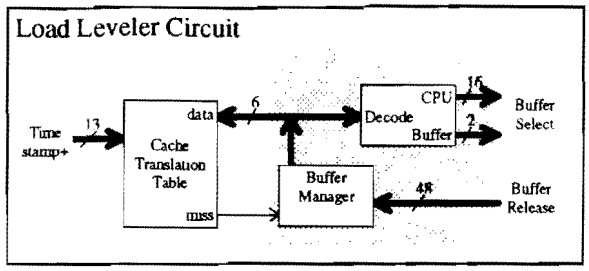

Each of the 16 processors in a cluster has three buffers for a total of 48 buffers. This allows a processor to be working on one event while two others are being transmitted. The buffer manager chooses processors based on the number of buffers in use. It fills all processors with one event before sending a second event to any processor.

Packaging

Clusters of 8 or 16 processors along with their combiner and/or load leveler are put on a single PC board. These boards plug into a switching backplane that contains the sorting switch as well as system bus and readout bus. The whole package of 8 processor boards and backplane make a box that is 8" x 16" x 16". This box is called a Phi-box, since it processes all of the tracks for one phi section.

Identical boxes are used to do vertex processing, but they are called Vertex-boxes. The trigger system is made up of 32 Phi-boxes and 4 Vertex-boxes. Track records from each of the four clusters within a Phi-box go to a different Vertex-box via a high speed optical link. A diagram of the entire system is shown below.

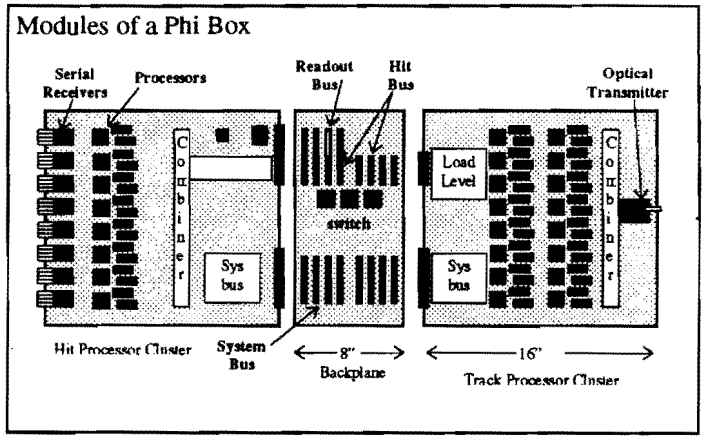

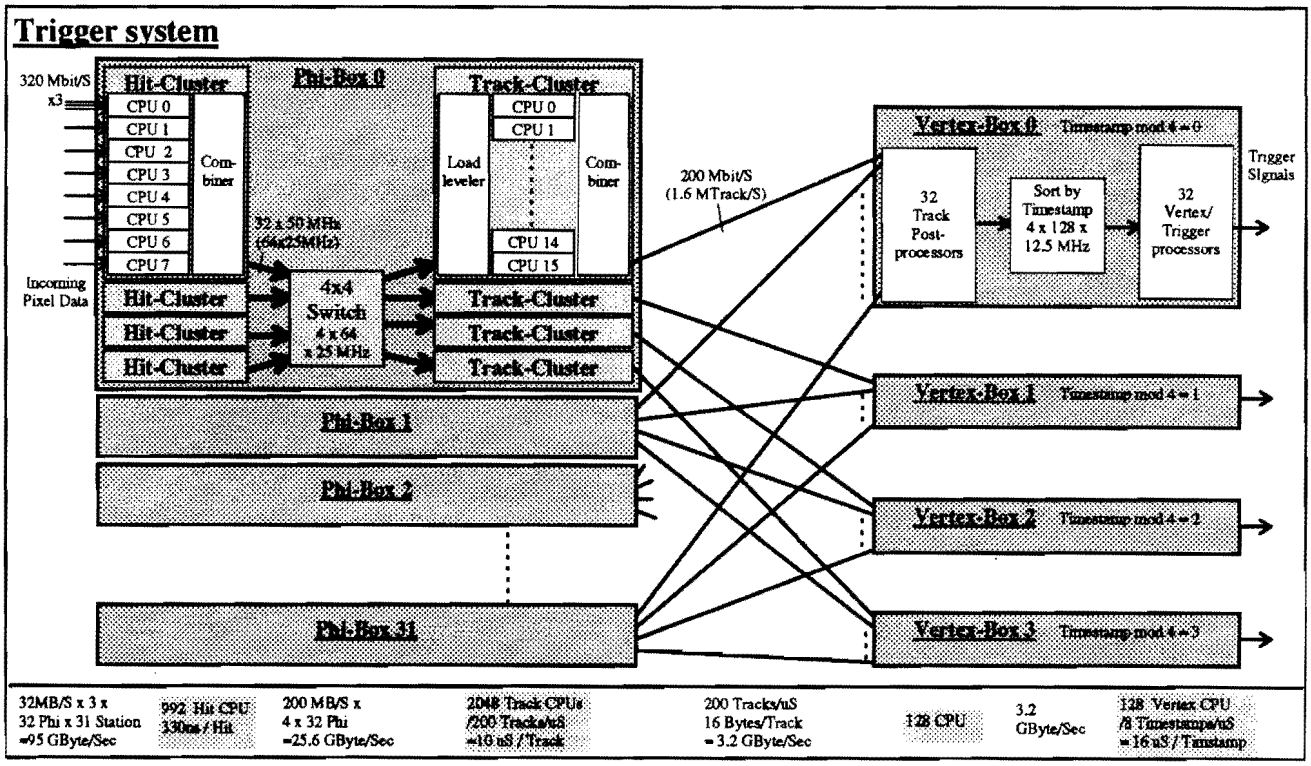

Prototype

A prototype phi box is being built by the Fermilab Electronic Systems Engineering group. The prototype will include a single hit cluster and track cluster and a full 8-slot backplane with VME or PCI system interface. It will be a full working prototype, except that it will not have the serial links or readout bus. Work on this prototype began in March, 1995 and is expected to be completed by early 1996.

The microprocessors used in the prototype are from the MIPS 4600 family. They are the same type of processor found in Silicon Graphic's Indy and Challenge systems. They can execute 133 Million instructions per second and 44 MFlops. Since this particular processor, the Orion 4650, is targeted for the video game market, it will sell for under $100 per chip.

The choice of processor will be re-evaluated when work on the production system begins. Since microprocessors are doubling their performance every two years, there is a significant advantage in waiting as long as possible before committing to a processor. One likely candidate is a processor currently being proposed by Dario Crosetto [3]. His processor is specifically designed for physics trigger applications. Very long instruction words and multiple function units make it well suited for the type of processing done by high speed triggers. In addition, his design includes four processors on a single chip, making it possible to fit thousands of processors in a reasonable amount of space.

ACKNOWLEDGMENTS:

We express our thanks to J. Butler and H.H. Williams for their assistance and support.
E. Barsotti, D. Crosetto, M. Johnson, S. Mani, and S. Shapiro have provided much helpful information.

REFERENCES:

* Paper presented at the LeCroy Conference on Electronics for Particle Physics, Chestnut Ridge NY, May 1995.

† Work supported by U.S. Department of Energy under contract DE-AC02-76CHO3000 at Fermilab and partially supported under contract DE-ACO2-76-ERO-3071 at the University of Pennsylvania.

[1] S. Erhan, *et al.* (COBEX Collaboration)
"COBEX, A Dedicated Collider B Experiment"
Nucl. Inst. Meth. **A333** (1993) 101.

[2] J. Ellett, *et al.* (P238)
"Development and Test of a Large Silicon Strip System for a Hadron Collider Beauty Trigger"
Nucl. Inst. Meth. **A317** (1993) 28.

[3] D. Crosetto
"Digital Programmable Level-1 Trigger with 3D-Flow Assembly"
SSCL-PP-445, August 1993. Submitted to Nucl. Inst. Meth.

[4] D. Crosetto, private communication on hit processor calculations.

EMAIL ADDRESSES:

Philip Chew	chew@upenn5.hep.upenn.edu
Don Hubsy	husby@fnal.gov
Walter Selove	ws@upenn5.hep.upenn.edu
Kevin Sterner	sterner@upenn5.hep.upenn.edu

Copies of this document and others related to the trigger electronics can be obtained via anon-ftp from:
esesrv0.fnal.gov:/pub/bphys/lecroy.ps

TEST-BEAM LAYOUT

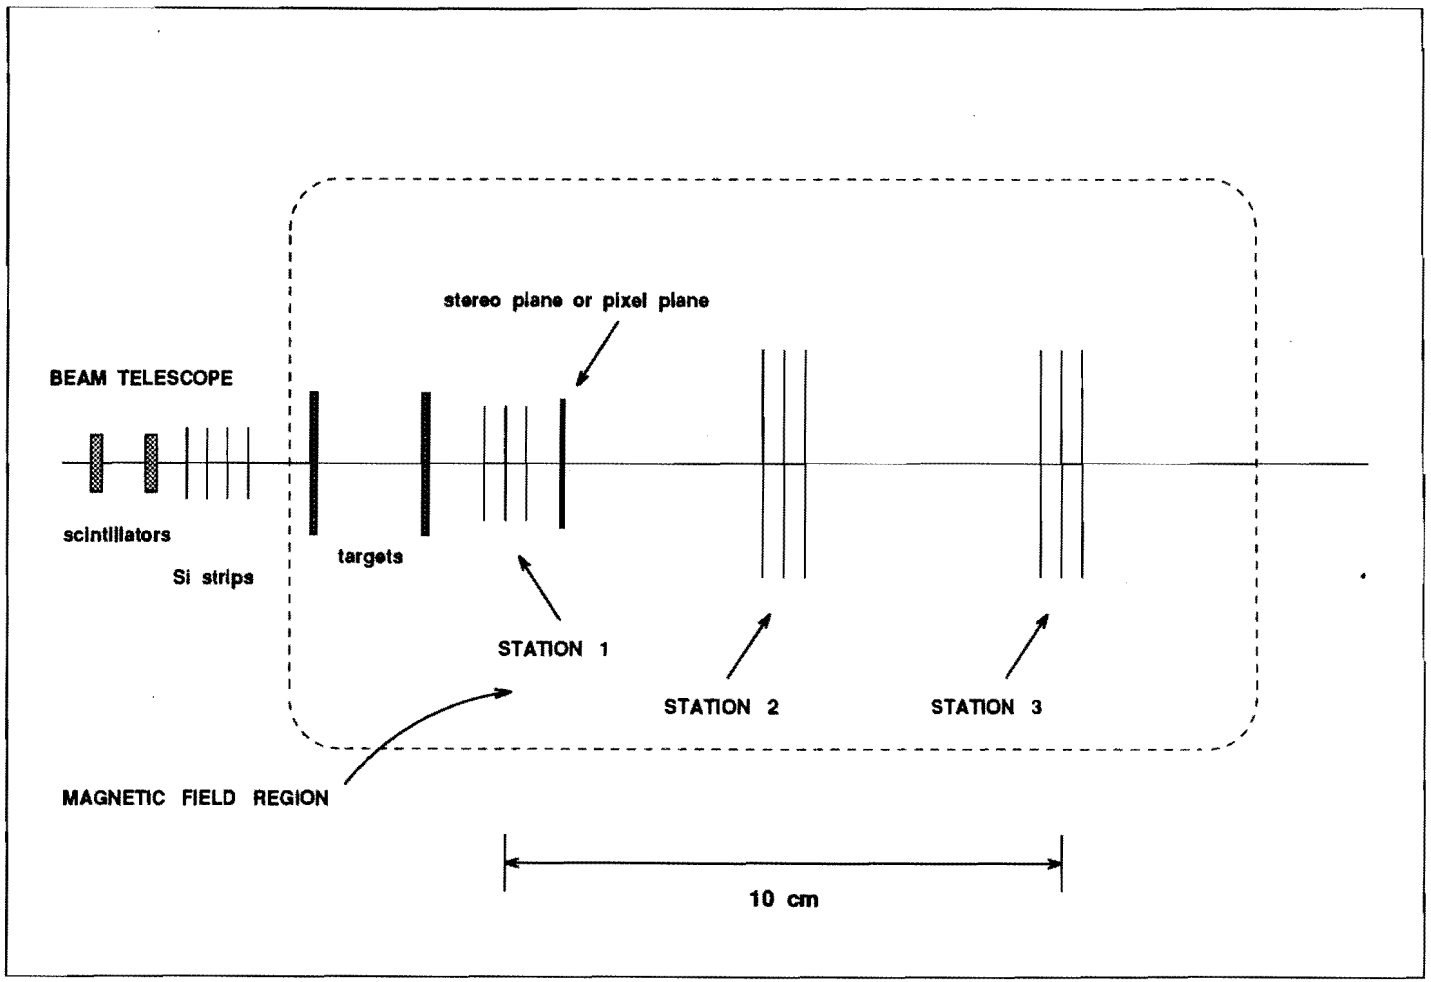

FIGURE 1

A QUADRANT TRIPLET OF ONE STATION

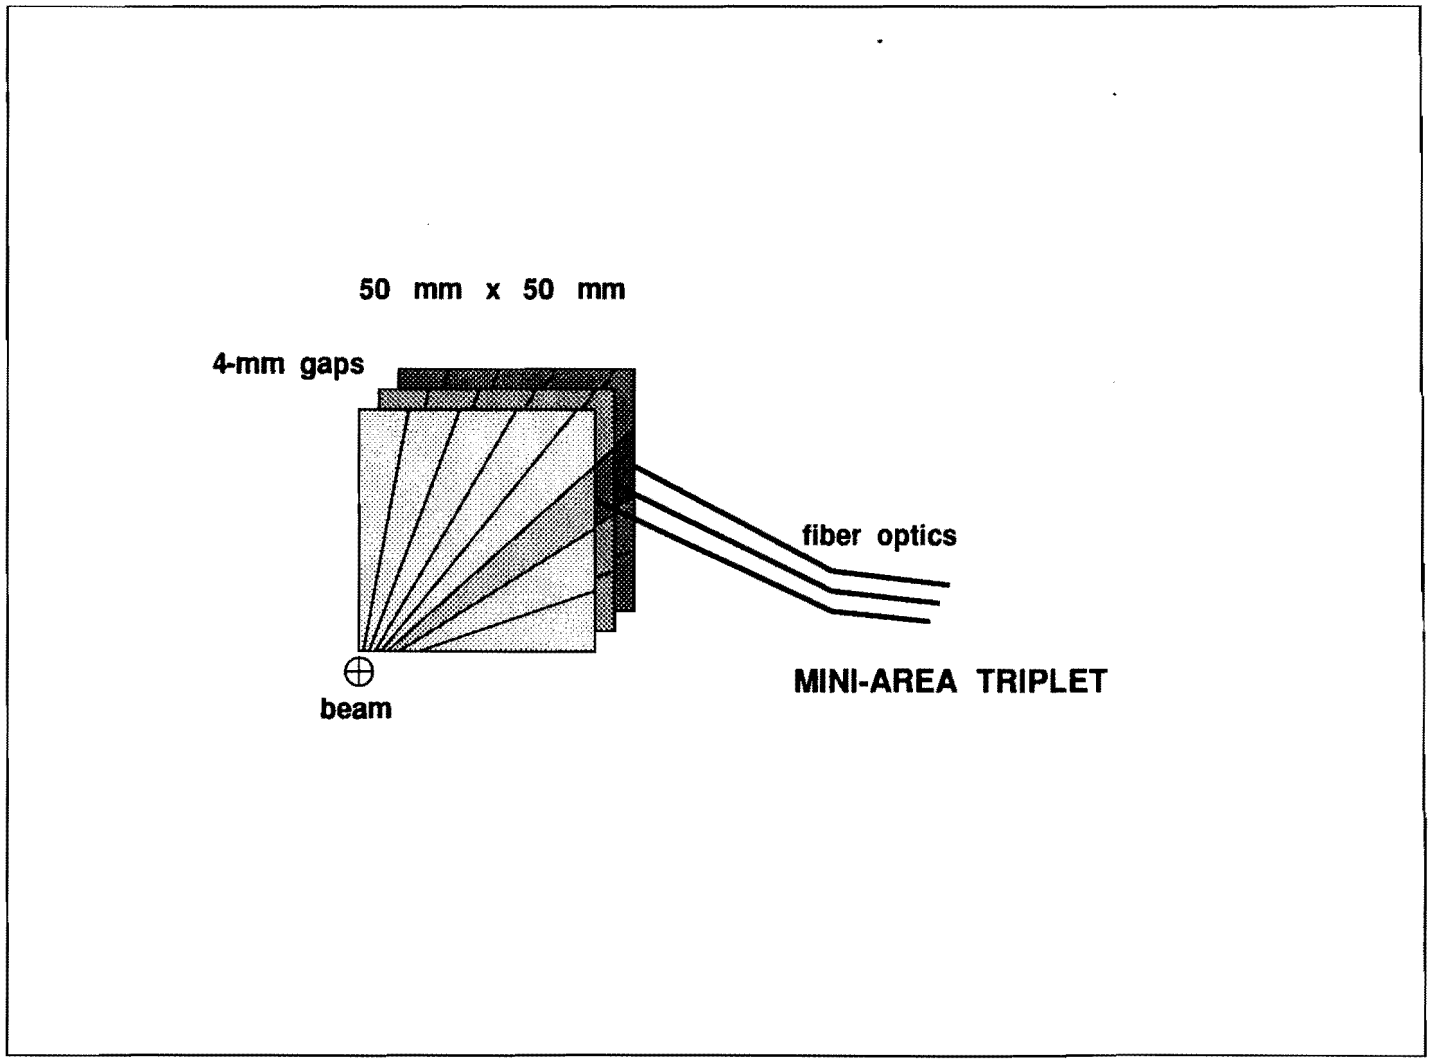

FIGURE 2

SINGLE PLANE HIT

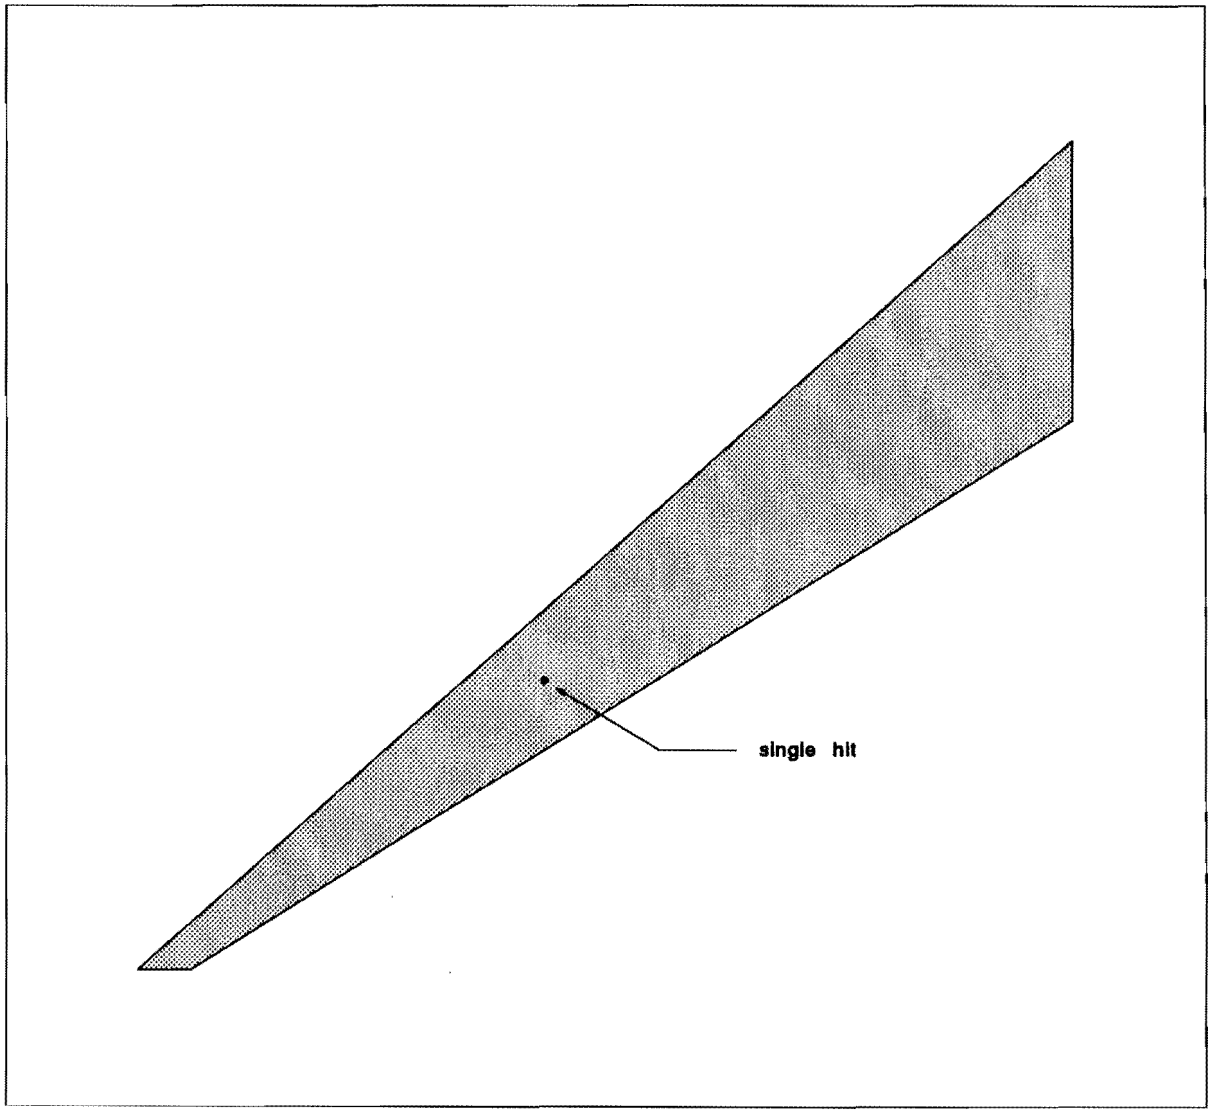

FIGURE 3

TRIPLET HIT

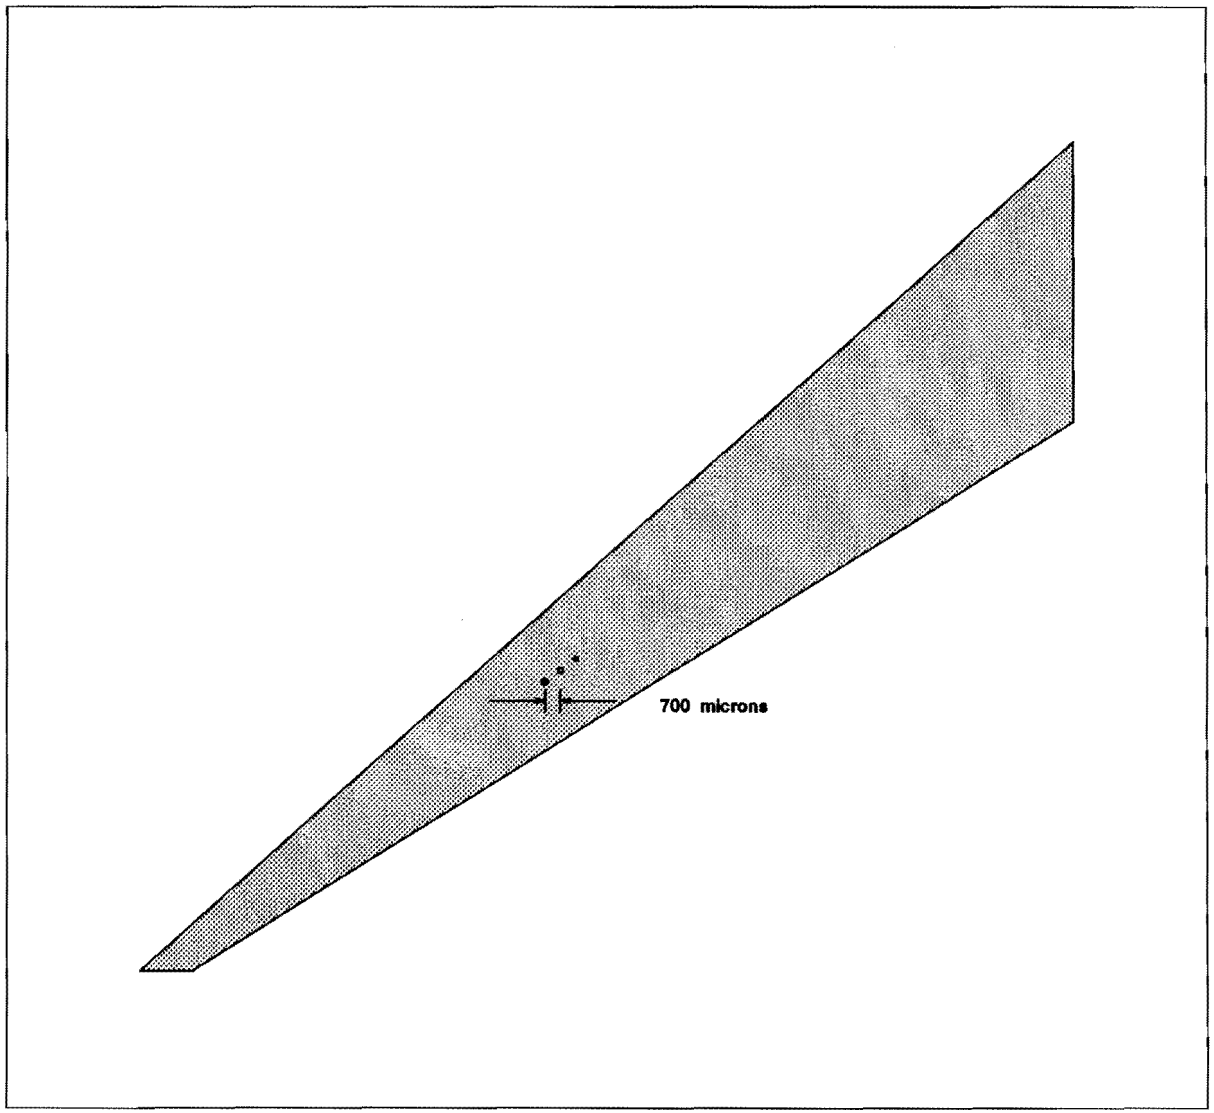

FIGURE 4

TRIGGER SYSTEM LAYOUT (Husby)

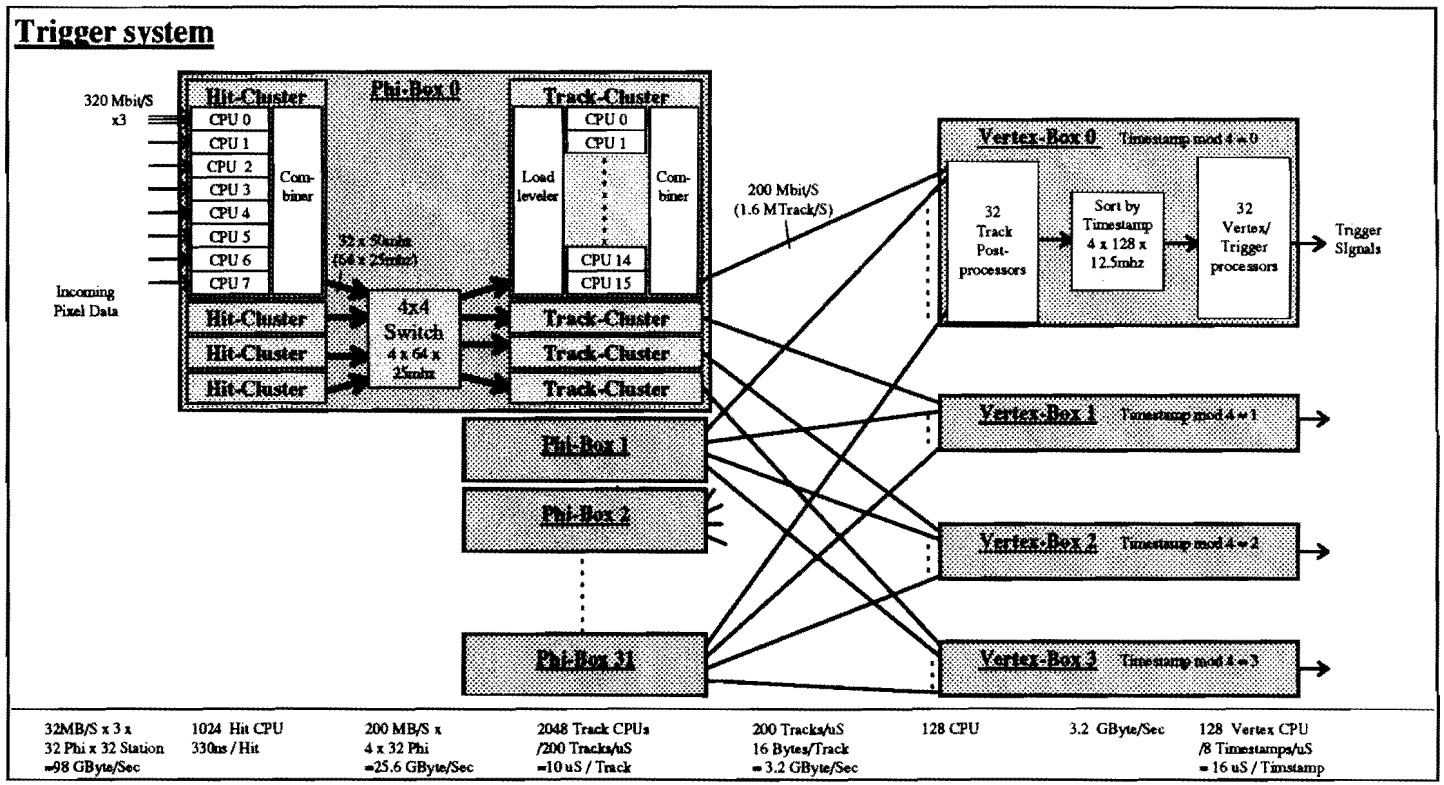

FIGURE 5