Fermi National Accelerator Laboratory

FERMILAB-Pub-87/191-E
[E-741/CDF]

A Fast Hardware Track-Finder for the CDF Central Tracking Chamber*

G. W. Foster, J. Freeman, C. Newman-Holmes, and J. Patrick
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

July-August 1987

*Submitted to Nucl. Instrum. Methods A

A FAST HARDWARE TRACK-FINDER FOR THE CDF CENTRAL TRACKING CHAMBER

G. W. Foster, J. Freeman, C. Newman-Holmes, J. Patrick

Fermi National Accelerator Laboratory[1], Batavia, IL

May 14, 1987

Abstract

We describe the design, construction, and preliminary operation of a hardware track-finder used in the trigger for the Collider Detector at Fermilab (CDF). The track-finder is a 19 stage digital pipeline which uses fast timing information from the Central Tracking Chamber (CTC) to find high momentum tracks. It then transmits information about the tracks to the CDF trigger system for a final decision. The track finder has 8 programmable thresholds between 2.5 and 15 GeV. A search for all high momentum tracks in the Central Tracking Chamber can be completed in an average of 2.5 μsec per event. The momentum resolution is $\delta P_t/P_t^2 = 3.5\%$, with a high efficiency that is independent of track density.

1 Introduction

We have built a hardware track-finder for the CDF Central Tracking Chamber [1]. The track-finder's function is to detect the presence of high transverse momentum (P_t) charged particles passing through the drift chamber, determine the momentum

[1]Operated by Universities Research Association under Contract with the U.S. Department of Energy, No. DE-AC02-76CH03000

of the particles, and present the results to the rest of the CDF multilevel trigger [2]. The results are then used either as a stand-alone trigger or in combination with calorimeter and muon detector information to construct triggers for electrons, photons, and muons. In the first data-taking run of CDF (Jan. - May, 1987), our track-finder was used in coincidence with the central muon [3,4] trigger to reduce the trigger rate to an acceptable level. In addition, it was used as a stand-alone trigger to extend the limits of the measured charged particle P_t spectrum.

1.1 Geometry of the Track Processing

The conceptual operation of the track-finder is intimately tied to the geometry of the Central Tracking Chamber [1]. The CTC is an axial wire chamber which operates in a solenoidal magnetic field. It is composed of nine "superlayers" of sense wires, alternating axial and stereo views, with a total of five axial and four stereo superlayers. Our track processor uses only the r-ϕ information of the axial wires to measure the transverse momentum of the tracks. Each axial superlayer consists of 12 layers of sense wires and the associated field wires. In total, the chamber has 4392 axial sense wires. "Supercells" of an axial superlayer are shown in Fig. 1. In this figure, the crosses correspond to the locations of the sense wires, and circles to the field wires. The sense wires of a supercell lie along a plane that is rotated 45° from the radial direction. Thus any high P_t (approximately radial) track will cross at least one sense wire plane in each superlayer (see Fig. 1). This generates a series of hits with small drift times, called "prompt" hits. These prompt hits are the basis of our device's operation. The track-finder begins operation for

each beam crossing (every 7.0 microseconds) by applying a coincidence latch to the digital stream of fast out data for each axial sense wire (4392 channels). The fast out data are time-over-threshold wire chamber signals passed through Lecroy 1879 TDCs. The coincidence gate is 80 ns long to ensure that each sense plane crossing generates at least one prompt hit. A coincidence gate was chosen instead of an edge-triggered device because of the danger that in events with high track density, hits which overlap in time would not register separate leading edges. This would generate an efficiency bias against e.g. tracks inside the tightly collimated jets at a proton collider. A coincidence gate with time-over-threshold chamber data does not depend on pulse edges, and avoids this problem. A second "delayed" coincidence gate is applied to record the presence of hits at a larger drift distance. This delayed gate occurs 500-650 ns following the beam crossing, which corresponds to $\sim 2/3$ of the cell size. The width of the gate is chosen to ensure at least two delayed hits in each superlayer. These "delayed" hits are processed on a second pass through the data, and are used to verify any high P_t tracks found while processing "prompt" hits. The geometry of delayed hits is shown in Fig. 1.

Fig. 2 shows a typical "road" or hit pattern that the hardware will recognize as a high P_t track. The road is defined by a prompt hit in the outer superlayer and a range of P_t. The short vertical lines along the track in Fig. 2 indicate the geometrical acceptance of the prompt and delayed hits expected to be present for a track in this momentum bin (4.75-5.5 GeV). The figure is compressed along the track direction in order to emphasize the curvature of the tracks. The vertical scale is 30 cm and the horizontal scale is 150 cm. A total of 32 roads are defined for

each sense wire in the outer superlayer. These are divided into eight P_t bins and two azimuth bins for each sign of curvature in order to cover the entire momentum range above 2.5 GeV.

1.2 Serial Processing

We have chosen an approach in which digital data from the chamber is processed sequentially in ϕ. It is data-driven in the sense that the first step in the processing is to form a list of hits in the outer superlayer of the chamber. This list is then used to look inward for the "roads" (hit patterns) of all possible tracks which lead to a given hit. This takes place in parallel on a single clock cycle for roads of all acceptable curvatures. We have chosen this mixture of serial and parallel processing, instead of a fully parallel approach, for a number of reasons. A fully parallel search for all road patterns is possible in principle in as few as two levels of logic. Such an approach might use the pattern of hits to directly access a RAM or PLA device. This rapidly becomes untenable as the number of channels in the road pattern increases. In our case, a RAM or PLA with 120 address lines would be required.

Another advantage of a serial processing scheme is that the wiring within the trigger processor is simplified considerably. A fully parallel approach generally requires complicated wiring topology since it involves correlation of signals which are physically separated on the chamber (and, in our case, in the TDC readout as well). Since we process the chamber data from all layers sequentially in ϕ, we are able to use a shift register readout for the wire hit bits. This greatly simplifies the inter-crate wiring, since only the bits required for processing a given ϕ must

be transported to the the central coincidence logic on each clock cycle. In order to obtain the desired readout bandwidth, the shift registers must be wide: during processing, the 96-bit wide shift registers on each of the five axial superlayers have a combined bandwidth of 15 Gbits/sec.

2 Hardware Description

2.1 Nineteen Stage Digital Pipeline

Our track processor makes heavy use of pipelining in order to maximize the processing speed. Pipelining is a technique in which information is processed gradually as it passes through a series of simple operations, with intermediate results at each stage residing in storage registers. These registers are synchronously clocked to advance each datum to the next stage in the pipeline. This allows processing throughput determined by a clock cycle time allowed by the slowest circuit element in the pipeline, which is typically much greater than could be achieved if all operations were performed in a single cycle. Table 1 lists the operations performed at each of the 19 pipeline stages in the track processor.

The minimum clock cycle time of the processor (35 ns) is set by the RAM table lookup time (25 ns) plus the setup and hold times of the pipeline registers used (10 ns). Other digital operations such as bit field encoding, data multiplexing, threshold comparisons, etc. are performed with combinatorial logic and PLA's which are faster than the RAM access time. Some pipeline stages consist simply of transfer of information from one module to another along a flat cable several feet

in length; in these cases, the pipeline clock is delayed by the same cable length to ensure that the clock and data arrive with the appropriate phase at the end of the cable. In cases where the data processing pipeline diverges (in one case spreading out to 13 modules in four different crates), care is taken to equalize the delay among the different branches so that the data is synchronized properly when the pipeline reconverges for the final trigger decision.

While most elements of our system have been successfully tested with a 32 ns clock (slightly above the design clock speed of 35 ns), data taken during the 1987 run were primarily at clock speeds of 40 or 50 ns.

2.2 Advanced CMOS/Advanced Schottky TTL

The track processor benefitted enormously from the choice of Advanced CMOS and Advanced Schottky TTL logic rather than ECL, the traditional choice for projects of this type. Fast (25 ns), dense (16 kbit), and low power (50 mw) CMOS RAMS were used for lookup tables and data storage. CMOS octal registers with 250 MHz maximum clock frequencies were used for time-critical pipline stages, and AS TTL used elsewhere. This technology also benefits from a rapid rate of advancement (for example, 12 ns CMOS RAMS are now available as drop-in replacement for the 25 ns parts initially used).

2.3 Mechanical Assembly

Of the 61 circuit cards in the track processor, 46 are coincidence latch/shift register readout cards (Fig. 3) which use surface-mounted components. The use of surface

mount components on both sides of the PC board allowed the entire circuit (over 160 IC's, resistor arrays, capacitors, and connectors) to be mounted on a single Auxiliary FASTBUS card of about 50 in^2. The remainder of the circuit boards are FASTBUS size cards which only draw power from the crates used for the TDC readout. These include: 1) Five "bitfield extractor" cards which are attached to the end of the readout shift register for each superlayer. These control the shift registers and extract, from the end of the shift register, the 16 bits which contain all possible roads leading to a given hit (see Fig. 2). They also record the full pattern of hits for inclusion in the event record. 2) Eight "subroad" cards containing the RAMs which define roads of four different curvatures. These cards also contain the majority logic which counts the number of hits for each road on each clock cycle. 3) One "master" card which produces the list of hits from the outer superlayer. 4) One "decision" card which chooses among all track candidates found on a given clock cycle, suppresses instances in which the same track might be found on successive clock cycles, generates the first level trigger signal, and contains the FIFO interface which transmits the track list to the second level trigger electronics.

2.4 Diagnostic Features and Software

Debugging a system with a large digital pipeline is complicated by the fact that data is processed not in a single step but gradually as it flows through the pipeline. In addition, the clock rate can always be increased to a point where one or more stages of the pipeline will fail to operate correctly. In such a case it is often difficult to tell at which pipeline stage the data is being corrupted. We responded to this by

developing the following tools: 1) Extensive diagnostic readback circuitry to enable the host CPU to read out (and force diagnostic data into) the contents of registers at many locations in the digital pipeline. 2) A detailed simulation of the hardware which predicts the contents of all pipeline registers at any point in time. 3) The ability to download test patterns: real or simulated track data, random data or data designed to excite "worst case" switching conditions. 4) The ability to generate, under program control, a selected number of clock cycles at a computer-selected frequency.

These tools enabled us, for example, to download track information from an event in which a discrepancy had been observed between the hardware trigger decision, and the trigger decision as simulated offline. The hardware was then stepped, at full clock speed, to the point at which the discrepancy was observed. The contents of the pipeline were then read out, and the fault (whether in the hardware or the simulation code) could then be observed and isolated. This procedure also provided a means of verifying the simulation code which is now used for offline analysis and Monte Carlo studies.

2.5 Computer Readout

The host processor for downloading, diagnostics, and simulation was a Motorola MC68020 CPU board from the Fermilab Advanced Computer Program [5]. The 68020 readout is independent of the CDF FASTBUS system. In order to get data into the data stream and into the event record, FASTBUS readout was performed using the simple TTL interface provided by the LeCroy 1879 TDC's used to record

the chamber data. Event related information recorded included: 1) The hit patterns of all prompt and delayed coincidence latch bits. This data alone was sufficient to permit offline simulation and verification of the trigger decision. 2) A list of found tracks, including momentum, sign, and azimuth information. 3) The contents of various registers at the end of processing, which could be compared to the simulation program to provide offline verification of the track processor's performance.

2.6 Central Muon Trigger Interface

Following each beam crossing, the trigger processor produces a list of found tracks. This list is entered into a FIFO and subsequently transmitted to a device (the "Muon Matchbox") which correlates these tracks with hits in the central muon chambers. The Muon Matchbox contains RAM-based lookup tables which use the track curvatures, with allowance for momentum-dependent multiple scattering, to predict the positions of muon chamber hits. This r-ϕ matching requirement reduces the muon trigger rate by a factor of $\sim$10. This factor was crucial in the 1987 run since the absence of higher level trigger electronics made the deadtime of the data acquisition system marginal even with this reduced muon trigger rate (0.2 Hz at a luminosity of $10^{29}\,\mathrm{cm}^{-2}\,\mathrm{sec}^{-1}$). It is expected that this rate will be reduced by a series of software improvements.

3 Performance

3.1 Performance with Central Muon Trigger

The CDF trigger is a three-level system characterized by a Level 1 trigger decision before the next beam crossing (7 μsec in the 1987 run); a Level 2 trigger decision taking up to several beam crossing times, and a Level 3 trigger in software. In the final system, it is planned to force a Level 1 trigger whenever 1) a high momentum track is flagged by the CTC track processor, or 2) the muon chambers record a track pointing to the event vertex, or 3) the hit pattern in the CTC is sufficiently complex that the track processor has not finished by the time of the next beam crossing. The matching of high-P_t tracks with muon chamber hits then takes place during the Level 2 trigger decision time. During the 1987 run, however, the Level 2 trigger logic was not available and a major effort was undertaken to make the hardware operate in the much faster Level 1 trigger decision time. This involved operating the track processor pipeline clock at the highest possible rate, and transfering the track list to the Muon Matchbox as fast as possible. The speed of the pipeline clock was gradually increased so that at the end of the run the pipeline was operable with a 40 ns clock, while retaining complete agreement with the offline simulation code. To permit operation of the Muon Matchbox in parallel with continued operation of the track processor, the track processor was modified to transmit the first entries in the track list FIFO before processing was completed. The readout protocol of the Muon Matchbox was modified to provide a quasi-pipelined readout on the 200 ns long interface cable, with the Matchbox asynchronously sampling the list of wires

produced by the trigger processor.

These improvements allowed the central muon trigger rate to remain manageable (less than 0.25 Hz) as the luminosity increased by over 2 orders of magnitude during the run. During this time there were trigger inefficiencies (typically totalling 10-25%) from a number of causes. High voltage on sections of the Central Tracking Chamber was shut down, and software did not yet exist to correct the downloaded RAM tables for these dead wires and cells. At the slower track processor clocking rates initially used, there was a significant probability that for busy events the track processing and ϕ-matching would not progress around all 2π of the chamber circumference before the Level 1 decision time. This resulted in an azimuthal dependence in the trigger efficiency. The changes to the Muon Matchbox protocol resulted in a certain probability of track candidates being lost in transmission. The prompt and delayed gate timings were not fully optimized with respect to the RAM lookup tables. During a period of the running, errors in the ϕ-matching tables of the Muon Matchbox effectively deleted 7 of the 24 ϕ-slices. Despite these difficulties, a preliminary estimate of the combined efficiency of the CTC track processor and central muon trigger is 70% (averaged over the run). By the end of the run, and in regions of ϕ where the above problems were not severe, the track finding efficiency of the track processor itself was approaching 100% (see sect. 3.3).

3.2 Execution Time

The track finder generates both a Level 1 trigger decision and information for the Level 2 trigger. Each has a distribution of times to completion, owing to the serial

nature of the device and varying event complexity.

A Level 1 trigger may be generated if a high momentum track is identified using the prompt hits alone. The dotted line in Fig. 4 shows the distribution of times for the Level 1 decision to be formed, for events with only a beam-beam counter "minimum bias" trigger requirement. The small peak at low time is due to events in which a high momentum track is found before all prompt hits are processed. The large peak at later time gives the distribution of times for completion of prompt hit processing. No track is found in most of the events in the second peak. In general, the Level 1 decision is made before 3 μsec.

For a Level 2 decision, track candidates found during prompt hit processing are validated by examining the delayed hit patterns. The surviving tracks are passed to the Level 2 decision logic where they are matched with hits in the central muon chambers (see Sec. 2.6). The solid curve in Fig. 4 shows the distribution of total execution times for minimum bias events. The Level 2 information is generally available in less than 8 μsec.

3.3 Efficiency and Momentum Resolution

To calculate the efficiency of the track processor, software track reconstruction was performed on a sample of events taken with an independent (calorimetry) trigger. Tracks used in the efficiency measurement were required to have reconstructed P_t greater than 2.0 GeV/c and were required to pass through all five axial superlayers of the CTC. Only tracks in one quadrant of the CTC (where chamber performance was optimal) were used.

The prompt and delayed hits from the event record were used as input to the processor simulation to obtain a list of tracks found by the processor. Next tracks were found using the offline track reconstruction program. For each track found by the offline program, the track list generated by the processor simulation was searched for a match. The efficiency is then the ratio of the number of tracks with a processor match to the total number of tracks found by the offline program. This procedure avoids overestimation of our efficiency due to the presence of multiple tracks in the event. The use of the simulation (which is in agreement with the hardware results in the event record) allows for the offline study of the effects of different P_t thresholds. Fig. 5 shows track processor efficiency vs track P_t for a nominal P_t threshold of 2.5 GeV.

The momentum resolution of the track processor was calculated for a sample of tracks identified as muons by the offline program. This event sample was chosen because it contains events which are ultimately of interest for the track processsor/central muon trigger. In addition, the tracks were required to be in one quadrant of the CTC as described above. Tracks found by the software were matched to those found by the track processor, and the quantity $\delta P_t/P_t^2$ was calculated. Fig. 6 shows this distribution. The HWHM of this distribution is 3.5%. Similar results can be obtained with other event samples. It is expected that this will be reduced by optimization of the RAM tables and of the coincidence gate widths and positions. For comparison, $\delta P_t/P_t^2$ from the software track reconstruction is about 0.2%. The asymmetry of the distribution is due to the fact that the processor always chooses the highest P_t track candidate from among all track can-

didates with the same coincidence level. This conservative choice is appropriate for a device used at the trigger level, but occasionally results in overestimation of track momenta when the track is accompanied by a large jet or a burst of chamber noise.

4 Conclusion

We have built and successfully operated a processor to find and identify tracks left by high momentum particles in the CDF Central Tracking Chamber. The processor triggers on high momentum tracks in typically 2.5 μsec with an efficiency of essentially 100% for P_t above 3.5 GeV. The processor output was used to form an efficient Central Muon trigger during the 1987 run.

References

1. F. Bedeschi et al., "Design and Construction of the CDF Central Tracking Chamber", submitted to Nucl. Instrum. and Methods.

2. D Amidei et al., "A Two Level FASTBUS Based Trigger for CDF", submitted to Nucl. Instrum. and Methods.

3. G. Ascoli et al., "The CDF Central Muon Detector", submitted to Nucl. Instrum. and Methods.

4. G. Ascoli et al., "The CDF Central Muon Level 1 Trigger Electronics", submitted to Nucl. Instrum. and Methods.

5. H. Areti, et al., "ACP Modular Processing System: Design Specifications", Fermilab Pub #FN-403, Revision April 2, 1984.

List of Figures

6. The transverse momentum resolution for the track processor.

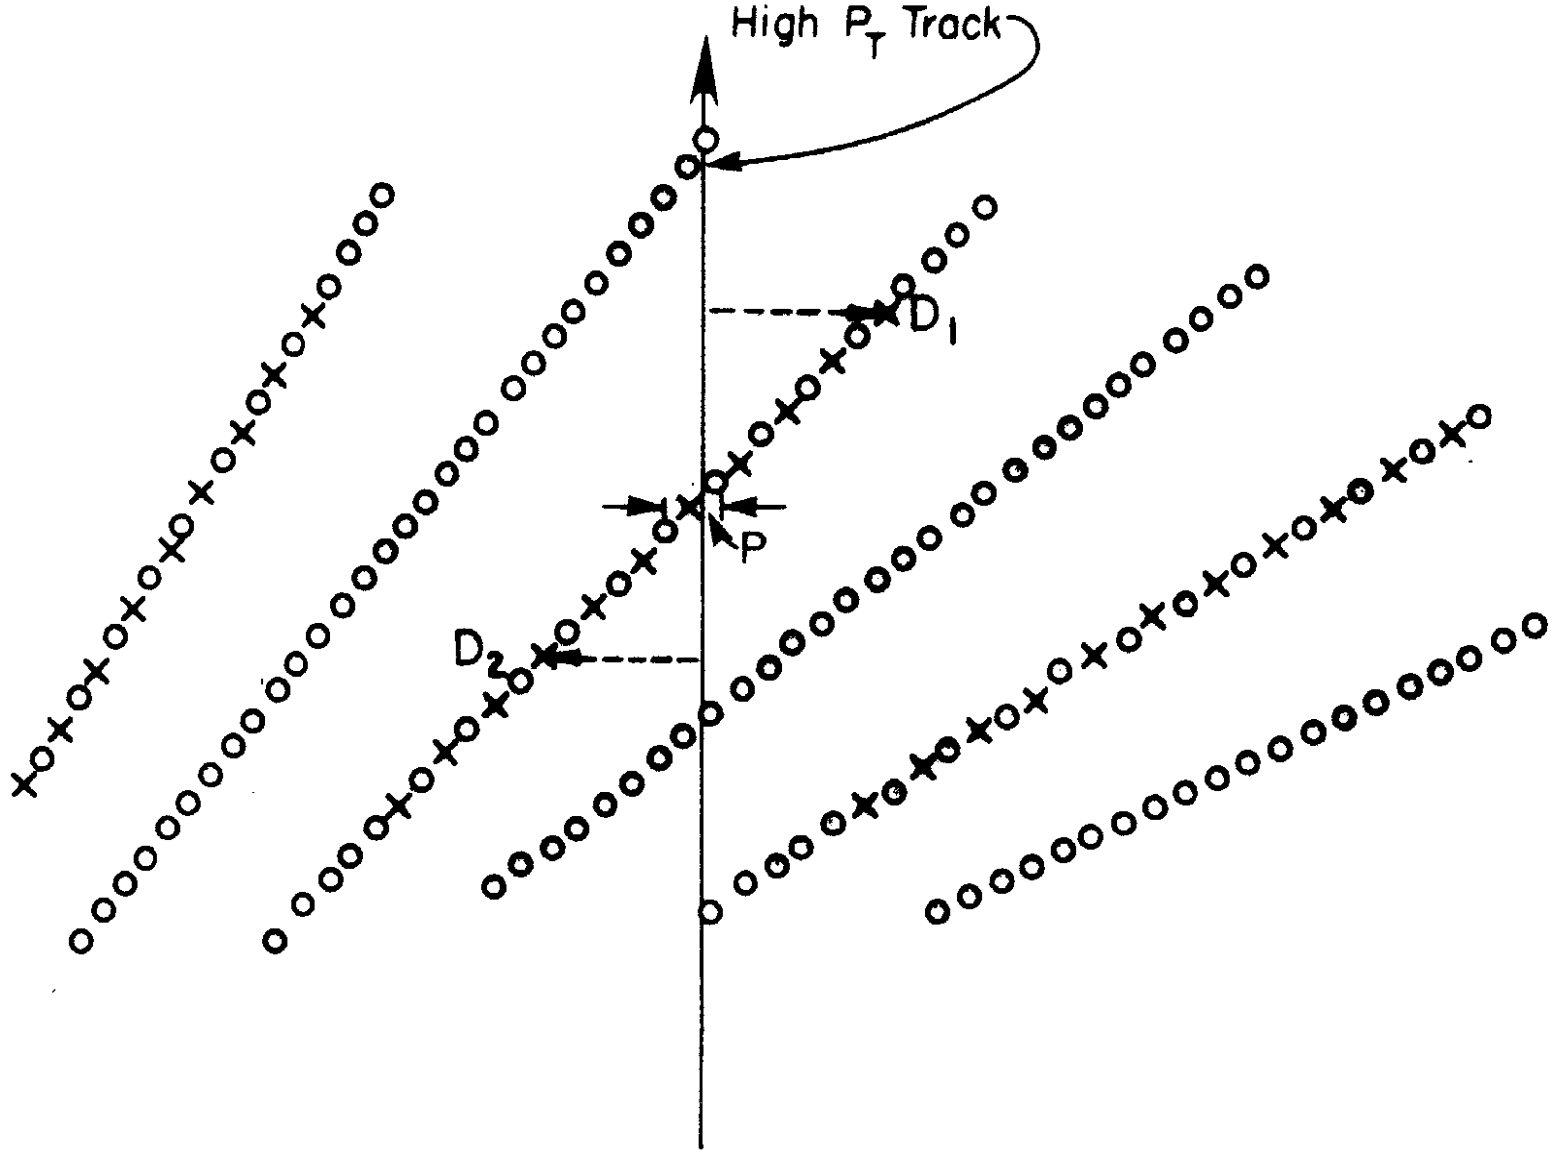

FIG 1

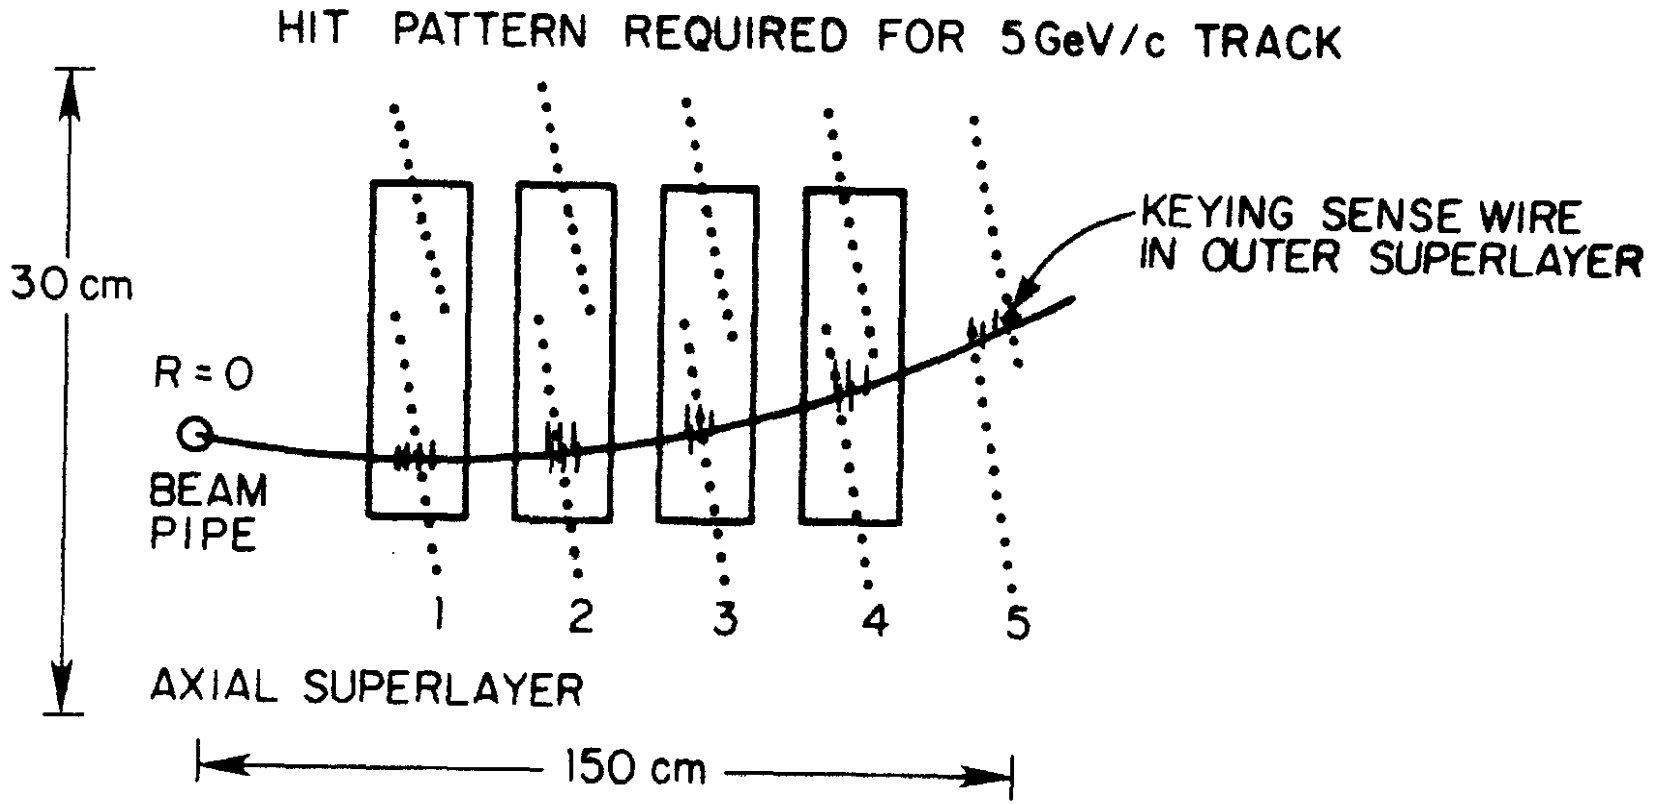
HIT PATTERN REQUIRED FOR 5 GeV/c TRACK
30 cm
R = 0
BEAM PIPE
KEYING SENSE WIRE IN OUTER SUPERLAYER
1
2
3
4
5
AXIAL SUPERLAYER
150 cm

Fig 2

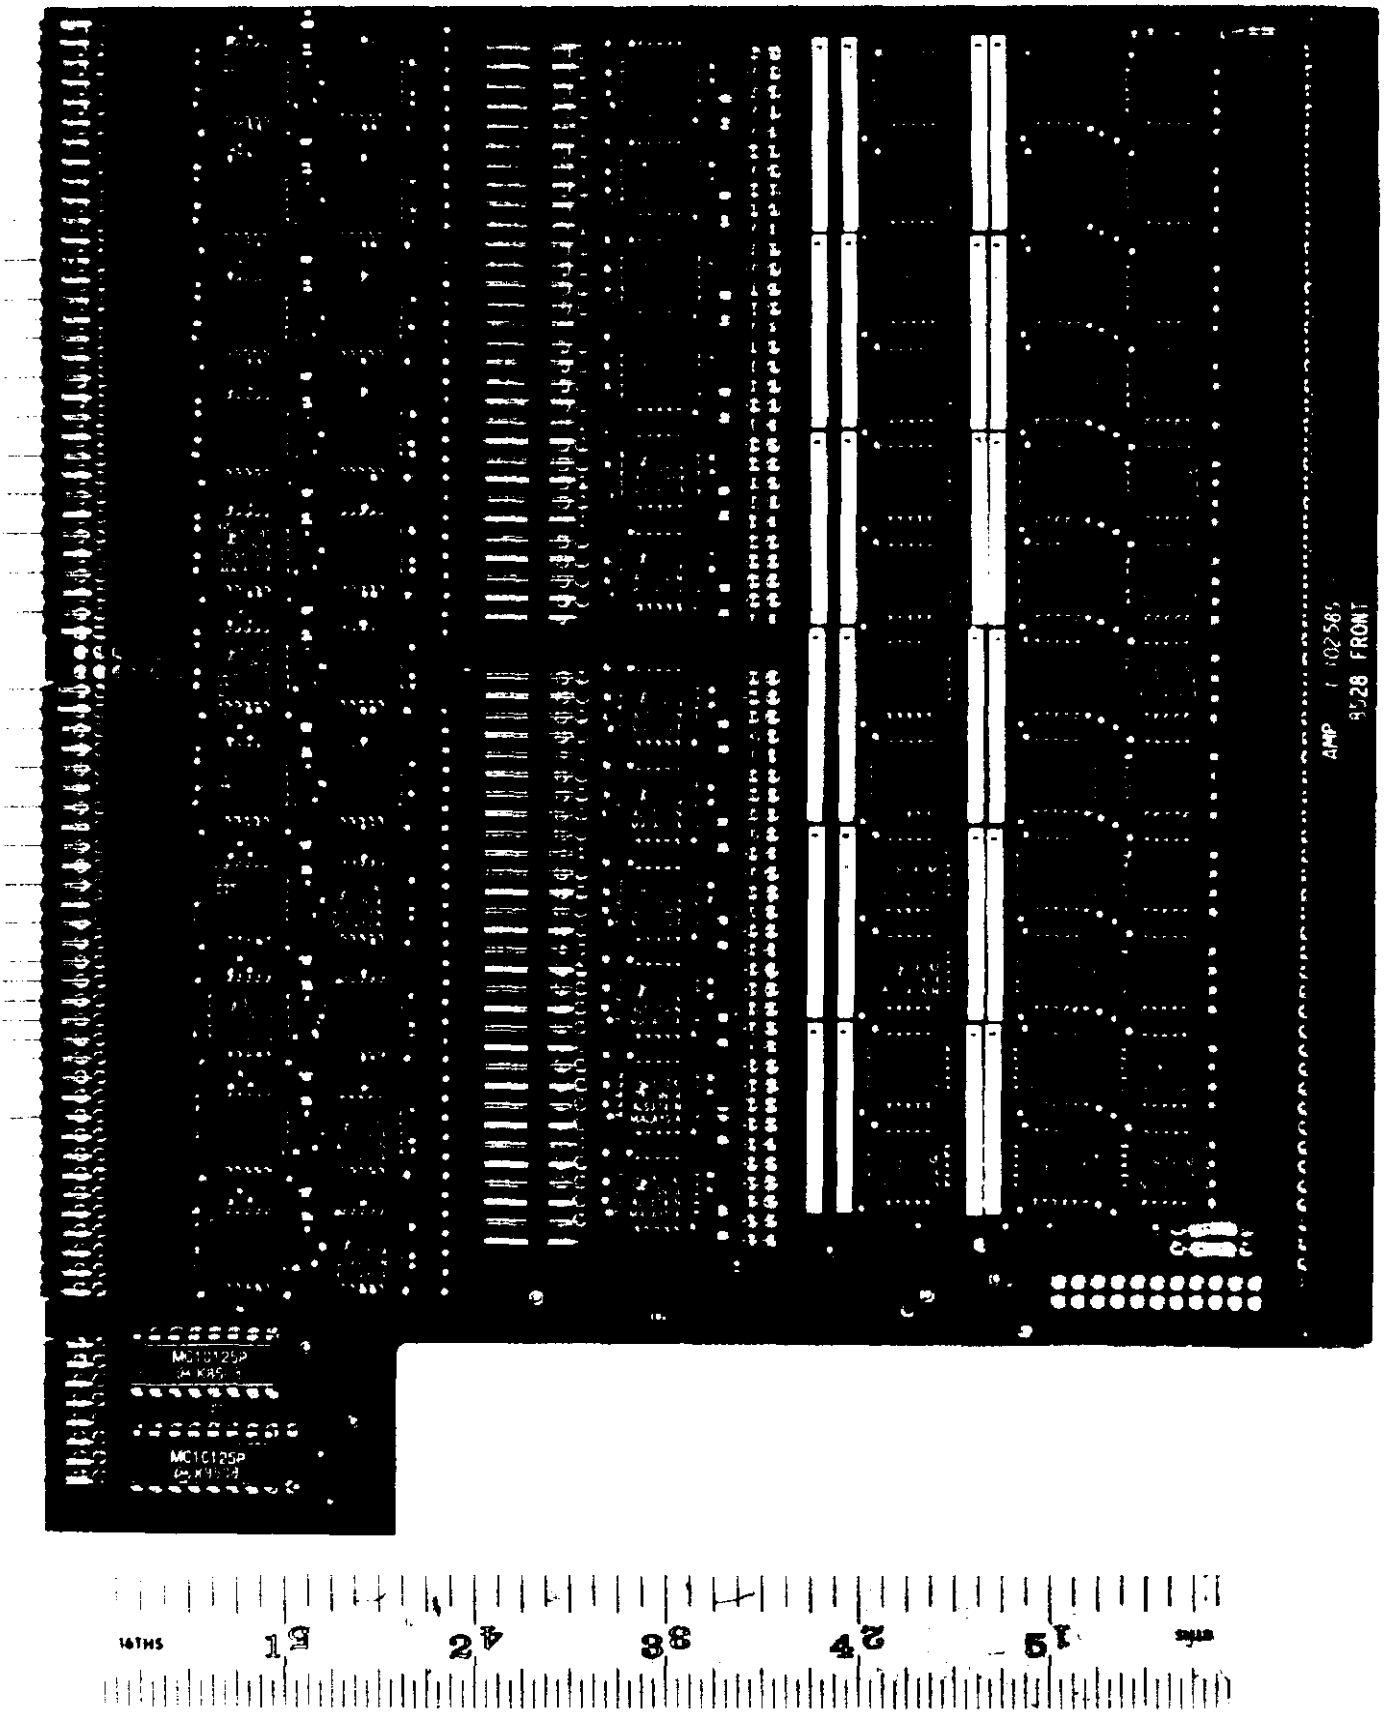

FIG 3

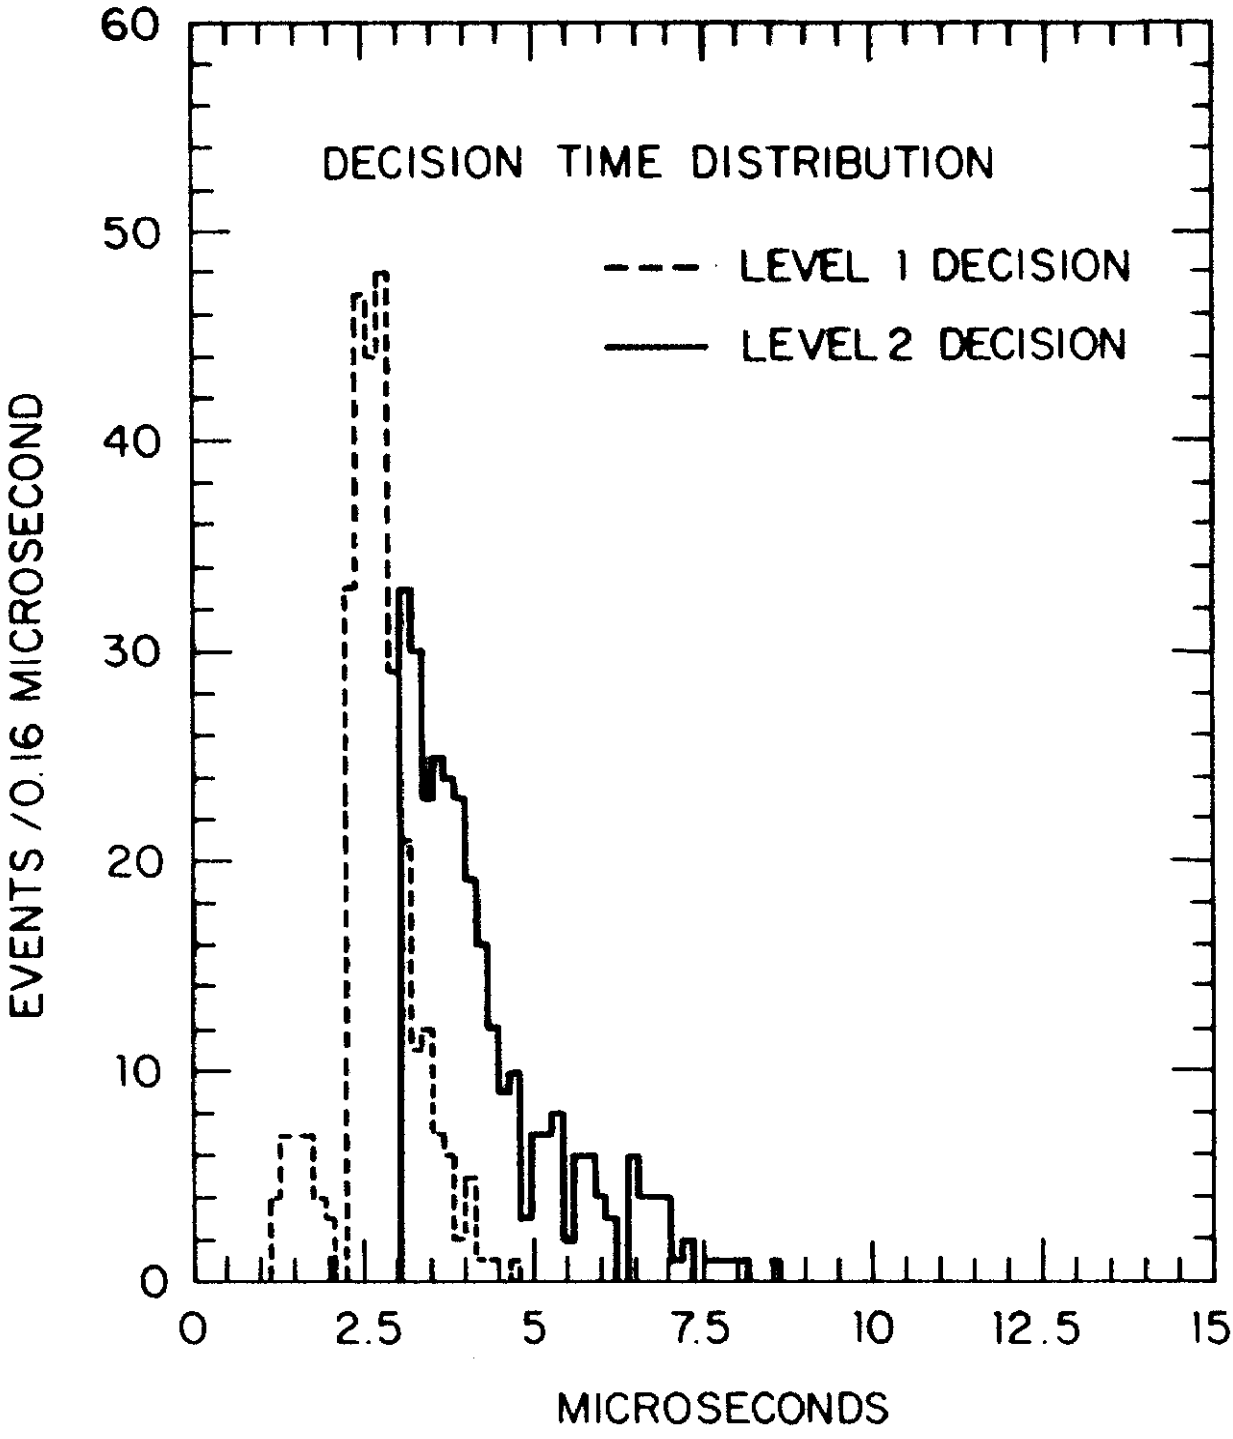
DECISION TIME DISTRIBUTION
LEVEL 1 DECISION
LEVEL 2 DECISION
EVENTS /0.16 MICROSECOND
MICROSECONDS
0
10
20
30
40
50
60
0
2.5
5
7.5
10
12.5
15

FIG 4

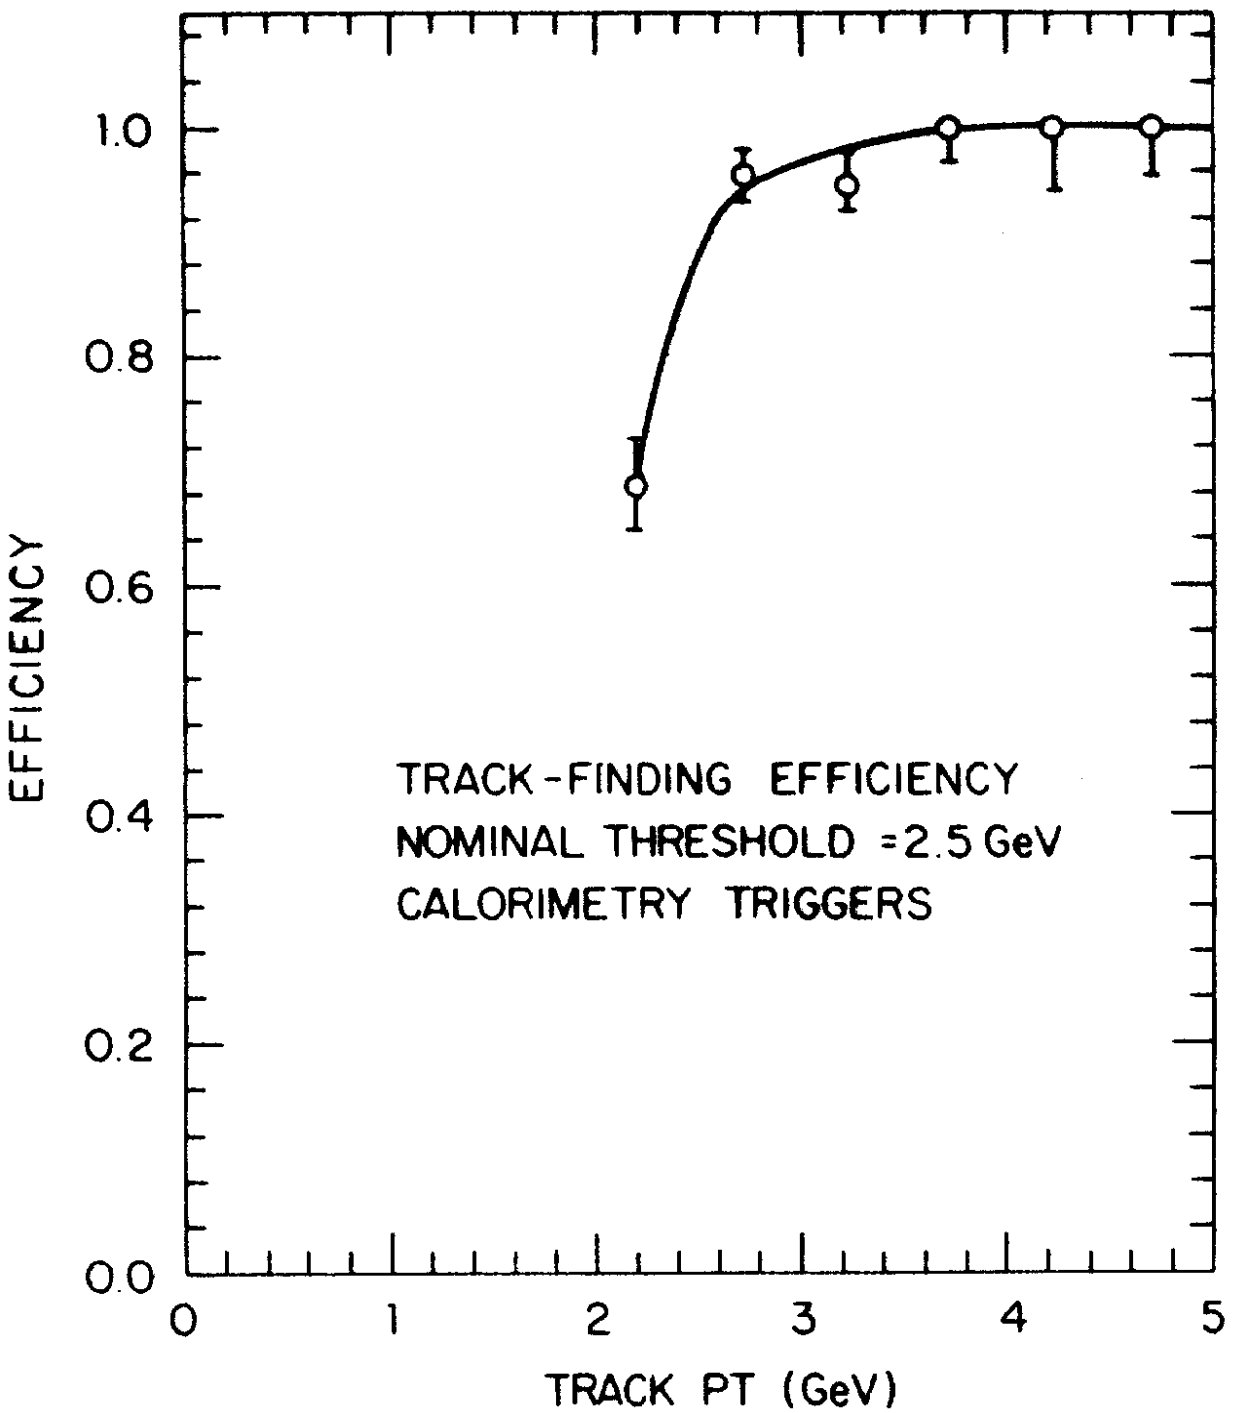
EFFICIENCY
1.0
0.8
0.6
0.4
0.2
0.0
0
1
2
3
4
5
TRACK PT (GeV)
TRACK-FINDING EFFICIENCY
NOMINAL THRESHOLD =2.5 GeV
CALORIMETRY TRIGGERS

Fig 5

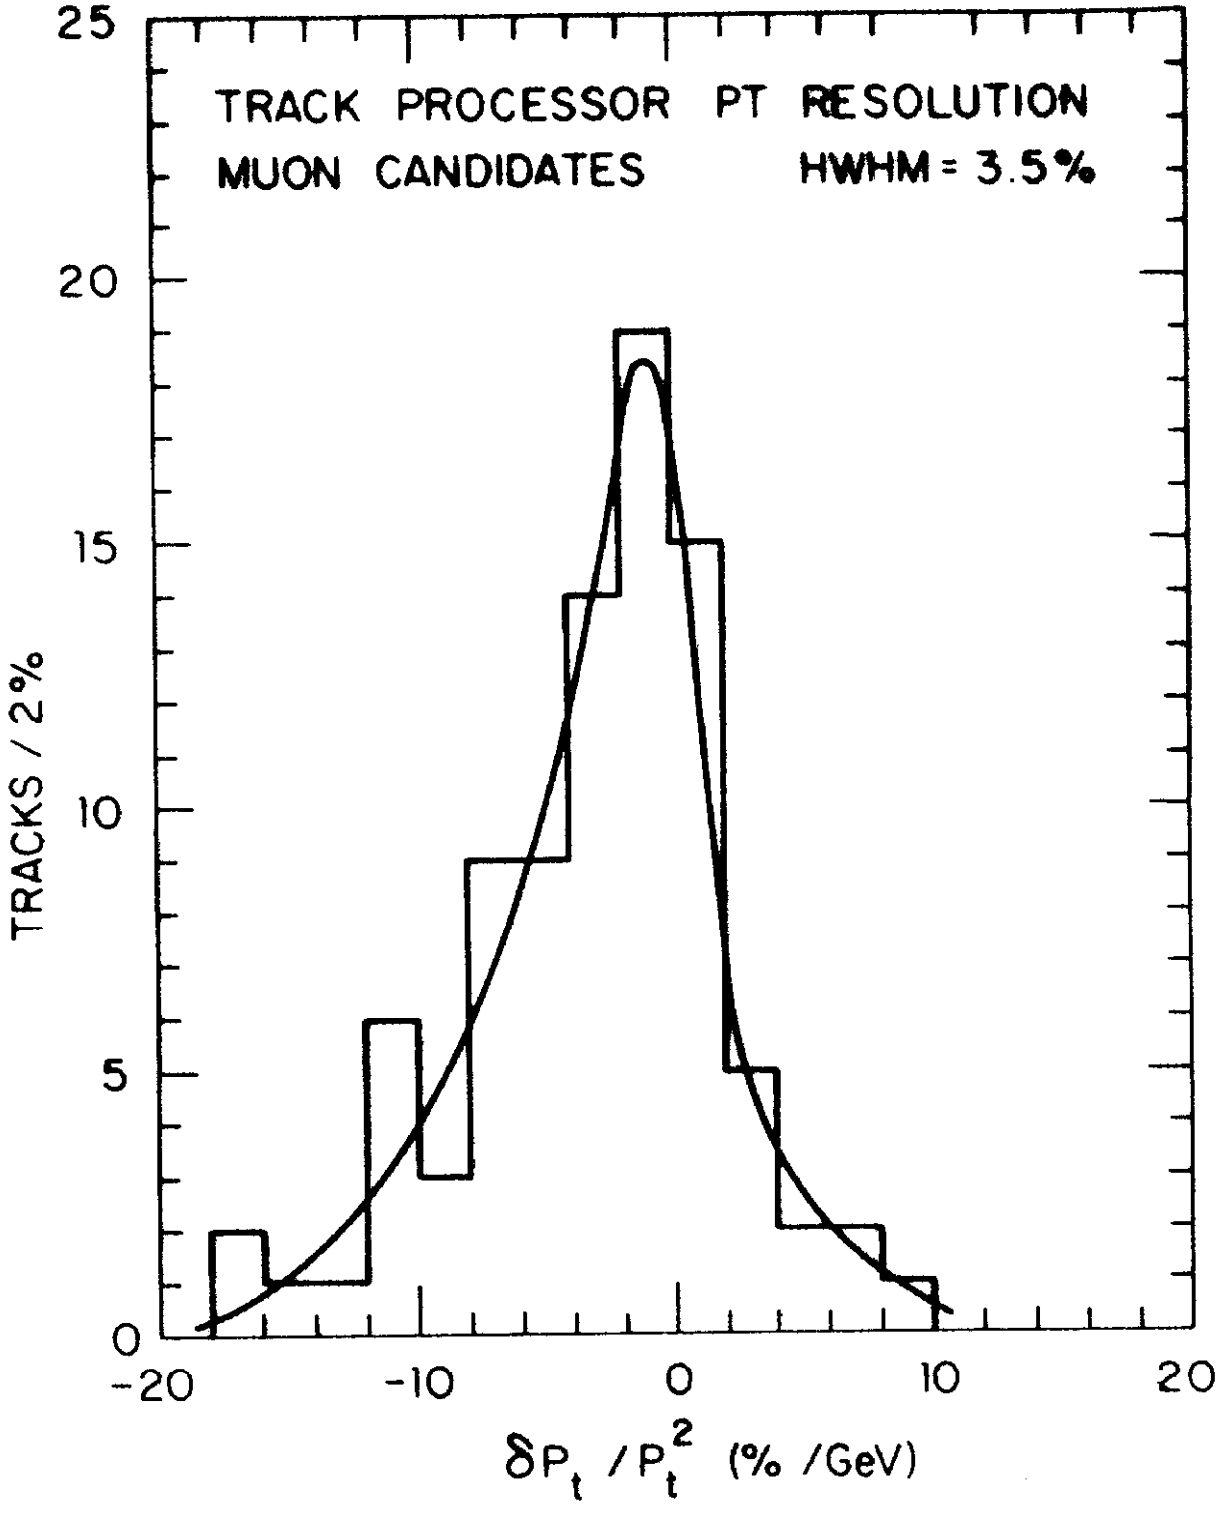
TRACK PROCESSOR PT RESOLUTION
MUON CANDIDATES
HWHM = 3.5%
TRACKS / 2%
25
20
15
10
5
0
-20
-10
0
10
20
δP_t / P_t^2 (% /GeV)

FIG 6

Table 1: Nineteen Stage Digital Pipeline for CTC Track Processor

PROCESSING STAGE	HARDWARE	FUNCTION & DESCRIPTION
1. Wire Hit Coincidence Latches	Open-collector ALS TTL octal buffers used as zero-catching latches.	Coincidence latches used to identify wire hits. The latches record a hit if the chamber signal is over threshold at any time during a "prompt" coincidence gate (0 to 80 ns following the beam crossing). These bits are then parallel-loaded into the readout shift register (stage 2 below) and the latches then record hits in a second "delayed" coincidence gate 500-650 ns following the beam crossing.
2. Readout Shift Register	ALS TTL octal registers	A 96-bit wide shift register to provide ϕ-ordered readout of each superlayer in the chamber. Each stage resides on a single card on the Auxiliary Fastbus Connector (AFC) associated with each 96-channel TDC used to record the chamber data. The length of the shift register is different in each superlayer to account for the different number of channels at different radii of the chamber. Each half of each register is clocked independently so that the end of each shift register contains data for the ϕ currently being processed.

3.-5. Outer Superlayer Hit Bit Encoder	PLA's, AS TTL, Advanced CMOS	Three-stage pipelined encoder to produce a list of hits in the outer superlayer of the chamber. The search for "roads" leading to these hits will be keyed on this list. This wire number list is recorded in RAM and "played back" during the second pass of "delayed" hit processing.
6. Extractor Control RAM Lookup	High Speed CMOS RAM's	RAM lookup tables which convert the wire number of an outer superlayer hit into control information necessary to "extract" the bits from the inner superlayer shift registers for the ϕ being processed.
7. Fanout of Control and Wire Number Cables	Terminated Flat Cables	Fanout of control signals to Bit Field Extractor cards on each of five superlayers, and of "Wire Number" cables to each of 8 Subroad cards and to the final Decision card.
8a-9a Bit Field Extractors	AS TTL registers and barrel shifter I.C.'s	The Bit Field Extractor cards use the control signals from stage 6 (above) to: 1) clock the readout shift register for its superlayer so that the head of the shift register contains the 96 bits appropriate for the region of ϕ being processed; and 2) select, from these 96 bits, the 16 bits which contain the envelope of all high-P_t roads leading to a hit in the outer layer (see Fig. 2).

10a. Transmission of Extracted Bit Fields to Subroad Cards	Terminated Flat Cable	The extracted bit fields from all superlayers are transported 2-3 m. to the crate containing the Subroad cards which search these bits for hit patterns from high P_t tracks.
8b-10b. Subroad RAM Table Lookup	High Speed CMOS RAM's	The subroad RAM tables (~640 kbytes) contain the hit patterns or "roads" expected for all possible high-P_t tracks leading to a hit in the outer superlayer. Hit positions are expressed as a 4-bit offset into the 16-bit field extracted in 8a. above. The RAM patterns also hold a tag bit which indicates whether the expected hit pattern is 1 or 2 bits wide.
11. Hit Bit Coincidence Mask	TTL Multiplexors	The extracted bit fields are searched for the bit patterns stored in the subroad RAM's. The hit patterns for 32 roads in each of five superlayers are checked in parallel on each clock cycle. In cases where the expected pattern is 2 bits wide, a hit in either bit position will cause the coincidence mask to fire.

12. Coincidence multiplicity calculation	High Speed CMOS RAM's	The 10 bits of hit/miss information from each road address a RAM which produces a 2-bit output majority code corresponding to zero, one, two, or many misses. Ten bits are available from five superlayers since there can be prompt hits from as many as two sense plane crossings per superlayer, and a track will typically generate one delayed hit for each drift direction.
13. Transfer of Majority Codes to Decision Card	Flat Cables on AFC Backplane	The 2-bit majority codes from each of 32 road circuits are transported to the Decision card which will choose among the track candidates and form a trigger decision.
14. Transfer of Majority Codes to/from Storage RAM's	High Speed CMOS RAM's	During the first pass of processing (prompt hits), the majority codes are written into a buffer RAM. This RAM is played back during the second pass (delayed hits) so that the total number of misses (prompt + delayed) is available to evaluate the track candidate.
15. Prompt plus Delayed total miss decoding	PLA's	Decoded outputs are produced representing 0,1,2,or 3 total (prompt plus delayed) misses on each of the 32 roads.

16. Priority Encoding of Best Track Candidate	AS TTL Priority Encoders	The "best" track candidate on each clock cycle is chosen from among the 32 roads as: 1) the road with the fewest number of misses, or 2) the road with highest P_t among those with an equal number of misses.
17. Comparison with track from previous clock cycle	High Speed CMOS RAM's	A four-bit code is produced which compares the quality of each track with those found on the previous clock cycle. The code indicates if: 1) the coincidence/P_t code is better 2) the curvatures indicates that it could be the same track 3) the outer layer hits leading to both tracks were at adjacent ϕ's 4) the pipeline contains a valid wire # (i.e. the processor was not marking time due to the absence of an outer-superlayer hit)
18. Identify Local Maxima of Track Goodness	High Speed CMOS RAM's	The four-bit comparison codes from three successive clock periods address a RAM which is programmed to identify peaks in the "goodness" of the track candidates, and thereby eliminate cases where the same track is found on successive clock cycles.

19. Track List FIFO Write and Trigger Output F/F	35 MHz TTL FIFO's	Track candidates surviving through stage 18 of the pipeline are written into FIFO's for transmission to the muon trigger and readout into the event record. On the first pass (prompt processing) the presence of a track above a programmable threshold sets a flip/flop which generates a Level 1 trigger.