Online Track Processor for the CDF Upgrade

E.J. Thomson, C. Ciobanu, J.Y. Chung, J. Gerstenslager, J. Hoftiezer, R.E. Hughes, M. Johnson, P.Koehn, C. Neu, C. Sanchez, B.L. Winer, J. Dittmann, J. Freeman, S. Holm, J.D. Lewis, C.J. Lin, T. Shaw, T. Wesson, K. Bloom, D. Gerdes, N. Goldschmidt, J. Dawson, W. Haberichter .

Abstract— **A trigger track processor, called the eXtremely Fast Tracker (XFT), has been designed for the CDF upgrade. This processor identifies high transverse momentum (> 1.5 GeV/c) charged particles in the new central outer tracking chamber for CDF II. The XFT design is highly parallel to handle the input rate of 183 Gbits/s and output rate of 44 Gbits/s. The processor is pipelined and reports the result for a new event every 132 ns. The processor uses three stages: hit classification, segment finding, and segment linking. The pattern recognition algorithms for the three stages are implemented in programmable logic devices (PLDs) which allow in-situ modification of the algorithm at any time. The PLDs reside on three different types of modules. The complete system has been installed and commissioned at CDF II. An overview of the track processor and performance in CDF Run II are presented.**

Keywords— **CDF, XFT, online track processor,trigger.**

I. Introduction

The CDF collaboration has upgraded the CDF detector for the next $p\bar{p}$ collider Run II, which started in March 2001. The upgrades are substantial, including a complete replacement of the charged particle tracking detectors, extension of the muon coverage, replacement of the plug and forward calorimetry, and almost complete replacement of the data aquisition system. The purpose of these upgrades is to both handle and take advantage of the new Tevatron running conditions planned for Run II. The Tevatron upgrades will bring approximately an order of magnitude increase in the instantaneous luminosity, and the detector upgrades will yield a large increase in acceptance, particularly for those analyses dependent on central tracking. This will greatly extend the physics reach for CDF, providing an unique opportunity to probe the standard model in great detail.

The eXtremely Fast Tracker (XFT), is a trigger track processor that identifies charged tracks in CDF's new Central Outer Tracker. The tracks are found in time for the Level 1 trigger decision, and are used for online identification of electrons and muons. This trigger is at the heart of much of the physics that CDF hopes to do in Run II. The XFT is needed to identify high momentum leptons for top, electro-weak, and exotic physics. It is also needed to identify low-momentum charged tracks for B physics analyses.

Manuscript received November 26, 2001. E. Thomson, C. Ciobanu, J. Chung, J. Gerstenslager, J. Hoftiezer, R. Hughes, P. Koehn, M. Johnson, C. Neu, C. Sanchez and B. Winer are with The Ohio State University, Columbus, Ohio. J. Dittmann, J. Freeman, S. Holm, J. Lewis, C. Lin, T. Shaw and T. Wesson are with Fermi National Accelerator Laboratory, Batavia, Illinois. K. Bloom, D. Gerdes and N. Goldschmidt are with University of Michigan, Ann Arbor, Michigan. J. Dawson and W. Haberichter are with Argonne National Laboratory, Argonne, Illinois.

II. CDF II Upgrade

The CDF goal for Tevatron Run II is to accumulate 2 fb^{-1} of integrated luminosity at $\sqrt{s} = 2.0$ TeV, using instantaneous luminosities up to $2 \times 10^{32} cm^{-2} s^{-1}$. This represents a factor of 20 increase in data sample size from the 0.109 fb^{-1} accumulated by CDF during Run I. The instantaneous luminosity increase is a factor of 10 over the maximum experienced during Run I and is made possible by both the new Main Injector, as well as the anti-proton recycler [1], [2]. To keep the number of overlapping interactions per crossing at a manageable level, the Tevatron will increase the number of proton and antiproton bunches in the machine to 108 (from 6 used in Run I), while decreasing the time between bunches to 132 ns (from 3.5 μs used in Run I). This new time structure requires major changes to the data aquisition and trigger systems, as well as replacement of the central tracking chamber. The upgrades to CDF are extensive, and are documented in detail elsewhere [3]. The specific parts of the upgrade relevant for this paper are listed below.

A. *The Central Outer Tracker (COT)*

The COT is an open cell drift chamber for charged particle reconstruction, occupying the radial region from 44 to 132 cm. The COT replaces the Central Tracking Chamber(CTC), which in addition to aging problems observed during Run I, would also suffer from degraded performance at $\mathcal{L} \geq 1 \times 10^{32}\ cm^{-2} s^{-1}$. The basic problem with the CTC is its maximum drift time (800 ns) relative to the expected bunch crossing time in Run II (132 ns).

To address this, the COT uses small drift cells ($\sim$ 2 cm wide – a factor of 4 smaller than the CTC) and a fast gas to limit drift times to less than 130 ns. Each cell consists of 12 sense wires oriented in a plane, tilted at approximately 35° with respect to the radial. A group of such cells at a given radius is called a superlayer. There are 8 alternating superlayers of stereo (nominal angle of 3°) and axial wire planes.

B. *Trigger and Electronics*

Just as in Run I, CDF will employ a 3 Level trigger system in selecting events to write out to mass storage. Each successive level in the trigger applies more stringent selection criteria, in order to eventually reach the Level 3 output rate of < 50 Hz. With only 132 ns between bunches, the Level 1 trigger electronics needs to be pipelined, so that there is enough time to make a decision on whether to pass the event onto Level 2. In addition, due to the factor of 10 increase in luminosity expected in Run II, the rejection rate of the trigger will also have to increase by almost a factor of 10.

III. The XFT Track Processor

In Run I, tracks were identified in the CTC in time for the Level 2 decision by the Central Fast Tracker (CFT) [4]. The processing time for the CFT was approximately 10 μs, and was strongly dependent on the complexity of the event. Due to the much smaller bunch crossing time, and the need to move track identification into the Level 1 trigger decision (more on this below), a new track processor, the eXtremely Fast Tracker (XFT), was needed [5]-[7].

A. Role of the XFT in the Trigger

The purpose of the full trigger system is to identify the subset of "interesting" physics events, reducing the raw collision rate from 7.5 MHz to a rate at which events can be stored for further offline analysis. The XFT plays a major role in this reduction.

The steps from a $p\bar{p}$ crossing to writing events to offline storage can be summarized as follows:

1. The bunch crossing rate is 7.5 MHz. This corresponds to a $p\bar{p}$ crossing every 132 ns. For an instantaneous luminosity of 2.0×10^{-32} $cm^{-2}s^{-1}$ one expects an average of 2 $p\bar{p}$ interactions every crossing.
2. The L1 trigger needs to reduce this rate to less than 50kHz. This assumes a Level 2 processing rate of 20 μs, and results in a CDF deadtime of $\leq$ 10 %. The primary purpose of the XFT is to identify tracks in the COT in time for the Level 1 decision. The XFT tracks at Level 1 are matched to electromagnetic-calorimeter clusters for electron identification, to stubs in the muon system for muon identification. The tracks are also used in a two-track trigger for events such as $B \rightarrow \pi^+\pi^-$.
3. The Level 2 trigger needs to reduce the Level 1 rate down to $\leq$ 300 Hz. This rate is set by the processing capability expected for the Level 3 trigger. To help accomplish this rate reduction, the XFT tracks will be linked to hits in the silicon vertex detector by a separate device called the silicon vertex trigger (SVT) processor. This will result in an improved determination of the track momenta, as well as identification of tracks displaced from the primary vertex, a first for a hadron-collider experiment.
4. Finally the Level 3 trigger reduces that rate to $\leq$ 50 Hz, at which point accepted data are transferred to permanent storage media.

B. Requirements on XFT Performance

Based on these considerations, the following design specifications were set [3]:

- XFT will need to be pipelined, and will need to present a new set of found tracks every 132 ns. The results must be ready in time for the Level 1 trigger decision, which means the XFT will need to find tracks in less than 1.5 μs.
- For physics analyses such as top, electroweak studies, and searches for new phenomena, high efficiency is critical. We set a goal that the XFT track-finding efficiency will be greater than 96% when the single-hit efficiency of the central tracker is greater than 92%.
- The momentum resolution reported to the trigger will be $\Delta P_T/P_T^2 < 2.0\%/GeV/c$. In Run I, the track processor momentum resolution was $\Delta P_T/P_T^2 \sim 3.5\%/GeV/c$. Better momentum resolution allows a lower threshold at high P_T, which in turn increases acceptance.
- The resolution on the extrapolated ϕ position at the origin (σ_{ϕ_0}) should be better than 8 mrad. This is set by the requirement that XFT tracks will be used in the Level 2 trigger and attached to hits in the silicon vertex detector.
- The fraction of tracks found by the XFT which are not associated with a real track (i.e. fakes) should be less than 10% for $P_T \geq 10$ GeV/c. This fake-track rejection is also needed to reduce the high-P_T muon trigger rate at high luminosity.
- The minimum track P_T will be 1.5 GeV/c. In Run I the minimum P_T was 2.2 GeV/c. The cut-off at 1.5 GeV/c is set by the fact that muons are stopped in the calorimeter for track momenta below this. In addition, a lower P_T cutoff allows greater acceptance for B decays that can be used to measure CP violation.

C. XFT Algorithm Overview

The XFT processor uses hit data from the 4 axial superlayers of the COT. As mentioned earlier, the superlayers are arranged in cells of 12 wires each, oriented at an angle of $\sim 30°$ relative to the radial. There are a total of 16,128 axial wires, and the data on each wire are classified as prompt and/or delayed, for a total of 32,256 bits of information. A charged track passing through an axial superlayer will generate 12 "hits" of prompt and/or delayed data. The definition of a prompt or a delayed hit will depend upon the maximum drift in the COT. For a bunch spacing of 132 ns, the maximum drift is $\sim$ 121 ns, and so a prompt hit occurs whenever there is a hit in the time window 0-44 ns, and a delayed hit is defined as a hit falling in the window 44-121 ns.

Track identification is accomplished in two processes by the Finder and the Linker. The Finder searches for high-P_T track segments in each of the four axial superlayers of the Central Tracker. The Linker searches for a four-out-of-four match among segments in the 4 layers, consistent with a prompt high-P_T track. If no track is found, the Linker searches for a three-out-of-three match among segments in the innermost 3 layers.

D. The Finder

The Finder is designed to look for valid track segments in a given axial superlayer. To do this quickly, each COT superlayer is divided into groups of four cells each, with each group processed by a single Finder PLD. Within a given layer, the design and operation of the PLD are all identical. The PLDs compare hit data to a predefined set of patterns, and all of the patterns are searched over simultaneously.

A track segment is defined by the cell and whether a prompt or delayed hit was generated in each of the 12 wire planes within a superlayer. A collection of the cell numbers and hit types for the 12 wires in an axial layer is called a mask. An example of a mask is shown in Fig. 1. The mask will change depending on the ϕ of the track, and its angle through the cell (or P_T). The Finder works by storing all possible masks for tracks with $P_T \geq 1.5$ GeV/c in a database (or equivalently, programmed on a chip). The mask set is determined for a given axial superlayer by a Monte Carlo program. Since all of the cells in a given layer are identical, only one set of masks is needed for each layer. Due to the fact that for a given P_T, the local slope is greater in the outer layers than in the inner layers, the outer layers will require a larger mask set than the inner layers. The inner layers 1 and 2 and the outer layers 3 and 4 require 173, 242, 290, and 337 masks respectively. The Finder compares incoming TDC information with all masks for the given axial superlayer, allowing a programmable (up to 3) number of missed wire planes.

In the inner two axial layers, valid segments are characterized by 1 of 12 pixel positions across the midpoint of the cell. In the outer two axial layers, valid segments are characterized by 1 of 6 pixel positions across the midpoint of the cell, and 2 bits of slope information. These two bits are characterized as follows:

- 00 = no mask found
- 01 = negative low Pt track segment
- 10 = positive low Pt track segment
- 11 = high Pt track segment (nominal $P_T \geq 8$ GeV/c)

The pixel bin size is approximately 1.5mm in the inner two layers, and approximately 3.0mm in the outer two layers.

Every 132 ns the Finder outputs 12 bits per cell (12 ϕ pixels for the inner two layers, 6 ϕ pixels $\times$ 2 slopes for the outer two layers) to the Linker. The Finder outputs *all* pixels corresponding to a valid track pattern - not just the pixel position of the best segment.

To reduce the total number of Finders, each Finder chip identifies segments for 4 adjacent COT cells. There will be only one set of masks stored per layer, and the inputs from the 4 separate cells will be multiplexed. The total number of Finder chips needed is then 336. The Finders are implemented using in-system reprogrammable logic chips. The Finder system takes in 2 bits of information for each of 16,128 axial wires every 132 ns, and output a total of 12 bits for each of 1,344 cells every 132 ns.

E. The Linker

The Linker is designed to look for valid tracks which cross either 3 or 4 axial superlayers. The input is the pixel and slope information transferred from the Finders. To locate tracks quickly, the COT is divided into 288 identical phi-slices, each of which is processed by a single Linker chip. The Linker chips are also implemented using in-system reprogrammable logic chips. The chips compare pixel and slope data to a predefined set of patterns, and all of the patterns are searched over simultaneously.

Each Linker chip is given all of the pixel information needed to find the tracks in a $\Delta\phi = 1.25°$ phi-slice of the tracking chamber. Due to track curvature, a significant fraction of the pixels needed to identify all tracks that key off the 3rd axial layer come from *outside* of the $\Delta\phi = 1.25°$ phi-slice. This is shown in Fig. 2. The Linker algorithm begins by searching (in parallel) a list of about 2400 roads, where a road is a group of 4 pixels, one from each axial superlayer, corresponding to a valid track with $P_T \geq 1.5$ GeV/c. The pixels in the outer 2 axial layers are required to have the same sign of P_T as the Linker track. The roads are defined by their P_T and the pixel position in layer 3. The roads are then combined with a logical OR, resulting in 96 P_T bins and 8 pixel locations. Found tracks are passed through a priority encoder to find the best track in the 1.25° region covered by the Linker. The information reported on the best track is 7 bits of P_T, 3 bits of ϕ, 1 bit indicating whether the track is isolated, and 1 bit indicating whether the track used 3 or 4 layers. If no track is found using all 4 layers, then the best track found in the innermost 3 layers is output. This allows a small increase in the rapidity coverage of the XFT.

F. XFT System Hardware

F.1 Overview

The system hardware, which is illustrated in Fig. 3, begins with the XTC mezzanine module residing on the COT TDC's in the CDF collision hall. This module classifies the hits on the COT wires as prompt or delayed and sends that information to a transition module at the back of the COT TDC crate. The transition module drives the data at 45.5 MHz, with Low Voltage Differential Signal (LVDS) technology, onto Level 1 Trigger cables that carry the COT wire data 220 ft to the Finder module crates. Finder transition modules receive the data and send it across a customized backplane to the Finder modules. Finder modules find track segments and report them to the Linker modules. Transition modules drive found track information from the Linker modules to the XTRP system, which extrapolates the tracks to the calorimeter and the muon chambers, applies trigger momentum thresholds and distributes the track data to the Level 1 and Level 2 trigger systems. The XFT data arrives at the XTRP 1.9 μs after the collision occurred, which is well within the 5.5 μs available for the Level 1 trigger decision. Listed below are the essential elements that compose the Finder and Linker modules.

F.2 Finder Modules

The Finder portion of the XFT system consists of two types of modules: the SL13 for COT axial superlayers 1 and 3, and the SL24 for superlayers 2 and 4. Each type of Finder module searches for track segments in a 15° region of ϕ. There are 24 SL13 and 24 SL24 modules in three 9U $\times$ 400mm VME [8] crates (eight of each module type per crate). Each crate has a commercial VME-based processor that controls communication via VMEbus. All Finder modules have Input Alignment PLDs that latch and align the COT wire data received from the Level 1 Trigger cables and send it to the Finder PLDs at 30.3 MHz. SL13 modules contain 2 superlayer 1 PLDs and 4 superlayer 3

PLDs. SL24 modules contain 3 superlayer 2 PLDs and 5 superlayer 4 PLDs. In addition to the track segment finding algorithm, the Finder PLDs hold the Level 1 pipeline and the Level 2 data buffers for COT hit information. The Finder PLDs report the found pixel data to a set of PLD's called the Pixel Drivers that duplicate this information, and send it to LVDS drivers. The data is driven over 10 ft of cabling to the Linker modules at 30.3 MHz. The Pixel PLDs also hold the Level 1 pipeline and the Level 2 data buffers for the found pixel information. A dedicated PLD on each Finder module provides the VMEbus slave interface. Clock circuitry generates on board 33 ns, 66 ns, and 132 ns clocks derived from the Master CDF clock feeding silicon delays. The Finder module also contains: RAM for loading PLD designs, circuitry for allowing boundary scans of all PLDs, and ports for loading PLDs from a serial port of a PC. The time between the rising edge of a data bit at the input of the LVDS drivers on the COT transition module to the rising edge of the data bit at the output of the channel link receiver on the linker module is about 875 ns. This measurement includes data propagation through the Level 1 trigger and the channel link cables. The finder module alone takes about 560 ns to process the data.

F.3 Linker Module

The Linker portion of the XFT system contains 24 9U modules in three crates (eight modules per crate). Each Linker module covers a 15° region of ϕ. Track segment information from the Finder modules are captured at LVDS receivers on the Linker module. Six Input Formatting PLDs latch the data from the Channel Link receivers and synchronize the data with the on board 33 ns clock. The 6th Input Formatter performs error checking at the input. There are 12 Linker PLDs that receive data at 30.3 MHz from the Input Formatters, and search for the best track. Each Linker PLD covers a region 1.25° in ϕ. The Linker PLDs transmit data at 7.6 MHz to 2 Output Formatter PLDs. The data is reformatted and passed to a transition module with LVDS drivers that sends the data over 50 ft of cabling to the XTRP system. A VME-Control PLD contains the functionality for the VME slave interface, a state machine that controls the response to trigger signals, and the loading of all the PLD designs on the module. As with the Finder, the Linker module also contains: on board clock generation from the Master CDF clock, RAM for loading PLDs, PLD boundary scan circuitry, and ports for external loading PLDs. The time between the rising edge of the data bit from the output of the channel link receiver on the linker module to the rising edge of the data bit from the output of the LVDS receiver on the XTRP transition module is about 820 ns. This measurement includes data propagation through the XTRP cables. The linker module alone takes about 730 ns to process the data.

IV. XFT Performance in CDF Run II data

CDF has accumulated 4 pb^{-1} since the start of Run II in March 2001. XFT was completely installed and operational before the start of Run II. XFT has been a very stable system, responsible for only 1 hour of downtime. The flexibility inherent in the design of XFT has been proven very important. The 2 miss Finder design has been used to obtain a satisfactory efficiency and fake rate. Since the beam position at CDF is offset by about 4 mm at an angle of 105^o, new Linker road designs were generated with this beam position. Important performance criteria for XFT include the efficiency for finding tracks, the measurement resolution on the transverse momentum and the direction of the tracks, and the fraction of fake tracks.

The performance of XFT is measured from a sample of events collected by a 10 GeV Jet trigger at Level 1. Tracks in this sample are recontructed by the offline tracking algorithm (without a beam constraint) and must satisfy the following requirements:

- Number of Axial hits used in fit > 24
- Number of Stereo hits used in fit > 24
- $p_T > 1.5$ GeV/c
- Pass through all 4 Axial superlayers
- Maximum p_T track in linker chip (1.25^o)

An offline track is defined to have been found by XFT if the distance between the extrapolated track positions is less than 10 pixels (about 1.5^o) in at least 3 of the axial superlayers. The efficiency for finding XFT tracks is measured to be $96.1 \pm 0.1\%$. Fig. 4 (a) shows the efficiency variation versus transverse momentum above the threshold of 1.5 GeV/c. Fig. 4 (b) shows that the efficiency is flat with respect to the azimuthal position of the offline track. XFT recovers the loss in efficiency from dead COT wires (there are 6 dead wires in one superlayer cell at 5 radians) by masking on these wires at the finder chip, without any significant increase in the number of fake tracks. There are less than 20 dead wires out of 16,128 axial superlayer wires.

The quality of the measurement by XFT of the transverse momentum and azimuthal position, with respect to the result of the full offline track reconstruction, is shown in Fig. 5. The momentum resolution is measured to be $1.74 \pm 0.01\%$ per GeV/c, better than the design specification of 2% per GeV/c. The azimuthal angular resolution is measured to be 5.45 ± 0.02 mR, better than the design specification of 8 mR.

The bias towards higher values for the transverse momentum is deliberate. The simplest case is where just 1 track has passed through a given linker chip. Since the Finder tends to report more than one pixel per layer even for isolated tracks, the Linker will tend to find multiple tracks, all in neighbouring p_T bins. The logic is designed to report the median p_T bin of this cluster of p_T bins. In addition, the logic must handle the case where more than one track has passed through a given Linker chip. The logic is designed to select the cluster of p_T bins which is higher in p_T and report the median of this cluster.

Fig. 6 show the efficiency for three different XFT momentum thresholds. The turn-on is very sharp at 1.5 GeV/c, which is important for the rate of B physics triggers based on combinations of tracks.

About 3% of XFT tracks at low transverse momenta are fake, where fake means no offline track was matched to a

XFT track. The effect of fake XFT tracks on the Level 1 trigger rate for electrons and muons with energies above 8 GeV/c is to increase the rate by about 6%.

V. Conclusions

In summary, Table I shows that the performance of XFT in Run II data meets or exceeds the design specification. The expectation from the simulation of the COT and the emulation of the XFT trigger is also shown.

TABLE I

XFT performance from design specification, simulation and Run II data.

Parameter	Specification	Simulation	Data
Efficiency	$\geq$ 96%	96%	96%
Percentage Fakes	$\leq$ 10%	1.2%	3%
Momentum Threshold	1.5 GeV/c	1.5 GeV/c	1.5 GeV/c
Momentum Resolution	$\leq 2.0 \frac{\%}{GeV/c}$	$1.0 \frac{\%}{GeV/c}$	$1.7 \frac{\%}{GeV/c}$
Angular Resolution	$\leq$ 8 mR	4 mR	5.5 mR

References

[1] K. Koba, "Status of the Fermilab Main Injector," in Proc. PAC-2001, vol. 1, pp 27-29.
[2] M. Hu, "The Fermilab Recycler Ring," in Proc. PAC-2001, vol. 1, pp.30-32.
[3] R. Blair *et al.*, "The CDF-II detector: Technical design report," FERMILAB-PUB-96-390-E.
[4] G.W. Foster, J. Freeman, C. Newman-Holmes, and J. Patrick, "A Fast Hardware Track-Finder for the CDF Central Tracking Chamber," Nucl. Instr. Meth. vol. A269, pp.93-116, 1988.
[5] P. Koehn *et al.*, "Online track processor for the CDF Upgrade," IEEE Trans. Nucl. Sci. vol. 46, pp. 933-939, 1999.
[6] S. Holm *et al.*, "System architecture and hardware design of the CDF XFT online track processor," IEEE Trans. Nucl. Sci. vol. 47, pp. 895-902, 2000.
[7] K. A. Bloom, "Fast track triggering for the CDF II detector," Int. J. Mod. Phys. A vol. 16S1C, pp. 1172-1174, 2001.
[8] VME International Standards Organization, Scottsdale, Az.

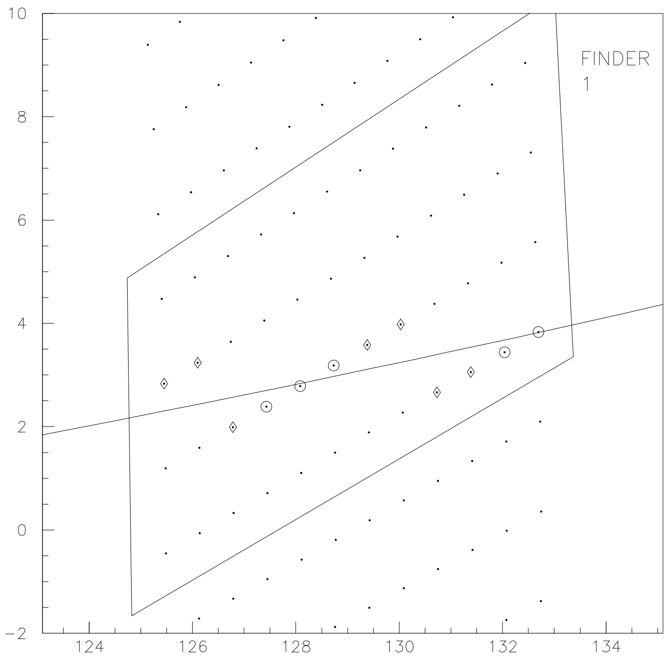

Fig. 1. A closeup of a track with $P_t = 1.5\ GeV/c$ in the 4th axial superlayer. The wires marked by the diamonds are those with "delayed" hits, while the wires marked with open circles are those with "prompt" hits. The collection of prompt and delayed hits on the given wires is an example of a mask.

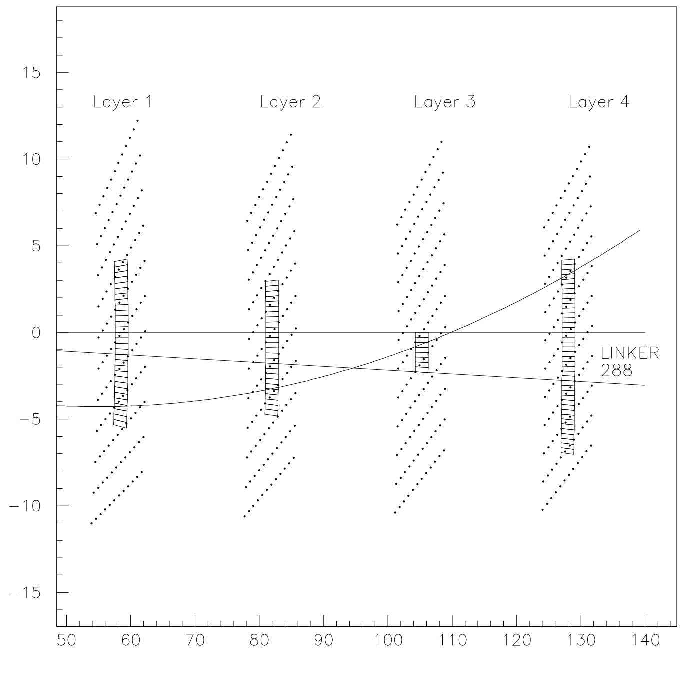

Fig. 2. A closeup of a track with $P_t = 1.5\ GeV/c$ showing the Linker the track should be found in and all of the pixels needed by this Linker.

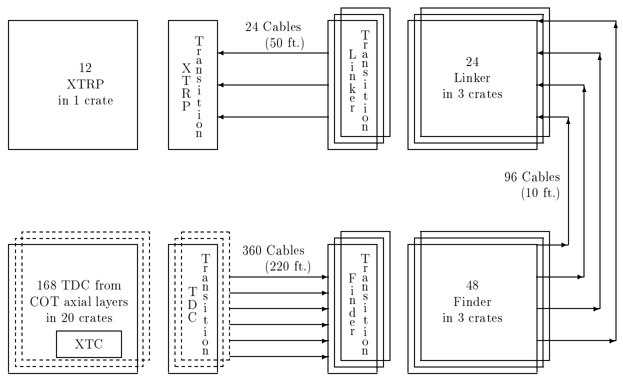

Fig. 3. Schematic of the XFT system

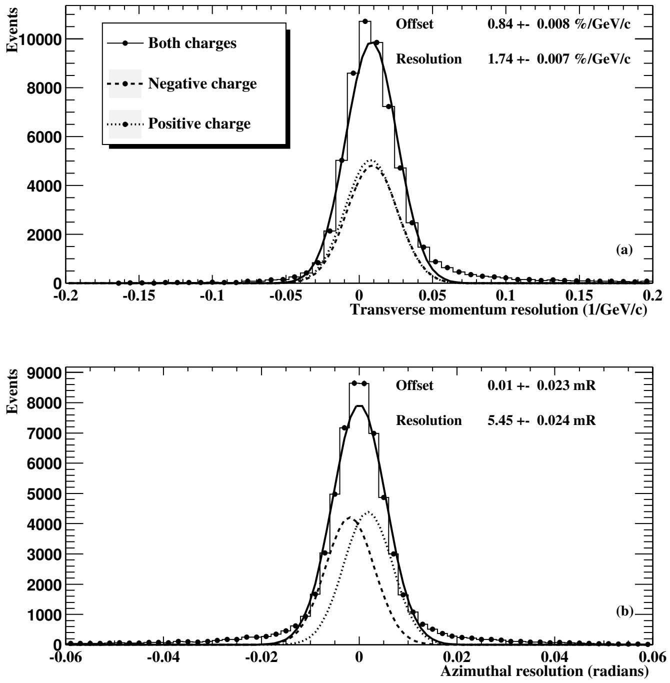

Fig. 5. (a) The transverse momentum and (b) the angular resolution of XFT tracks, with respect to tracks found by the full offline tracking algorithm.

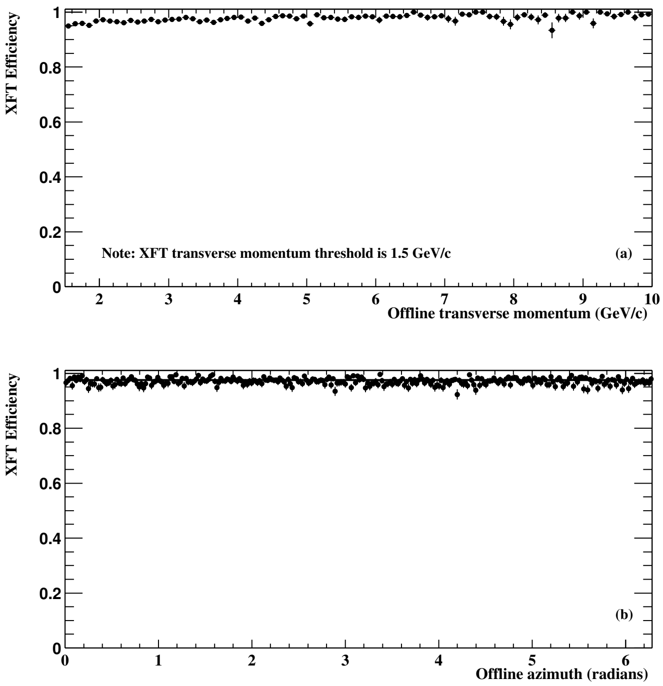

Fig. 4. XFT track-finding efficiency versus (a) the transverse momentum and (b) the azimuthal position from the offline tracking algorithm.

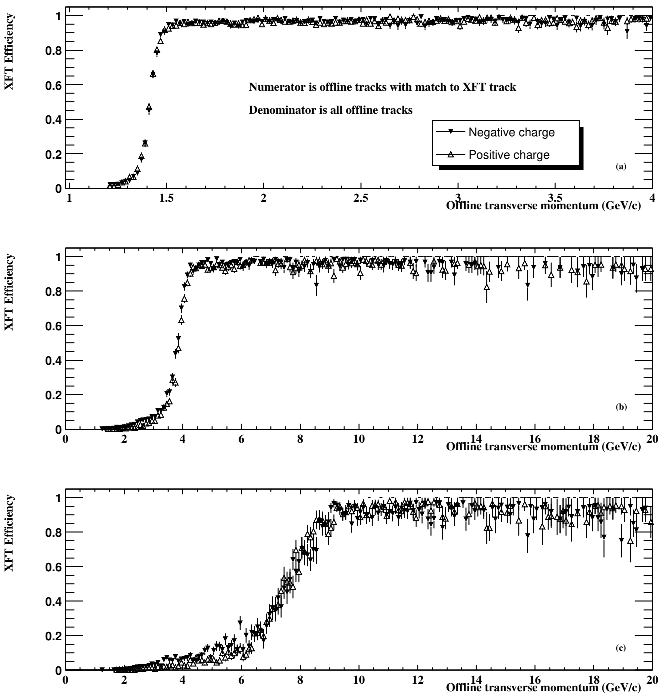

Fig. 6. Efficiency of XFT versus the offline transverse momentum for three nominal trigger thresholds. The XFT transverse momentum must be in or above the bin centred at (a) 1.50 GeV/c, (b) 4.09 GeV/c and (c) 8.35 GeV/c.