Simulation of the XFTII Trigger

Version 0.1

Richard Hughes, Phillip Koehn, Brian Winer
Ohio State University
Greg Veramendi
University of Rochester
June 16, 1997

1 Introduction

During Run I, CDF used an trigger track processor (the CFT) to identify high-momentum charged tracks in the Central Tracking Chamber (CTC). The tracks found were linked with clusters in the EM calorimeter and muon stubs in the muon chambers to identify electron and muon candidates. These were then used in the Level 2 trigger decision.

Since the CTC will be replaced by the COT for Run II, the CFT will also need to be replaced. In addition, the change in bunch timing for Run II (396 or 132 nsec) implies that this replacement must be pipelined. This new device, the XFT (eXtremely Fast Tracker), will be highly parallel and will process the data from each bunch crossing. The tracking results from the XFT will be available in time to be used in the Level 1 trigger decision.

The minimum design specifications for the XFT are set by the requirement that it perform at least as well as the CFT. The design specifications are:

- The 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 $\Delta P_T/P_T^2 < 2\%$.
- The resolution on ϕ_0 will be better than 8 mrad.
- The fake-track rejection will be at least twice as good as the CFT.
- The minimum track P_T will be 1.5 GeV/c.

The goal of high track efficiency is set by the desire to be efficient for high-P_T physics. Good momentum resolution allows a lower threshold at high P_T, particularly for muons, since after matching tracks to muon stubs, there is nothing but momentum to cut on. The improved fake-track rejection is also needed to reduce the high-P_T muon trigger rate at high luminosity. The lower P_T cutoff is set by the desire to maximize the acceptance for B decays that can be used, for example, to measure CP violation. The resolution on ϕ_0 is set by the requirement that XFT tracks will be used as the seed for the SVT at Level 2.

2 Algorithm Description

A block diagram of the XFT is shown in Figure 1. The processor works off of hit data from the 4 axial layers of the COT. 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 32,356 bits of information. The definition of prompt or delayed will depend upon the maximum drift in the COT. Assuming that this maximum is $\sim$ 121 nsec, a prompt hit occurs whenever there is a hit in the time window 0-44 nsec, and a delayed hit is defined as a hit falling in the window 44-121 nsec.

Track identification is accomplished in two processes, the Finder and the Linker. The Finder searches for high-P_T track segments in each of the outer-four axial superlayers of the Central Tracker. Each found segment is characterized by a mean-ϕ position in the axial superlayer, called a pixel. The Linker searches for a four-out-of-four match among segments in the 4 layers, consistent with a prompt high-P_T track. An example of a track traversing the COT is shown in Figure 2. The Finders and the Linker used in identifying this track are highlighted in Figures 3 and 4.

A block diagram of the Finder logic is shown in Figure 5. TDC information is brought from the detector to the XFT on Ansley cables. These cables can carry up to 24 signal pairs. The TDC data is multiplexed 4:1, such that prompt and delayed hit information for 4 neighboring COT cells comes in on a single cable. The Finder is designed to look for valid track segments in each of these 4 cells. A track segment is then defined by the cell it traverses for each of the 12 wire planes in a superlayer, and whether it generated a prompt or delayed hit. To keep the total number of masks relatively low, on a subset of the wires, the prompt and delayed bits are "or"ed together. A collection of the cell numbers and hit types for the 12 wires in an axial layer is called a mask. 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, hard wired on a chip). 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. Every 132 nsec the Finder outputs 12 bits per cell (12 pixels for the inner two layers, 6 pixels times 2 slopes for the outer two layers) to the Linker.

Since all of the cells in a given layer are identical, only one set of masks is needed for each layer. At present it is planned to perform the finding for 4 adjacent cells in one Finder. There will be only one set of masks stored, and the inputs from the 4 separate cells will be multiplexed one at a time. The masks required for each of the axial layers is slightly different, so there will be a separate Finder design for each axial layer.

A block diagram of the Linker logic is shown in Figure 6. Each Linker is given all of the pixel information from the Finders needed to find the tracks in a $\Delta\phi = 1.25°$ phi-slice of the tracking chamber. The Linker 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 roads are defined by their P_T and the pixel position in layer 3. The roads are then "or"ed down to 48 P_T bins and 8 pixel locations. Found tracks are passed through a priority encoder to find the

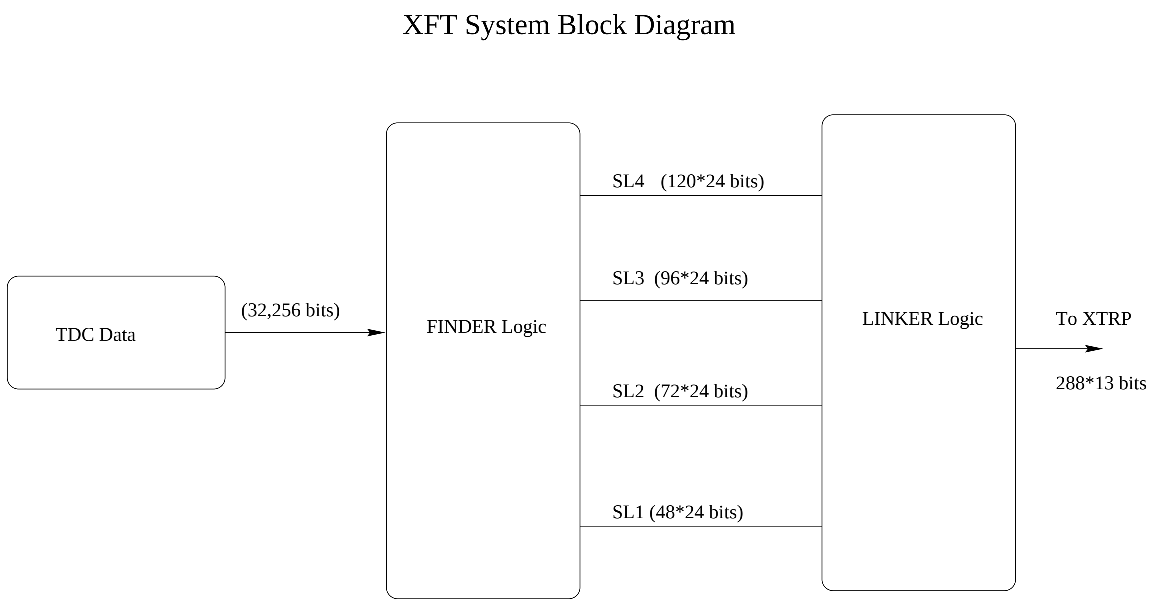

Figure 1: Block diagram of the XFT trigger.

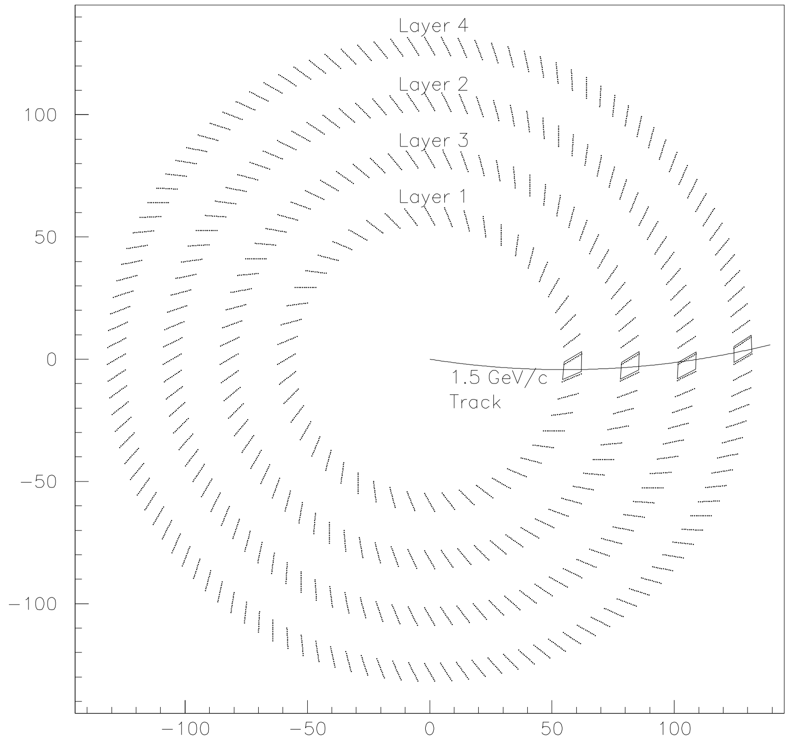

Figure 2: A schematic drawing of a track with $P_t = 1.5\ GeV/c$ in the COT. The four axial layers are shown, and only every 4th cell is shown in each layer.

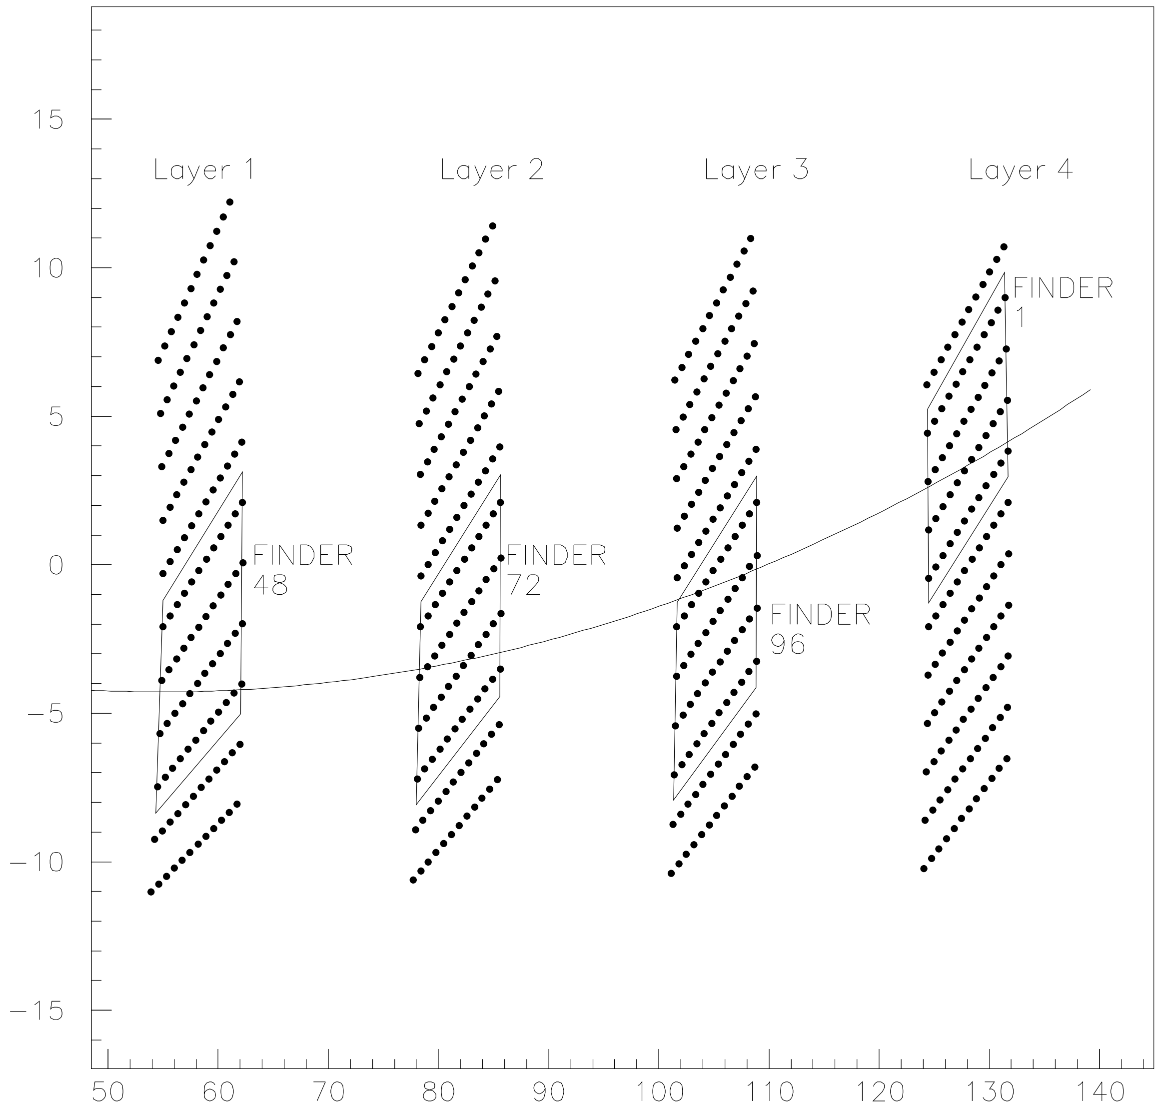

Figure 3: A close up of the previous figure of a track in the COT. All cells in all four axial layers are shown. The relevant Finder in each layer is highlighted.

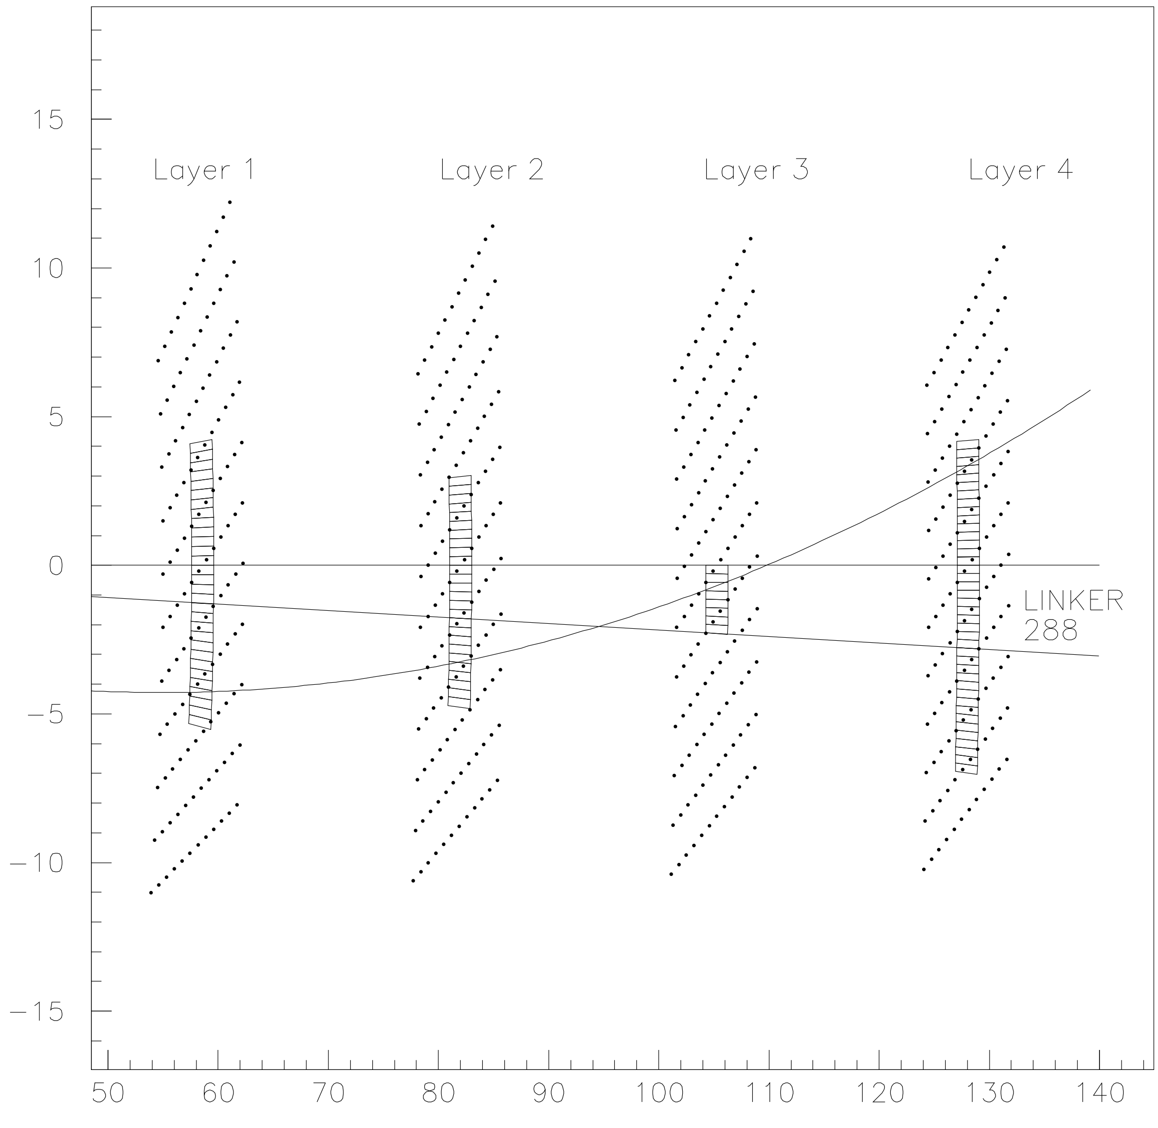

Figure 4: Same as the previous figure, but now showing the Linker the track should be found in. Also shown are all of the pixels needed by this Linker.

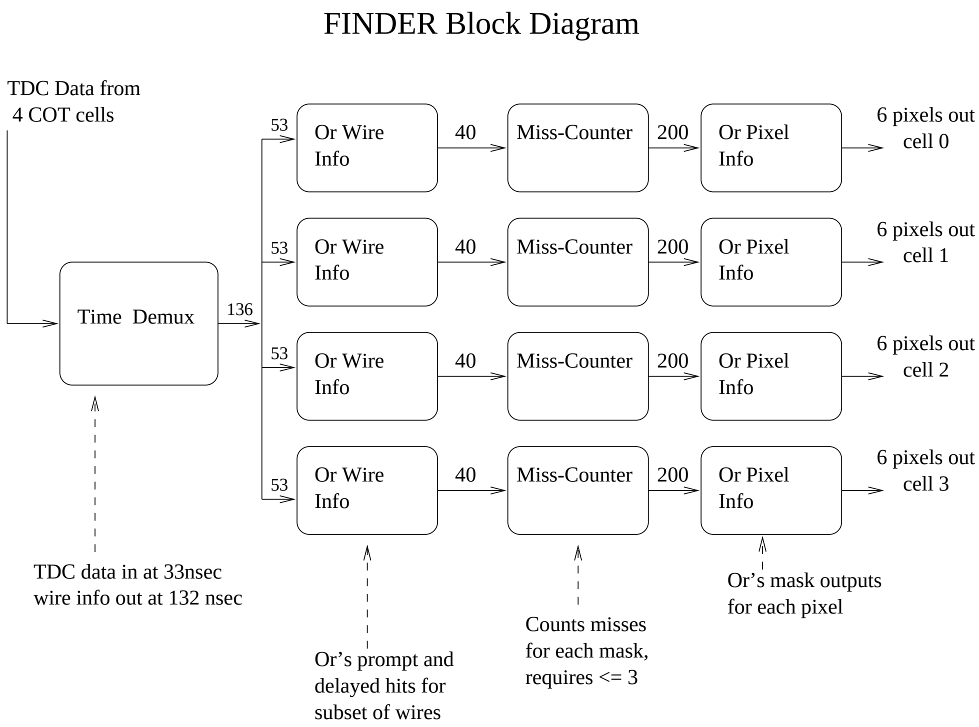

Figure 5: Logic flow for the Finder part of the XFT trigger.

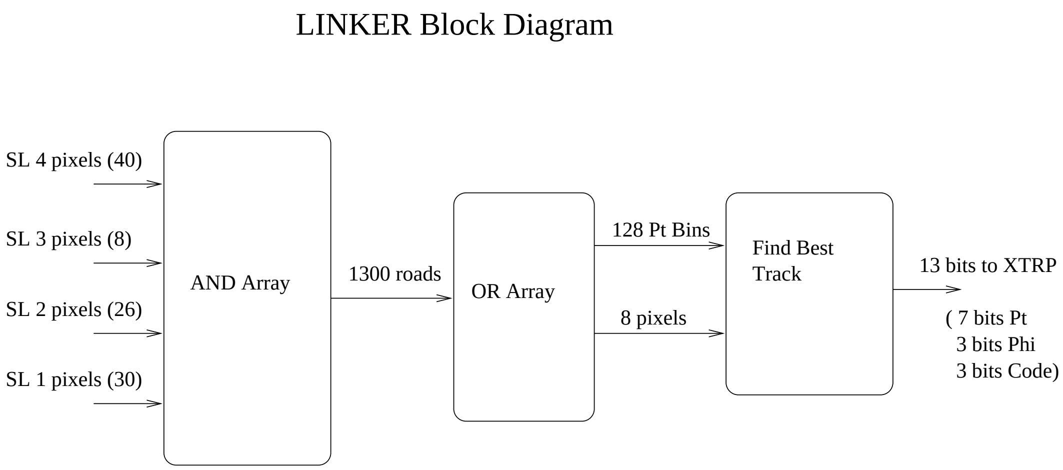

Figure 6: Logic flow for the Linker part of the XFT trigger.

highest-P_T track in the 1.25° region covered by the Linker. The linker then finds the median pixel of the pixels that have been lit, and reports the P_T and the pixel found.

3 XFT Simulation

3.1 Introduction

A software simulation of the XFT algorithm has been developed using FORTRAN. This simulation models the COT geometry, the expected pulse width ($\sim$ 40nsec) of hits in the COT, and the occupancy expected under Run II conditions (we assume 132 nsec bunch spacing). The pulse with incorporates the same pulse width modeling used by the standard COT simulation, including the aspect angle variation. The simulation was run on a data sample from Run I. This data sample was collected requiring a muon stub as defined at the trigger level. A muon stub is defined as a pattern of hits in the muon chambers consistent with the passage of a track. To simulate the performance under Run II conditions, additional hit data from minimum bias events from Monte Carlo are added, so that the effective luminosity of this sample can be dialed from 2×10^{32} $cm^{-2}sec^{-1}$ to 10×10^{32} $cm^{-2}sec^{-1}$ (at 132 nsec bunch spacing).

The number of Finder masks used for the simulation varies from 173 masks for the first superlayer to 337 for the fourth superlayer. Both the Finder masks and the Linker roads are specified using a database.

3.2 Default XFT Algorithm

We begin by presenting simulation results from the default algorithm: This algorithm has the following features:

- There are 12 pixels/cell in the inner two axial layers, and no slope information.
- There are 6 pixels/cell in the outer two axial layers, with 2 slopes/pixel: positive and negative (high momentum tracks will fire both slopes).
- There are two time bins used per wire: prompt (0-44nsec) and delayed (45-132nsec).
- There is a "veto" time bin used (44-88nsec) to prevent long pulses from firing both the prompt and delayed time bins).
- The number of P_T bins reported is 96. Bins for tracks with $P_T \leq 3$ GeV/c are twice as wide in curvature than those with $P_T > 3$ GeV/c. The momentum resolution is better at high P_T.

The features were arrived at by exmining XFT performance under a variety of configurations. These studies are described the following section.

To assess the expected performance of the XFT, we study the following benchmarks:

- Muon trigger rate
- Fake Rate for $P_T \ \geq \ 5 \ GeV/c$
- Momentum resolution
- ϕ_0 resolution
- Track finding efficiency

The results for the default XFT configuration are shown in Tables 3.2 and 3.2.

We begin by studying the track finding efficiency. First we determine which of the XFT tracks are associated with a real track with $P_T^{offline} \ \geq \ 1.5 \ GeV/c$, as determined by our standard offline algorithm. Next we ask that this track be fully within the geometrical acceptance of the COT. The efficiency has an average error of $\pm.3\%$ for 2 overlaying Minimum Bias events, and $\pm.2\%$ for 10 Min. Bias.

The momentum resolution is defined as $(P_T^{Offline} - P_T^{XFT})/(P_T^{Offline})^2$ where $P_T^{Offline}$ is the momentum using our offline tracking algorithm and P_T^{XFT} is the momentum returned by the XFT. The momentum Resolution has an average error of $\pm.03\%$ GeV/c for 2 Min. Bias, and $\pm.04\%$ GeV/c for 10 Min Bias. To calculate the ϕ_0 resolution, we use the P_T and pixel position at the third layer, and extrapolate the track to the origin. This gives an estimate of the ϕ_0 of the track and can be compared with ϕ_0^{true} of the track. The Phi Resolution has an average error of $\pm.04$ millirad for 2 Min. Bias and $\pm.06$ millirad for 10 Min. Bias.

We next examine muon trigger rates. A muon trigger at Level 1 (for Run II) is defined as a muon stub matched with an XFT track. Matching is then done by extrapolating XFT tracks to the muon chambers and looking for the presence of a muon stub within 5°. This may overestimate the true rate since in Run II we expect to be able to tighten the matching requirement to 2.5°. To calculate the muon trigger rates, we assume a muon stub cross section of 10 μb (which comes from Run I data and is independent of instantaneous luminosity). We use the simulation to determine what fraction of XFT tracks match with muon triggers. We then increase the resulting rate by a factor of 1.3 to account for the increased coverage of the muon chambers for Run II. In Figure 7 we show the muon trigger rate vs the XFT P_T threshold imposed. This is the rate expected at a luminosity of 2×10^{32} $cm^{-2}sec^{-1}$ and 132 nsec bunch spacing. For a P_T threshold of 10 GeV/c the rate is approximately 30 Hz. This is clearly well within the Level 1 budget of 50 kHz, but more importantly is also well under the Level 2 budget of 300 Hz. Recall that there is very little to reduce the muon rate at Level 2 besides increasing the P_T threshold of the XFT track.

Finally, each XFT track is classified as either fake or real, by comparing it with the list of tracks found by our standard offline algorithm. The XFT track is required to match an offline track to within 2 pixels at 3 of the 4 layers to be considered as a real track, otherwise it is classified as fake. We then report the percentage of the tracks that are fake, for both tracks that have been matched to a muon stub, and for all tracks. The error for the percentage of muon trigger tracks is about $\pm5\%$ and the error for all tracks is about $\pm3\%$.

Min Bias	EFF	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
2	99	27.4	1.2%	1.3%	0.80	2.28	96.1%
10	92	151.7	3.8%	34.3%	1.23	4.16	96.5%
spec		50	50 %	50 %	2.00	8.00	96.0 %

Table 1: XFT performance using the standalone simulation.

The performance of the XFT is dependent upon the both the wire efficiency of the COT and the luminosity. The lower the true chamber efficiency, the more misses the Finder needs to allow to attain high segment finding efficiency. Unfortunately, the XFT performance degrades as more misses are allowed. We study the performance at

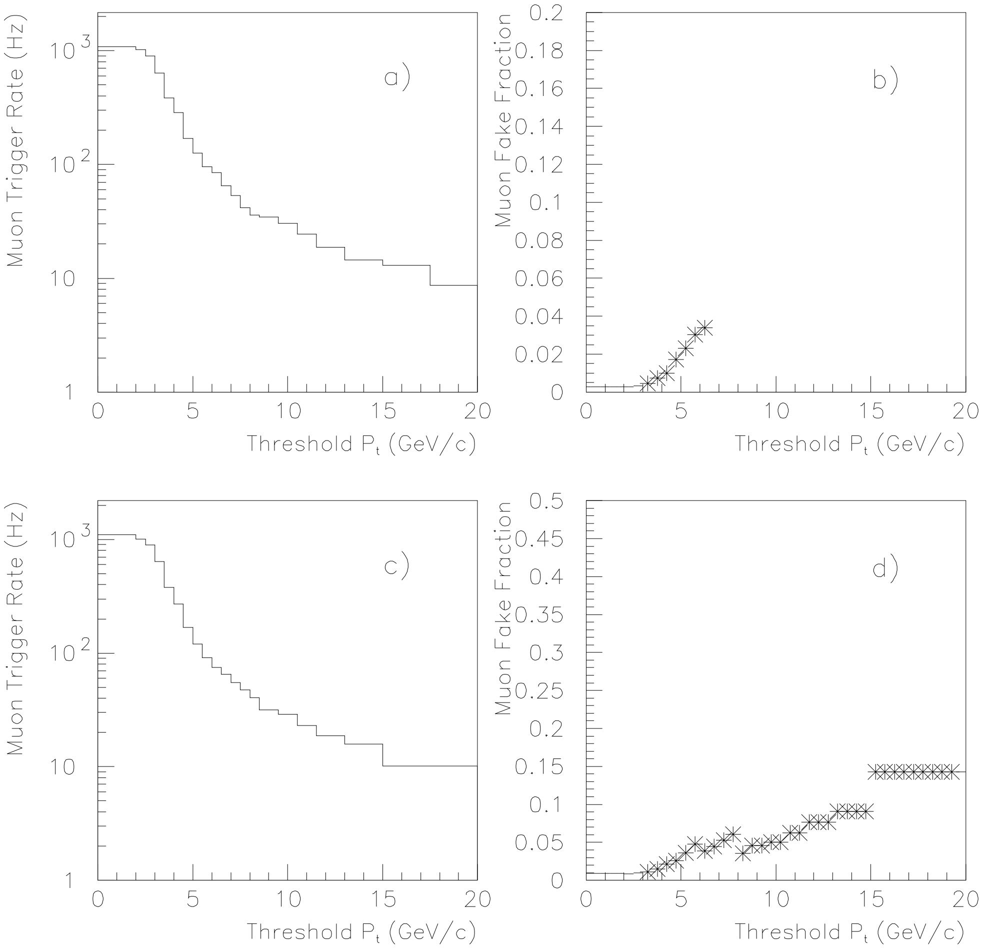

Figure 7: XFT simulation results for the muon trigger rate vs the XFT P_T threshold imposed. The simulation assumes a luminosity of $2{\times}10^{32}$ $cm^{-2}sec^{-1}$ and 132 nsec bunch spacing.

a wire efficiency of 92% and 99%. For these two cases, the number of misses allowed is 3 and 1 respectively. We also study the performance at various luminosities. As to be expected, the performance degrades when we increase the luminosity. Table 3.2 shows the results for true chamber efficiencies of 99% and 92% respectively.

Min Bias	EFF	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
2	99	27.4	1.2%	1.3%	0.80	2.28	96.1%
2	92	23.1	0.0%	3.2%	1.03	3.11	95.9%
4	99	54.9	0.0%	0.7%	0.82	2.44	96.5%
4	92	63.6	2.4%	7.7%	1.00	3.42	95.4%
6	99	82.3	0.0%	2.6%	0.86	2.46	96.1%
6	92	82.3	6.0%	11.3%	1.09	3.57	95.9%
10	99	137.2	1.3%	2.0%	0.92	2.74	96.7%
10	92	151.7	3.8%	34.3%	1.23	4.16	96.5%

Table 2: XFT performance versus efficiency and number of overlaid minimum bias events.

3.3 Optimizing the XFT Algorithm

As mentioned above, the XFT performance varies depending on exactly how we implement the algorithm. One would expect better performance if more information from the COT is used. At the same time, one would also expect an increase in complexity, and difficulty in realization, with more information. We have tried to balance these two competing effects in designing the XFT. In this section we present a number of studies which examine different XFT configurations.

In most of the following studies we use the cases of 2 Min. Bias and 99 % efficiency, and 10 Min. Bias and 92 % efficiency. Although the first study predicts expected Run 2 conditions, it is difficult to see changes in the fake rate, since the fake rate is smaller than it's error. The second study demonstrates how well the XFT will perform in harsher conditions, and changes in fake rates are more easily seen.

3.3.1 Number of Time Bins Used In The Finder

The size of the mask database can have a large effect on XFT performance. For the types of chips we are considering for the Finder protion of the XFT, we can fit about 400 masks per chip. However, there are larger chips available (at considerable increase in cost). The number of masks in the database is determined by how many time bins we have on each wire. We studied simulations that used two or three time bins on each wire. The 3 basic configurations we studied were:

- NOR2: This mask set did not "or" any wire information together, and used two time bins (prompt and delayed) for each wire. This mask set fits on the chip we intend to use.
- NOR3: This mask set did not "or" any wire information together, and used three time bins (prompt, middle, and delayed) for each wire. The outer 3 axial layers would require a larger chip to implement.
- M400: This mask set used either 2 or 3 time bins per wire. If a layer had much larger than 400 masks, some time bins were "or"ed together to reduce the mask set to about 400.

Table 3.3.1 shows the number of masks generated versus superlayer for these 3 schemes.

Mask Set	Configuration	SL1	SL2	SL3	SL4
NOR2	2-bins no "or"ing	173	242	290	337
M400	3/2 bins 400 masks	407	402	412	405
NOR3	3-bins no "or"ing	407	555	665	777

Table 3: Number of masks versus axial superlayer, for different number of time bins per wire.

As you can see the fourth superlayer has the greatest amount of masks, and also needs the most "or"ing to get a desired sized mask set.

Tables 3.3.1 and 3.3.1 show the results for the 3 configurations. In fact there is little difference among the 3. The reason for this is that the average pulse width is around 40nsec, which is comparable to the size of the time bins (which are around 44nsec for 3 time bins). Basically, the chance that a single pulse is going to fire more than one time bin is quite high.

3.3.2 Spacing of the Prompt and Delayed Bins

We also studied the spacing of the time bins, and the logic that the TDC uses to fire each bin. After we implemented the pulse-width, it was obvious, that some kind of leading edge logic would have to be employed. The two bins are set up so that the first one is from 0 to 44 nsec, and the second bin is from 44 to 132 nsec. This split up the cell into 3 equal bins. The only problem with this is that if a pulse started in the first time bin, it almost always ended in the second time bin. The current logic

Mask Set Set	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
NOR3	27.4	0.0%	0.0%	0.79	2.12	95.6%
M400	28.9	1.2%	2.1%	0.80	2.15	96.3%
NOR2	27.4	1.2%	1.3%	0.80	2.28	96.1%

Table 4: Comparison of 2 versus 3 time bins at the mask level. The wire efficiency is 99%, and the average number of minimum bias events overlaid is 2.

Mask Set	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
NOR3	137.2	2.3%	29.3%	1.08	3.55	95.8%
M400	159.0	7.7%	29.4%	1.15	3.64	95.9%
NOR2	151.7	3.8%	34.3%	1.23	4.16	96.5%

Table 5: Comparison of 2 versus 3 time bins at the mask level. The wire efficiency is 92%, and the average number of minimum bias events overlaid is 10.

employed splits up the drift time into 3 intervals, 0-44,44-88,88-132. The second time bin can be referred to as a "veto" bin, for the following reason. If a pulse is found in the first and second bin, but not the third bin, then the TDC outputs only a prompt hit. For all other cases the TDC will output hits where ever there is a pulse. The only penalty for this logic, is the loss of hits, which in turn is a loss of efficiency, because whenever there is a hit in the first and second bin, and the second pulse does not extend beyond 88 nsec, then the second hit is lost. This loss occurs only when there is more than one track hitting the same wire. At low luminosity the loss is almost unrecognizable, but at higher luminosities it can decreases the efficiency as much as 1%. Then we studied the placement of the boundary between the Prompt and Delayed hits, and at the size of the second time bin. Table 3.3.2 and 3.3.2 shows the results for these studies.

3.3.3 Number of Pixels Reported by the Finder

We studied the XFT performance versus the number of pixels reported by the Finder and used by the Linker. We varied the number of pixels/cell reported by the Finder from 6 to 12, and also investigated whether slope information reported by the Finder could be useful. The slope information is a simple +/- tag, indicating whether the track segment identified by the Finder was positive or negative P_T. High P_T segments would fire both slope tags. The results are shown in tables 3.3.3 and 3.3.3. The number of roads needed by the Linker for these configurations is shown in Table 3.3.3.

2nd prompt bin	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
30-85	26.0	2.5%	0.7%	0.80	2.20	97.0%
33-88	27.4	2.5%	1.4%	0.78	2.19	96.6%
35-90	27.4	1.2%	0.0%	0.76	2.13	96.3%
40-95	28.9	2.4%	1.3%	0.78	2.13	96.5%
44-88	27.4	1.2%	1.3%	0.80	2.28	96.1%
45-100	28.9	1.1%	0.0%	0.85	2.24	96.7%

Table 6: XFT performance versus spacing of the "veto" time bin, The wire efficiency is 99%, and the average number of minimum bias events overlaid is 2.

2nd prompt bin	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
30-85	130.0	3.6%	7.2%	0.89	2.59	96.5%
33-88	158.9	3.7%	5.6%	0.89	2.67	96.6%
35-90	130.0	3.8%	3.3%	0.88	2.46	96.7%
40-95	115.6	2.4%	3.9%	0.93	2.57	96.3%
44-88*	137.2	1.3%	2.0%	0.92	2.74	96.7%
45-100	130.0	2.4%	2.7%	0.99	2.59	96.4%

Table 7: XFT performance versus spacing of the "veto" time bin, The wire efficiency is 92%, and the average number of minimum bias events overlaid is 10.

The best performance with a reasonable number of Linker roads is given by the 12-6S configuration, which has 12 pixels reported per Finder Cell in the inner two axial layers, and 6 pixels with slope information reported by the outer two axial layers.

3.3.4 Number of Pt Bins Reported by the Linker

Finally, we also studied the XFT performace versus the the number of "reported" P_T bins. In this studied, the roads used by the Linker algorithm were all the same. The difference is in how these roads are "or"ed down to the reported P_T. The more the roads are "or"ed down, the smaller the number of reported P_T bins, and the worse the reported P_T resolution. The gain is in reducing the complexity of the Linker chip. The results are shown in Table 3.3.4. Not that for 128 and 64 P_T bins, the performance is similar. For the 48 bin case, we kept the Pt bin width above 3 GeV the same as the 64 and 128 bin cases (which is why the high P_T resolution is similar in these 3 cases). The design choice we have made is 48 P_T bins.

No. Pix	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
6	26.0	0.0%	1.4%	1.04	2.93	96.4%
12-6	28.9	0.0%	2.0%	0.86	2.32	96.2%
12-6S	27.4	1.2%	1.3%	0.80	2.28	96.1%
12	27.4	0.0%	1.4%	0.72	2.04	96.6%

Table 8: XFT performance versus the number of pixels reported by the Finder and used by the Linker. 6= 6 pixels/Finder Cell for all layers, 12 = 12 pixels/Finder Cell for all layers, 12-6 means 12 pixels in the inner 2 axial layers and 6 pixels in the outer 2 axial layers, and 12-6S means 12 pixels in the inner 2 axial layers and 6 pixels with two slopes per pixel in the outer 2 axial layers. The wire efficiency is 99%, and the average number of minimum bias events overlaid is 2.

No. Pix	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
6	209.4	14.3%	51.2%	1.47	4.98	95.7%
12-6	173.3	10.1%	44.2%	1.25	4.21	95.5%
12-6S	151.7	3.8%	34.3%	1.23	4.16	96.5%
12	195.0	10.8%	42.6%	1.25	4.02	95.6%

Table 9: XFT performance versus the number of pixels reported by the Finder and used by the Linker. 6= 6 pixels/Finder Cell for all layers, 12 = 12 pixels/Finder Cell for all layers, 12-6 means 12 pixels in the inner 2 axial layers and 6 pixels in the outer 2 axial layers, and 12-6S means 12 pixels in the inner 2 axial layers and 6 pixels with two slopes per pixel in the outer 2 axial layers. The wire efficiency is 92%, and the average number of minimum bias events overlaid is 10.

No. Pixels	Number of Roads
6	1290
12-6	2419
6-12	2937
12	4853

Table 10: Number of Linker roads versus number of pixels reported by the Finder.

No. road bins	Mu Rate	% fakes(Mu) $P_T \geq 5$	%fakes(all) $P_T \geq 5$	P_T Res (%GeV/c)	Phi Res (mRad)	Eff
128	26.0	2.6%	1.4%	0.75	2.10	96.6%
64	27.4	2.5%	1.4%	0.78	2.19	96.6%
48	27.4	1.3%	1.4%	0.89	2.48	96.5%
32	28.9	2.2%	1.8%	0.99	2.58	96.6%
16	17.3	3.3%	1.2%	1.41	3.68	96.7%

Table 11: XFT performance versus number of Pt bins used for the out Pt. The wire efficiency is 99%, and the average number of minimum bias events overlaid is 2.