CDF/DOC/TRIGGER/CDFR/DRAFT/4300

XFT3D: A Proposal for Adding Stereo Segment Finding and 3D XFT tracking in the Trigger

John Hoftiezer, Richard Hughes, Mark Johnson, Phillip Koehn, Brian Winer
Ohio State University
Jim Freeman, Scott Holm, Jonathan Lewis, Terri Shaw, Intea Yu
Fermilab
David Saltzberg
Univerity of California at Los Angeles
September 22, 1997
Version 1.0

1 Introduction

The baseline trigger design for Run II incorporates tracking in the Level 1 trigger decision. The tracks are found by the eXtremely Fast Tracker (XFT) and the resulting track list is sent to the extrapolation box (XTRP) for distribution and matching with other Level 1 primitives. The XFT track list is also sent to the SVT for identification of displaced tracks at Level 2. The XFT identifies tracks in $r\phi$ only, using the 4 axial layers of the COT. However, since both the electron and muon subsystems are segmented in z, the ability to extrapolate and match XFT tracks in the z view could prove extremely useful in reducing fakes. In addition, a single z measurement allows a coarse estimate of $cot(\theta)$, which in turn allows a multi-track mass calculation at the trigger level.

In this document we describe a method for adding a z measurement to 2D XFT tracks. The z measurement is accomplished by finding segments in the outermost stereo layer of the COT, and then associating found segments with previously found 2D XFT tracks. As we will describe below, the stereo segment finding can be accomplished using a straightforward extension of the current XFT design. Matching these segments with 2D XFT tracks will be done at Level 2 using the Alpha processors. We call the resulting track list XFT3D tracks.

In the following sections, we describe the method for identifying stereo segments in the COT. We detail the extensions that need to be made to the existing XFT design to accommodate stereo segment finding and the changes needed to transport the information to Level 2 processors. We present preliminary performance studies of the stereo segment finding, and the resulting impact on trigger rates. Finally, we present preliminary estimates of cost, scheduling, and manpower.

2 XFT Stereo Segment Finding

2.1 Baseline XFT Design

We will depend heavily on the existing design of the XFT for the segment finding in the stereo layers. Documentation on the XFT can be found on the web at

http://www-cdf.fnal.gov/upgrades/daq_trig/xft/xft_home.html.

We give a short summary of those parts of the XFT that are relevant for this discussion.

The XFT processor works off 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 ~ 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. The naming convention and some geometrical constants for the COT layers are shown in Table 1.

Track identification is accomplished in two processes: Finder and 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.

The input data for the XFT is TDC information carried on Ansley cables from the collision hall to the 1st floor. These cables can carry up to 24 signal pairs. The TDC data is multiplexed 6:1, such that prompt and delayed hit information for 4 neighboring COT cells comes in on a single cable, every 132 nsec. 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. A "mask" is defined by a collection of cell numbers and hit types for the 12 wires that are consistent with a good track segment. 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 number (up to 3) of missed wire planes. Every 132 nsec the Finder outputs 12 bits per cell (12 pixels for the inner two axial layers, 6 pixels times 2 slopes for the outer two axial layers) to the Linker.

The Finder chips are implemented using Altera's FLEX 10K-series programmable logic devices. The number of Finder chips needed for each of the axial superlayers is shown in Table 1. The Finder subsystem of the XFT consists of 48 9U VME cards, of two types. These 48 cards are housed in 3 VME crates. Finders from layers 2 and 6 are grouped together (2-6 boards) and finders from layers 4 and 8 are grouped together (4-8 boards). There are 6 finders on 2-6 boards, 2 from layer 2 and 4 from layer 6. There are 8 finders on 4-8 boards, 3 from layer 4 and 5 from layer 8. The hit data for each of the

Finder boards comes primarily from Ansley cables which plug into a transition card in the back of a crate behind the slot for a given board. However, each card also requires some information from the Ansley cables which feed the neighboring slot. Since there are 3 crates for the Finder cards, the neighbor information needed across crates is obtained from cables with duplicate hit data sent to the different Finder crates. Note that the mechanical properties of the Ansleys (they are bulky and relatively inflexible) and the limited space in the back of a VME crate means that 8 Finders per board is an effective maximum.

Each Finder chip has a Level 1 pipeline and Level 2 buffers for both the input wire data as well as the output pixel information. The Finder boards send output pixel information every crossing to the Linker subsystem over cables through the front panel. The driver/receiver pair is the National Channel Link set. Each channel Link drives data from 2 Finder chips, and each Finder chip needs to be sent to 2 Linker boards.

The Linker is the second stage of the XFT. It has the responsibility of linking pixels *between* the axial superlayers of the COT. The pixels themselves are found by the Finder. The Linker is separated into 288 ϕ slices of $1.25°$. Each slice has a dedicated chip with the task of finding tracks within the slice. The results of the Linker are passed to the XTRP for extrapolation to other portions of the detector and used in the Level 1 trigger decision. The information sent to the XTRP includes 7 bits of track P_T and 3 bits of local ϕ, measured at the 3rd axial layer (SL6) of the COT.

Layer	Radius	Number of Cells	Finders	Finder masks
1	46	168	–	–
2	58	192	48	170
3	70	240	–	–
4	82	288	72	230
5	94	336	–	–
6	106	384	96	272
7	119	432	108	300
8	131	480	120	310

Table 1: Some COT and XFT information vs COT superlayer.

2.2 Some Stereo Math

The measurement of the z coordinate of a track in a stereo layer of the COT requires 2 things: an identified track in the $r - \phi$ plane and a measurement of the ϕ position of the track segment in the stereo layer. If the track passes through the stereo layer at a z coordinate other than z=0, there will be an apparent difference in ϕ between the extrapolated position of the $r-\phi$ track at the stereo layer (denoted ϕ_{exp}), and the measured ϕ position of the track segment in the stereo layer(denoted ϕ_{fnd}). The z coordinate is given by the following expression:

$$z = \tan\left(\phi_{exp} - \phi_{fnd}\right) \cdot \frac{1}{S} \tag{1}$$

where S is a constant dependent on the "stereo angle". For COT stereo layer 7, $S = 2.75 \times 10^{-4}$.

2.3 Stereo Segment Finding Algorithm

The stereo segment finding works almost exactly like the segment finding for the axial layers. In defining the masks, we assume that the stereo angle is zero. The total number of masks needed to identify track segments with P_T greater than 1.5 GeV/c is then 300. As for the outer 2 axial layers (SL6 and SL8), we divide each cell into 6 pixels. Each pixel is characterized by 2 bits. The definition of the bits is:

- 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)

We then need to associate stereo segments to found XFT2D tracks. Each XFT2D track passing through a pixel in SL6 can be matched to a range of pixels in SL7. This range depends on the P_T (or curvature) of the track and the stereo angle. Assuming that the track comes from the beam line, and assuming zero stereo angle, the extrapolated position is SL7 is given by:

$$\phi_{SL7} = \phi_{SL6} + \sin^{-1} K \cdot \frac{R_7 - R_6}{P_T}$$

where ϕ_{SLn} and is the ϕ position in radians corresponding to the XFT pixel position in superlayer n. R_n is the radius in centimeters of superlayer n, and K is a constant $(= 0.5 \cdot c \cdot B = 0.00014989 \cdot 14.116)$. Given that we can go down to a P_T of 1.5 GeV/c, the range of pixels in SL7 we need to search over just due to curvature is ± 8 pixels. We then need to determine the number of additional pixels we need to search over due to the stereo angle. This angle is achieved by stringing the COT wires to cells which are offset by 3 cells at one endplate and by 3 cells in the opposite direction at the other endplate. Since each cell is split into 6 pixels by the Finder, one needs to look ± 18 pixels on either side of the *expected* pixel position (which assumes $z = 0$) at SL7.

2.4 Stereo Segment Finding Implementation

We define Finders which look for valid track segments in each of 4 neighboring cells. Given that the number of cells in SL7 is 432, and that each Finder covers 4 cells, the total number of Finders needed is 108. As mentioned above, due to the size and inflexibility of the Ansley cables, we are limited to a practical maximum of 8 Finders per board. We plan to use 6 Finders per board. This makes each board identical, and also allows room for the circuitry needed to format the data before sending it to the L2 interface board. Using 6 Finders per board, we need a total of 18 stereo Finder boards. Since each VME crate has 21 slots, of which 2 are taken up by the controller and tracer, we have 19 slots remaining.

The stereo Finder boards can then fit into a single crate. However, the Finder boards on either end of the crate will need neighbor wires from Ansleys that feed the opposite end of the crate. Rather than drive these along the full length of the backplane, we can simply use a separate cable with duplicate hit information. The two slots on either end of a stereo Finder crate will then receive 7 cables instead of 6.

Next we need to consider which parts of the XFT Finder design need to be modified to implement stereo segment finding.

- TDC mezzanine card, TDC transition module: These modules provide the prompt and delayed hit information, and drive this data onto the Ansley cables. There are no changes necessary.

- Ansley cables: A total of 110 ansley cables would be required (not including spares). The current allocation of Ansley cables has enough spares to accommodate this requirement.

- Finder transition module: The transition module is the board that plugs into the back of the VME crate behind each Finder board. It has the connectors and LVDS receivers for the Ansley cables. The current design has connectors for up to 8 Ansley cables and should be fine for the stereo Finder design.

- Finder Backplane: The Finder backplane is responsible for routing neighbor wires from slot to slot, so this will be new.

- Finder chip: This will be the same. Remember that we implement the axial Finders using the Altera FLEX 10K series cPLDs. These are fully reprogrammable devices. The Finder chip design incorporates logic for latching data on input to the chip, Level 1 pipelining and Level 2 buffers, dead wire removal and of course the algorithm for finding masks. It is this last section that is different for each axial layer, and it is only this section of the chip that will be different for the stereo layer as well. However, this section is actually implemented using Altera's software logic package (called TDF's) rather than schematic capture, and this software logic is in turn generated automatically from the Finder databases.

There is no additional work needed to implement a Finder chip for the stereo layers.

- Finder Board: This will require the largest number of changes. The input section from the Ansley cables is identical and needs no changes. The Finder chip section will also remain the same. The output section of the axial Finder boards use National channel links to drive data to the Linker boards. The Linker boards are not designed to take input from stereo layers (and it would not help anyway). This section should be replaced with a design to drive the data to a Level 2 board (more on this later).

Some parameters of the stereo Finder system are listed in Table 2.

Parameter	Value
Finder Chips	108
Finder Cards	18
Finder Transition Modules	18
TDC Transition Modules	54
TDC Mezzanine Cards	54
VME Crates	1
Ansley Cables	110

Table 2: Some parameters of the stereo Finder system.

2.5 3D tracking performance

We have implemented a version of the stereo segment finding algorithm in the XFT standalone simulation. For both the axial layers and the stereo layer, the Finder is implemented using a database of masks. The association of SL7 segments to XFT2D tracks is done by considering all segments in a valid window around the XFT2D track. As will be shown below, the fake background is small enough that this is a relatively minor detail (meaning it does not matter if we took the 1st segment, the average of all segments, or the closest segment).

We have generated 2 stereo segment databases, one for tracks with $P_T > 1.5$ GeV/c, and one for tracks with $P_T > 8.0$ GeV/c. The efficiency for finding a stereo segment vs the P_T of the track for these two cases is shown in Figure 1. The difference in turn on in the two cases is due to the fact that momentum resolution degrades as P_T increases since the resolution is proportional to P_T^2. In both cases the plateau efficiency is $\sim 97\%$. For the following studies, we use the stereo database with $P_T > 1.5$ GeV/c, and no cut on the P_T of the segment is made.

The difference between the true z position and the XFT determined z position is shown in Figure 2 for a sample of simulated tracks. In this simulation, the track P_T was flat from 1.0 to 10.0 GeV/c. Tracks from an average of 2 minimum bias events were overlaid on these Monte Carlo events to simulate the effect of occupancy in the COT in Run II. A similar plot is shown in Figure 3, where in this case the average number of minimum bias events is 10. The z coordinate is given by equation 1. The uncertainty on the z position is determined by how well we measure the two phi positions. In this case, the uncertainty is dominated by the extrapolation uncertainty on ϕ_{exp}, which comes from the XFT2D track measurement errors on P_t and ϕ_{SL6}. The z position resolution is $\sim$ 6cm for the overlap of both the 2 and 10 minimum bias events.

We have also investigated how often there is an additional XFT stereo segment within a valid region around an XFT track. In Figure 4 we show the number of valid stereo segments found per XFT track, for events with an average of 2 minimum bias events overlaid. In Figure 5, we show the same plot for events with an average of 10 minimum bias events overlaid. This distribution peaks at 2 because 1 miss is allowed in finding masks (a COT wire efficiency on 99.0% is assumed), and the 2 segments are in fact at neighboring pixels. The true background can be gauged by looking at how often 3 or more stereo segments are found, which is about 4% of the time with an average of 2 minimum bias events overlaid, while it is 12% with an average of 10 minimum bias events overlaid. (This is actually an overestimate, since some fraction of the cases with 3 stereo segments have all 3 segments as neighbors.) In Figures 4 and 5, we also show a plot of the fraction of XFT tracks with at least one good (or bad) stereo segment. This is plotted as a function of the number of found stereo segments. A good segment is defined as one within 20cm of the true z position of the track, while a bad segment is one outside of this range. As expected, when there are few stereo segments, most of them are good. And even when there are many segments, at least one of them is good.

A summary of the stereo segment finding performance can be found in Table 3.

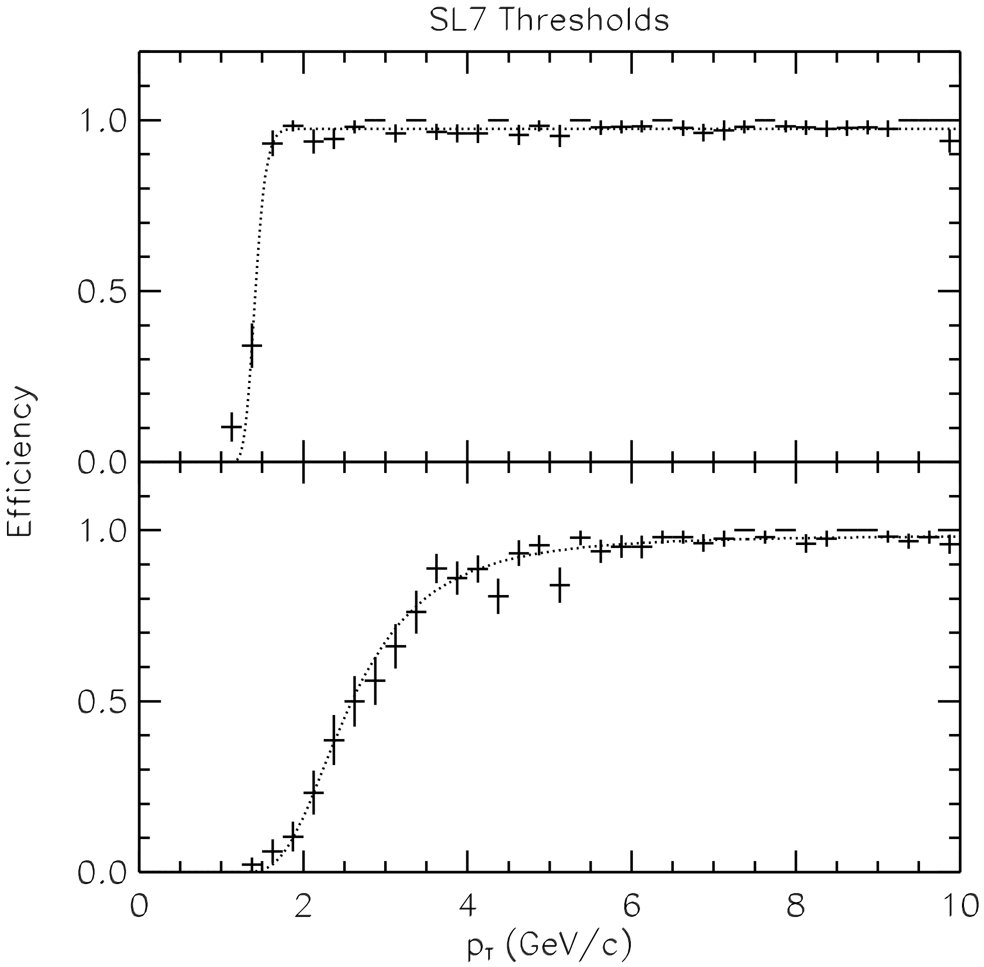

Figure 1: Efficiency of SL7 segment finding for low- and high-P_T segments. Dotted line is a fit to the normal frequency function in $1/P_T$. For the high-P_T masks the 50%-efficient point is at $2.68\,\mathrm{GeV}/c$ with a resolution of $0.11\,(\mathrm{GeV}/c)^{-1}$.

Performance	MB2	MB10
Efficiency	99.5%	98.9%
z Resolution	5.8cm	5.9cm
Comb. Backg	3.5%	12.7%

Table 3: Summary of the stereo segment finding performance using the XFT simulation. Note that the z resolution includes the extrapolation uncertainty of the XFT 2D track into the stereo layer.

3 Level 2 Segment-Track Matching

The basic strategy for the Level-Two (L2) trigger analysis of the stereo tracking data is to send a map of found SL7 segments directly to the L2 decision crate, which handles linking and further processing. There will be an azimuthal offset of the SL7 segments in proportion to the z position of the track as given in Equation 1. For each axial track, the L2 software will determine at which 6 cells of the 432 total cells in SL7 an associated track segment may exist. The total segmentation of these 6 cells will be 18 bins. The segmentation is a factor of two coarser than that used for $r - \phi$ tracking in order to reduce the required bandwidth into the L2 decision crate. Higher resolution would not add much, however, due to primary z-vertex smearing. The position of the SL7 segment in these 18 pixels yields its z position.

The L2 processor will calculate the z position for all possible SL7 segments associated with an axial track using a simple algorithm. As described in section 2.4, there will be 18 Finder boards, each with 6 Finder chips, corresponding to a total of 144 ϕ bins. For each ϕ bin the Finder boards will return 2 bits corresponding to the track p_T as described in section 2.3. The ϕ of the track at SL7 will be determined from the 2D track parameters (the p_t and SL6 ϕ position) by the L2 processor. This will be done using a lookup table. As noted earlier, depending on the z position of the track (as well as the small effect of the curvature), there are a range of possible z positions (and thus, ϕ bins) which are valid. The lookup table returns the ϕ bin corresponding to the minimum z position. The next 18 ϕ bins thus correspond to the range of possible z positions of the track at SL7. Each bin will hold the 2 p_T bits. Only high-p_T found segments will be associated with high-p_T tracks. The association is similar for low-pt tracks with the added condition that the slope of the segment be consistent with the

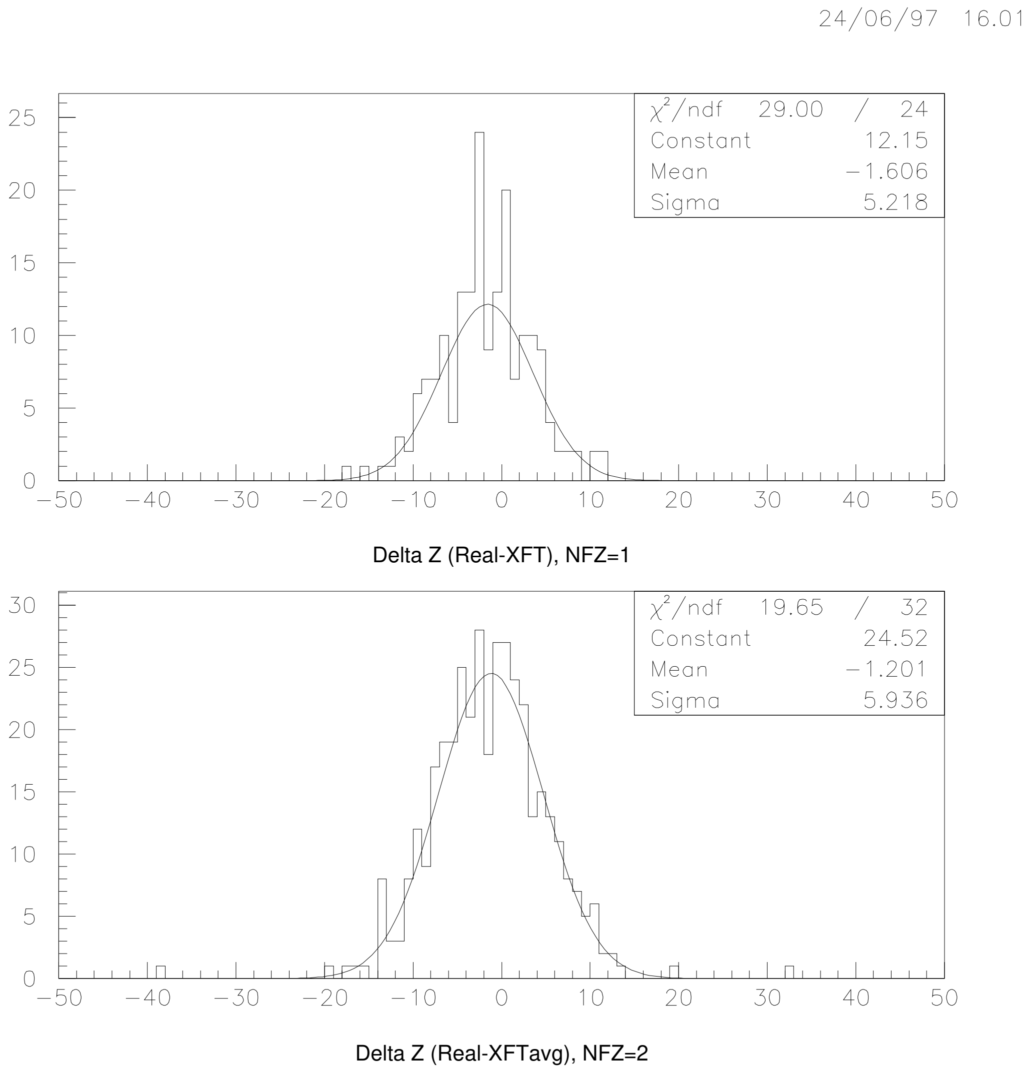

Figure 2: Top plot: Difference between the true z position and the XFT z position at outer stereo layer for one and only one valid stereo segment near XFT track. Bottom plot: Difference between the true z position and XFT z position at outer stereo layer for two valid stereo segments near XFT track. For both plots the average number of overlaid min-bias events is 2.

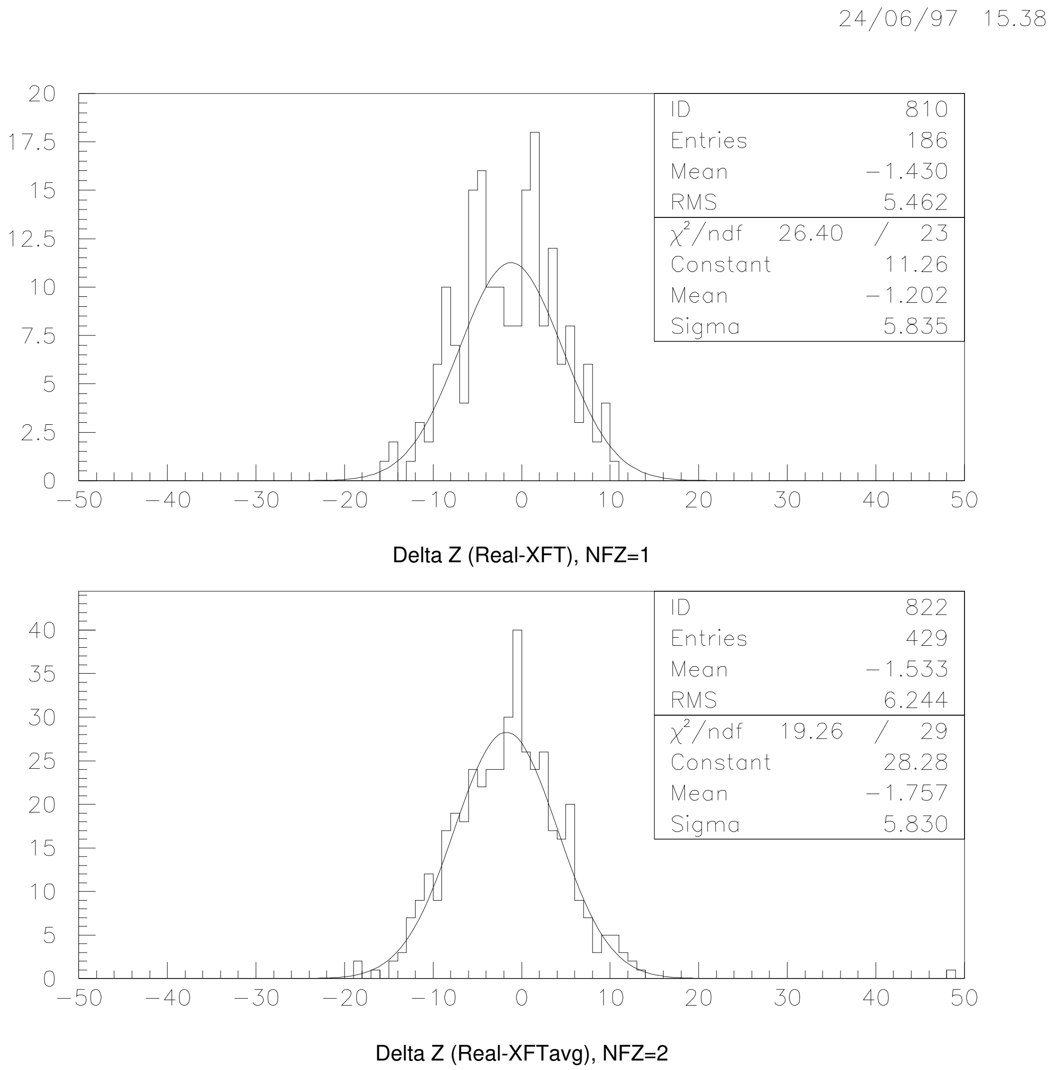

Figure 3: Top plot: Difference between the true z position and the XFT z position at outer stereo layer for one and only one valid stereo segment near XFT track. Bottom plot: Difference between the true z position and the XFT z position at outer stereo layer for two valid stereo segments near XFT track. For both plots the average number of overlaid min-bias events is 10.

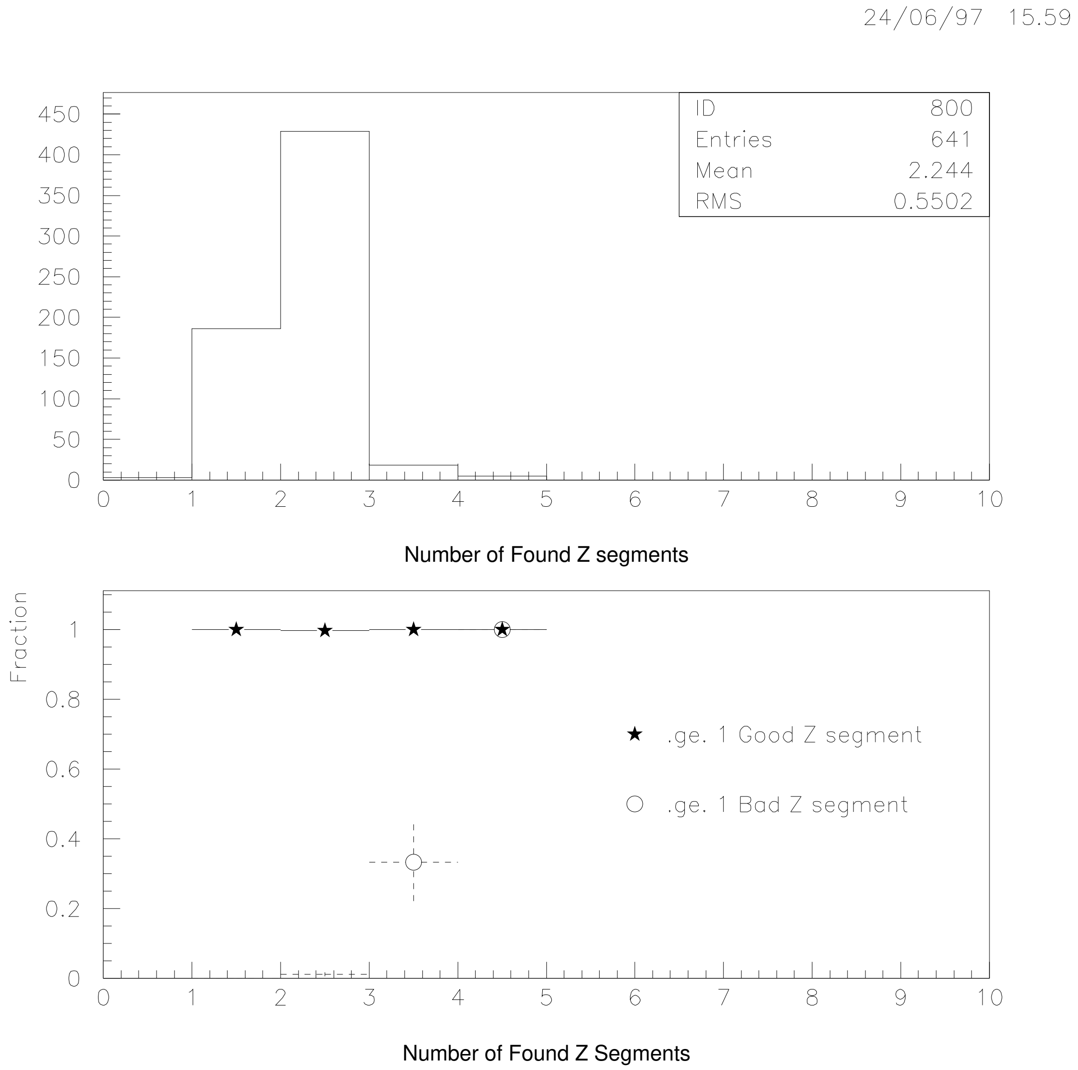

Figure 4: Top plot: Number of valid stereo segments per found XFT track. Bottom plot: Fraction of above tracks with at least one good (or bad) stereo segment. For both plots the average number of overlaid min-bias events is 2.

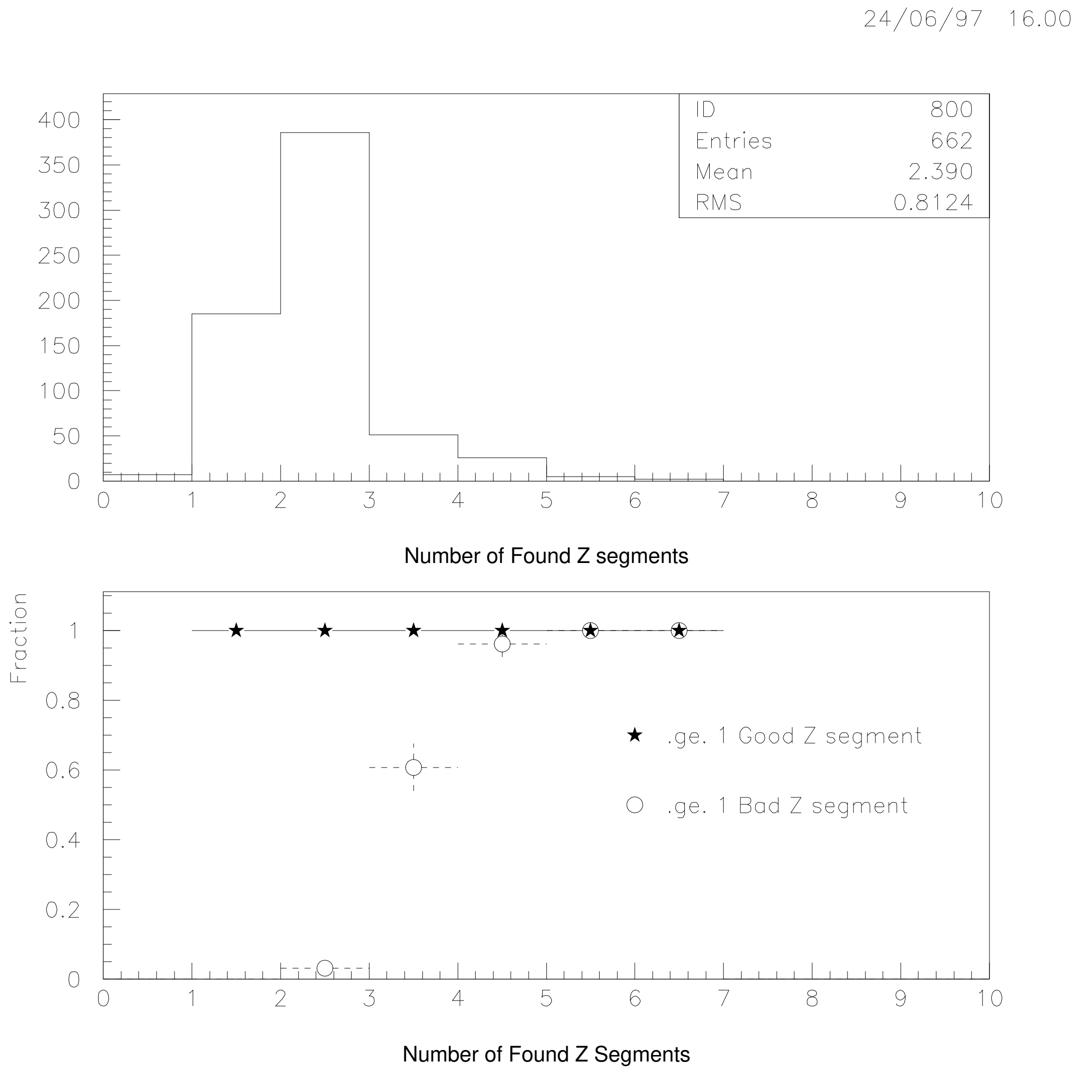

Figure 5: Top plot: Number of valid stereo segments per found XFT track. Bottom plot: Fraction of above tracks with at least one good (or bad) stereo segment. For both plots the average number of overlaid min-bias events is 10.

charge of the track. If more than one segment consistent with a 2D track is found, all will be considered for z matching and mass window cuts.

In order not to incur extra deadtime, the data must be transmitted before other systems have completed uploading into L2, or approximately 1.5 μs. Hence we need to accommodate a data transmission rate of 192 Mb/s per board. The data will be loaded through the front panel and stored in memories on the receiver board. The loading of the receiver board is accomplished in parallel with other processes hence should add no deadtime to the trigger. High-speed copper links such as LVDS Channel Link by National Electronics are limited to approximately 10 meter transmission lines [1]. Since the Finder boards and the L2 decision crate are on different floors, and further than the 10 meter cable limit imposed by the National LVDS channel links, we plan to use optical links. We have decided to send the data serially over low-speed optical links, such as the HotLink by Cypress which can run at 330Mb/s. The extra cost of optical relative to copper is about $100 per transmitter/receiver pair. Each Finder board will transmit its data over a single optical fiber to L2, so that the SL7 receiver board will receive 18 optical fiber links. There may not be enough room on the front panel of a 9U VME board to accept these individually but special multi-fiber connectors can be used.

Once loaded through the front panel, the data must be transmitted to the L2 processor boards via the L2 backplane, the 128-bit "Magic Bus". [2] Here speed is critical to reducing deadtime since only one board may use the backplane at a time. There are two options, for which both functionalities should exist. In one option, the data would be sent to the L2 processors as a block of 2592 bits over the Magic Bus, which would take 21 cycles or ~800 ns. Alternatively, the board may be polled via the Magic Bus to return in a single word the relevant 36 bits corresponding to a given track. Polling would be faster for low-multiplicity events. To separately address each track will require that each bin's 2 bits be independently and sequentially addressable via the Magic Bus. It will also require that the first 18 bins in ϕ be appended to the end of the list, for ϕ continuity. Which transfer protocol is used can be programmed into the L2 software and could even depend on the size of the event.

The SL7 segment receiver board will actually be a general-purpose L2 interface and so will also be available to receive data from any new, as-yet undetermined, system which may need to send data to L2. Since it will use optical receivers there will be little constraint on the location of such a system. We should make three copies of the final board. One used in the system, one hot spare and another in reserve for possible new application.

4 Impact on Trigger Rates

The addition of stereo information can reduce trigger rates for many trigger signatures. We consider two examples: CMX muons and $J/\psi \rightarrow \mu^+\mu^-$. For CMX muons, we require the association of a SL7 segment with an XFT track that is consistent with pointing to the CMX. Further, we can require the SL7 segment to pass the high-P_T threshold described in section 2.3. The same approach can be used for η-gap muon triggers and for electron triggers, where an SL7 segment must be consistent with a track projecting to the seed tower of an electromagnetic cluster. Similar pointing cuts might also be used for tau triggers. In all cases, the 30 cm r.m.s. spread of the z position of the primary interaction vertex sets a limit on the required precision on the z measurement. In principle, trigger rates can be further improved if Tevatron upgrades reduce the length of the luminous region. For $J/\psi \rightarrow \mu^+\mu^-$ triggers, we consider CMU-CMU pairs and use all the SL7 segments that can be associated with each XFT track to form mass hypotheses. Candidates consistent with the J/ψ mass would be accepted. Since the tracks are typically low-P_T, SL7 segment slope pointing may also improve the rejection. Good mass resolution can be achieved because the uncertainty in the primary-vertex z position does not translate into a large uncertainty in the opening angle between the two muons.

To simulate Level 2 stereo reconstruction, we use the RUN2TRS package[3]. XFT tracks are simulated using the hit-level XFT-CTC simulation. This probably gives poorer resolution and a higher fake rate than that of the XFT-II in the COT, but for trigger-rate estimates, it provides a worst-case analysis. Run 2 tight muon matching is simulated in the RUN2TRS package. To simulate SL7 XFT segments, we use offline CTC tracks. Because segment fake rates in the outer layers are low [see Section 2.5], [4] we argue that the additional background from fake SL7 segments will be small. For each track with $P_T > 1.1\,\mathrm{GeV}/c$, we use Eqn. 1 to translate the position of the track in SL7 to the effective position at $z = 0$. This position can be translated into a pixel number where there are

$$432 \text{ cells} \times 3 \text{ pixels/cell} = 1296 \text{ SL7 pixels}^*$$

We store two bit-arrays, one for positively-charged tracks, and one for negative. For each track, we set the bit of the corresponding position and charge. A parameterization of the turn-on curve for high-P_T bits (Fig. 1) was used to determine if both charge bits would be set, indicating a high-P_T segment.

We test the algorithm for z position reconstruction and determine the mass resolution in a sample of 10000 CMU-CMU $J/\psi \rightarrow \mu^+\mu^-$ events from the PSIA stream. To avoid threshold effects, we require the offline-measured P_T to be greater than 2.1 GeV/c

*For the trigger studies we used 3 phi pixels/cell, instead of 6 as described in section 2.3. This was to reduce the amount of information transferred to the L2 processor. Studies are in progress to determine if the full resolution of the stereo segment phi position is beneficial and the transfer implementable.

for each muon and dimuon mass to be between 3.0 and 3.2 GeV/c^2. We also make the position-matching cuts $\chi_x^2 < 9$ and $\chi_z^2 < 12$. Rather than using a hit-level XFT simulation, for this study, we simulate XFT tracks by smearing offline CTC tracks by $\delta(1/P_T) = 0.015\,(\mathrm{GeV}/c)^{-1}$ and $\delta\phi_6 = 1.5\,\mathrm{mrad}$. We do not attempt distinguish between real J/ψ events and background. Figure 6 shows the resultant reconstructed mass peak. If there is more than one candidate in an event, we select the one closest to the true J/ψ mass. The distribution is fit to a Gaussian with width 0.16 GeV/c^2. Since the sample includes some background, this is a slight overestimate of the mass resolution. In events with both muons passing the level 2 matching cuts, the efficiency for finding a mass within 0.5 GeV/c^2 of the J/ψ mass is 93%.

In Figure 6 we also show similar mass peaks for three alternative mass-reconstruction triggers that one may consider.

M_T: Although a transverse mass cut would be possible without new hardware, it has significantly lower efficiency. Also, the rate for a transverse mass cut will grow quickly with the muon threshold reduced to 1.5 GeV/c, half the J/ψ mass, and would likely grow with luminosity. Obviously, this option for dimuon triggers could have no impact on inclusive leptons.

CHA: In principle it would be possible to use the hit towers from the hadron calorimeter TDC (HTDC) trigger signals to find a set of valid z positions for muons to form mass combinations. This would also require a new readout path to Level 2 so it would not be cost-free. Again, since the hadron calorimeter is already used in forming CMX and η-gap muon triggers and is a veto for electrons, HTDC information would not help inclusive leptons.

TOF: The RUN2TRS simulation simulates the z information that could be obtained from a time-of-flight system of 216 bars with 1 ns time-difference segmentation. The overall segmentation, resolution and performance would be similar to the proposed SL7 system. However, first CDF would have to buy a TOF system.

Table 4 compares the performance of the 4 mass trigger alternatives.

Using test data from runs 70579, 70601, 70602 and 71021, we have estimated the cross section for CMU-CMU dimuon triggers with and without a mass cut. For these test runs, the timing gate for CMU trigger stubs was widened to 200 ns compared to the standard Run 1 value of 100 ns. We considered only those events with two bin-0 CFT tracks matched to muon stubs. Given the 95% efficiency of the CFT and the low fake rate for the XFT, this requirement should not create a large bias for trigger rate estimates. For this study we used the hit-level CTC-XFT simulation. Figure 7 shows the trigger cross section as a function instantaneous luminosity for an opposite-charge, 2 GeV/c dimuon trigger with tight matching and the four mass trigger options. The trigger cross section is reduced from 100 nb to less than 15 nb at a luminosity of $10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Type	J/ψ Mass Resolution (GeV/c^2)	J/ψ Efficiency	Cross Section at $\mathcal{L} = 10^{32}$ (nb)
No cut	N/A	1.00	94 ± 6
XFT SL7	0.16	0.93	14 ± 2
M_T	$^{+0.11}_{-0.41}$	0.73	10 ± 2
HTDC	0.16	0.99	12 ± 2
TOF	0.16	0.90	9 ± 2

Table 4: Comparison of J/ψ mass trigger alternatives. The cross sections are all for a muon P_T threshold of 2 GeV/c.

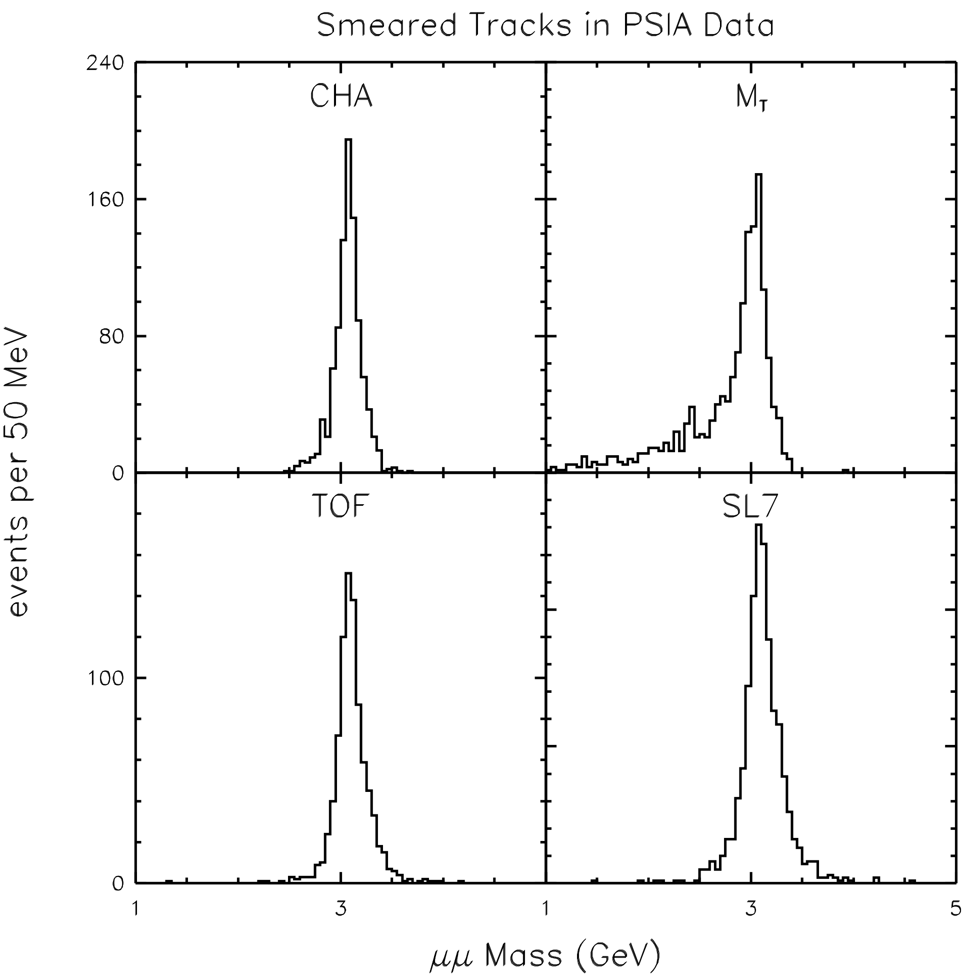

Figure 6: Reconstructed $J/\psi \rightarrow \mu^+\mu^-$ mass in PSIA events with a simulation of the possible mass triggers discussed in the text.

In order to extrapolate to the Run 2 muon threshold of 1.5 GeV/c, we have fit a power law to the inclusive dimuon trigger and SL7-mass trigger cross sections as a function of muon-track threshold. The stub timing cut is fixed at 200 ns. The result is shown in Figure 8. Although the extrapolation error is surely large, one draws the clear impression that the cross section for an inclusive dimuon trigger with a threshold 1.5 GeV/c may be enormous. The extrapolated cross sections are 160 nb for the opposite-charge trigger and 20 nb for the mass trigger. Based on Run 1 experience, we can estimate that a fully inclusive dimuon trigger without the charge cut and including CMX would probably have about a factor 4 higher cross section (640 nb) which corresponds to 125 Hz at $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The mass trigger rate including CMX is probably a factor of 2 higher than the CMU-CMU trigger estimates here, or 8 Hz at $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Thus, the cross section for a trigger with a mass cut should be well-controlled, and it will be possible to have a highly efficient $J/\psi \rightarrow \mu^+\mu^-$ trigger at the lowest possible momenta without the bias introduced by the SVT.

We have estimated the improvement in trigger rates for inclusive 12 GeV/c CMX muons using test data from runs 76399 and 76529. We consider only those events that a trigger with a bin 2 (nominally 3.4 GeV/c) CFT track matched to a high-P_T CMX trigger stub. The estimate of the Run 2 Level 2 trigger rate is estimated using XFT-CTC simulated tracks with the RUN2TRS tight matching requirement. The cross section estimates have been scaled for the increased CMX coverage in Run 2. Figure 9 shows the results of the study. Requiring a SL7 segment consistent with a track projecting to the CMX lowers the trigger rate by about a factor of 2. Requiring that it be a high-P_T segment gives another factor of 2. However, some of the reduction may be due to the fact that the SL7 segment comes from a good offline track its existence may in and of itself reduce the fake rate. Also, improved shielding of the CMX from the beam line should lower trigger rates as well, so the additional factor from the track z matching may not be as large as is estimated here.

The two studies described above highlight the capabilities of stereo segments. The prediction for the CMX at 12 GeV/c corresponds to a rate of 20 Hz at a luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ which would ensure that the CMX trigger could be fully efficient and would not require a prescale at the highest luminosities. The factor of 8 reduction in opposite charge dimuons with the J/ψ mass cut is highly efficient and the rate is comparable to other possible mass trigger schemes. The SL7 mass trigger should be well-behaved at the lowest possible muon threshold and the highest luminosities, ensuring full efficiency for $J/\psi \rightarrow \mu^+\mu^-$ decays.

5 Cost, Schedule, and Manpower

5.1 Cost Estimate

Table 5 details the R&D cost for the XFT3D system. We assume that we will go through a prototyping step for both the stereo Finder board as well as the L2 interface

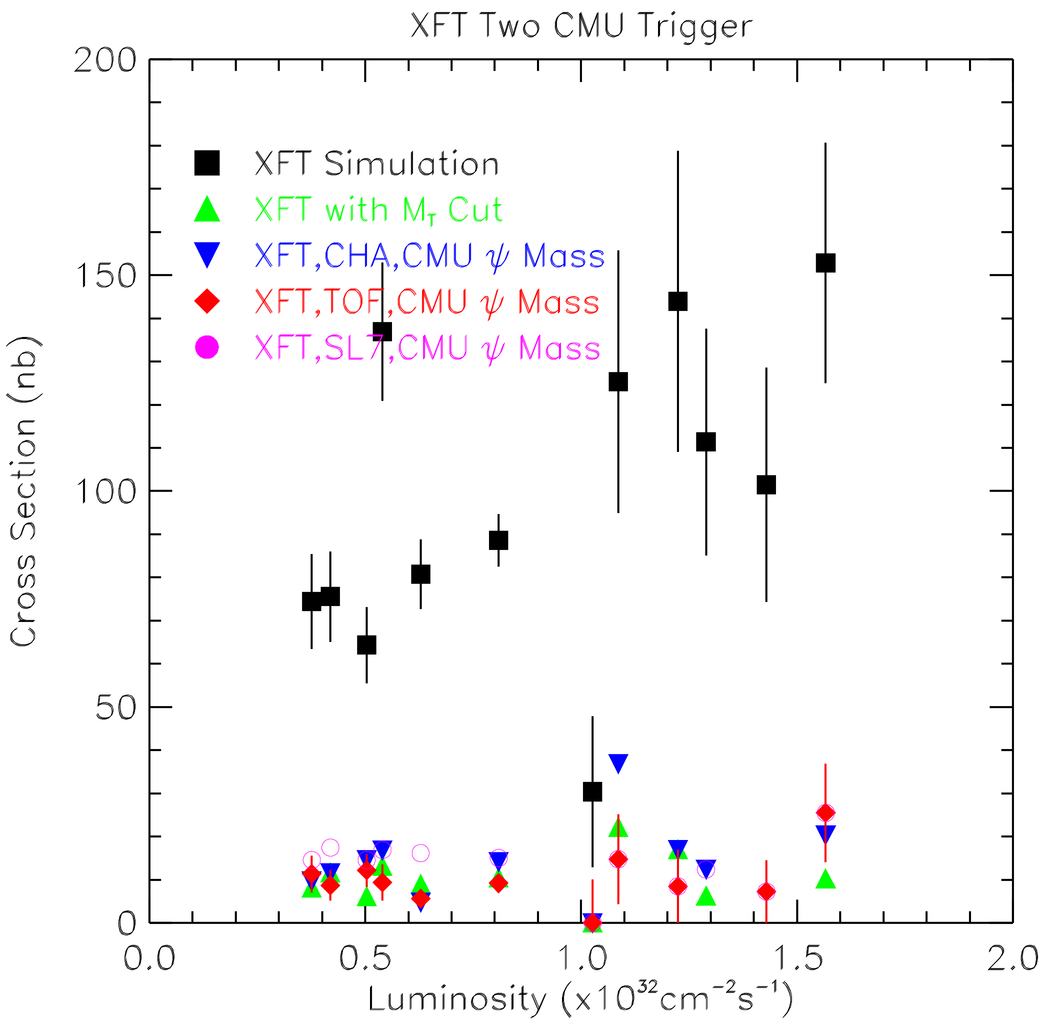

Figure 7: Cross section vs. instantaneous luminosity for an opposite-charge, 2 GeV/c dimuon trigger (squares) and the four mass trigger options discussed in the text. The circles show the results for the SL7 segments. The trigger cross section is reduced from 100 μb to less than 15 μb at a luminosity of $10^32\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

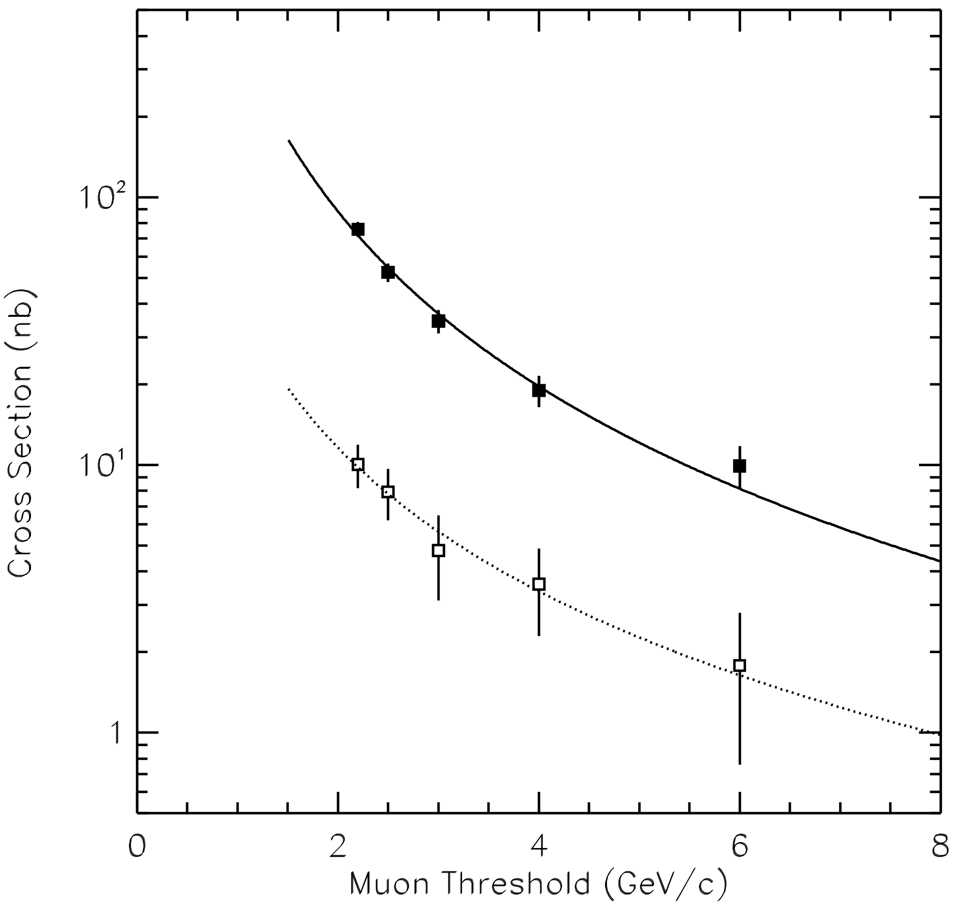

Figure 8: Dimuon trigger cross section vs. instantaneous muon threshold. Solid points are for tightly matched opposite charge pairs. The open squares are the SL7 mass trigger rates. The fit is to a power-law form. The points are not strictly independent, but the trend is clear.

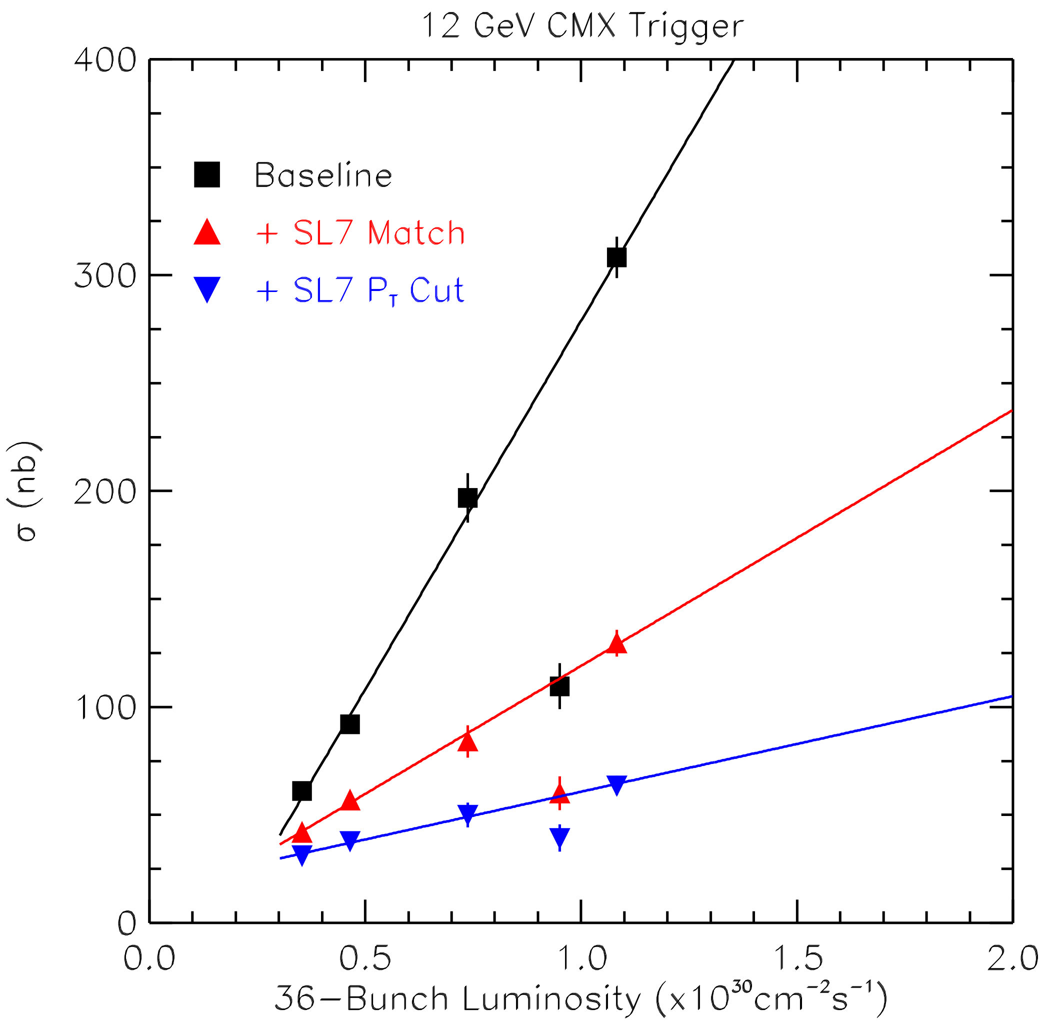

Figure 9: Estimated trigger cross section for 12 GeV/c CMX muons. Upper line and points do not include SL7 z information. The lower sets include all segments or only high-P_T segments. The points just below 1×10^{32} are excluded from fit because of a probable luminosity counting problem.

board. Given that much of the Finder board will be tested using the axial prototype, it is possible that the prototype will in fact be the production board. However, this is not assumed in either the R&D cost table or in the production cost table described below.

Table 6 details the cost for the production stereo Finding system. Most of the cost is in the purchase of boards already designed for the XFT2D system. The number of spares includes those needed for test stands. Contingency has not been added to the cost in the table, but note that most of the chip costs use present value quotes.

Desription	Quantity	Unit Cost	R&D Cost
VME Crate	1	$17,000	have
Finder Card Unloaded	2	$1000	$2,000
Finder Discrete Components	2	$300	$600
Finder Altera chips	12	$145	$1,740
Input section chips	12	$100	$1,200
Output section chips	6	$100	$600
Finder Backplane	1	$1,800	$1,800
Finder Transition module	2	$600	$1,200
TDC Transition module	2	$600	$1,200
TDC Mezzanine Card	2	$600	$1,200
L2 Receiver Test board	1	$5000	$5,000
L2 Interface board	2	$5000	$10,000
Connector/Cables			$500
Total			$27,040

Table 5: R&D cost estimate for the XFT3D system.

5.2 Schedule Estimate

Development of an axial Finder prototype is currently underway, and we expect that a fully loaded board will be available for testing in November, 1997. Allowing sufficient time for testing, we expect to begin work on the production board no earlier than

Desription	Production Quantity	Spares	Production Unit Cost	Production Cost
VME Crate	1	0	\$17,000	\$17,000
Finder Card Unloaded	18	6	\$1,000	\$24,000
Finder Discrete Components	18	6	\$300	\$7,200
Finder Altera chips	108	36	\$145	\$20,880
Input section chips	108	36	\$100	\$14,400
Output section chips	54	18	\$100	\$7,200
Finder Backplane	1	2	\$1,800	\$5,400
Finder Transition module	18	6	\$600	\$14,400
TDC Transition module	54	10	\$600	\$38,400
TDC Mezzanine Card	54	10	\$600	\$38,400
L2 Interface board	1	2	\$5000	\$15,000
Connector/Cables				\$10,000
Total				\$212,280

Table 6: Production cost estimate for the XFT3D system.

spring 1998. Only after the axial production board has been sent out for fabrication do we expect to begin work on the stereo Finder board. However, note that all of the functionality of the stereo Finder boards, with the exception of sending data to the L2 interface board, will be tested using the axial Finder board prototype. Based on these assumptions, it is likely that the first stereo Finder board will be the production board, and that it will be available sometime around fall, 1998.

The development of the the L2 interface board will be done in parallel by the UCLA group. We envision a 3 step process for this:

1. Design and build a simple VME board that transmits and receives optical signals. (controllable/readable by VME).

2. Design and build a prototype interface board with full L2 functionality.

3. Final production of L2 interface boards.

Item 1 can be finished by winter 1998, while item 2 probably will not be done before the end of summer, 1998. This is reasonably well matched to when the stereo Finder board will be available.

5.3 Manpower Estimate

The manpower needed for production of the stereo Finder boards will come from the existing XFT effort. The Finder/Linker subsystems are a joint effort of the Ohio State University and Fermilab groups, which total 5 physicists, and 4 engineers, and a small number of technicians. The manpower needed for the development of the L2 stereo interface board(s) will come from UCLA. This effort will involve 2 physicists.

6 Conclusion

The addition of stereo segment finding and the linking of these segments to XFT2D tracks in the Level 2 processor provides a means of reducing the trigger rates at Level 2, by factors ranging from 4-8 depending on trigger type. In addition, the implementation of the hardware involves very little R & D work. Since the stereo Finder boards would be nearly identical to the axial Finder boards, much of the prototyping and testing will naturally take place with the axial Finder boards. The cost of the system is quite modest, with more than 50% of the cost allocated towards the purchase of cards already designed for the axial Finder boards. This project has the rare qualities in that it offers large benefits for little cost and effort.

References

[1] National Semiconductor, "LVDS Owner's Manual and Design Guide" (1997); T. Shaw, S. Holm and R. Klein, "Cable Tests," unpublished note (1996).

[2] M. Campbell and C. Murphy, "The Magic Bus," see the following web page for details:
http://umaxp1.physics.lsa.umich.edu/cdf/documents.htm#Upgrade

[3] J.D. Lewis, M. Tecchio, K. Byrum, CDF/4321, "Guide to Run 2 Trigger Simulations", (in preparation).

[4] See the XFT web pages for details:
http://www-cdf.fnal.gov/upgrades/daq_trig/xft/xft_home.html.