

L2 XFT Stereo Upgrade Overview

C. Cox, D. Cox, R. Erbacher, R. Forrest, A. Ivanov, A. Lister, A. Soha
University of California, Davis

S. Holm, T. Liu, P. Wilson
Fermilab

Abstract

This document describes the XFT stereo upgrade for the L2 trigger system, starting from the Stereo Finder boards, up to the XFT stereo track algorithm implementation in the L2 PC. This note also describes the simulation tools available for developing triggers using stereo tracking information and discusses the implementation of the versions of the high p_T lepton triggers that use stereo tracking at L2.

Contents

1	System Overview	1
2	L2 Pixel Design	3
3	Data Sparsification	5
4	T2FD Bank Format	5
5	Data Format for the Transfer Between Finder and Pulsar Boards	7
6	XFT Pulsar Boards	10
6.1	Stereo Finder to Pulsar Mapping	11
6.2	TP2D Bank Format	12
6.3	L1 Abort	12
6.4	Event Truncation	13
7	Data Transfer from the Pulsar to the L2 PC	14
8	TL2D Bank Format	15

9 Stereo Tracking Algorithm Implementation	17
9.1 Initialization	18
9.2 Unpacking	18
9.3 Clustering and Masking	18
9.4 Combinatorics	19
10 Algorithm Performance and Stereo Track Triggers	20
10.1 Reconstruction Algorithm Performance	20
10.2 ‘Best Track’ versus ‘Valid Track’	21
10.3 CMX	22
10.4 CMUP	22

1 System Overview

The XFT Stereo upgrade provides many benefits over the purely axial triggering system used previously. One of the aims of this project is to reduce the rate of fake tracks in many triggers. Fake rates increase very rapidly with luminosity, much faster than the real track rates. By removing as many of the fake tracks as possible at trigger level, it is possible to keep these triggers, without a pre-scale, up to much higher instantaneous luminosities. The L1 path is used to confirm the existing XFT track, reconstructed with the axial COT layers only, goes through the stereo layers at the expected locations. At L2 the segmentation is much finer than at L1 allowing a better fake rejection rate and also providing information about the position of the track. In particular it is possible to measure the angle θ of the track with respect to the beam axis as well as the distance z from the center of the detector along the beam axis and use this information to point the track to other detectors. This 3D tracking opens up several additional capabilities such as trigger level multi-track mass calculations or isolation requirements and z-vertex reconstruction at L2.

Because of the slight angular inclination with respect to the beam-pipe of the stereo layers, the $\cot(\theta)$ of the track can be calculated from the distances between where the track is expected to cross the stereo layers and where it actually does cross. This is illustrated in figure 1. SL3 and SL7 have the same stereo angle whereas SL5 has the opposite angle. The $\cot(\theta)$ of the track is measured using the difference in the displacements between the expected track position were there no stereo angle and the measured one as follows:

$$\cot(\theta) = \frac{\Delta(SL7) - \Delta(SL5)}{r_7 - r_5} = \frac{\Delta(SL5) - \Delta(SL3)}{r_5 - r_3} = \frac{\Delta(SL7) - \Delta(0)}{r_7} \quad (1)$$

where r_7 , r_5 and r_3 are the radii of the 3 stereo superlayers, $\Delta(SL7)$, $\Delta(SL5)$ and $\Delta(SL3)$ are the deviations between the expected and measured track positions for the 3 stereo superlayers and $\Delta(0)$ is the location with respect to the center of the detector where the track would cross the beam-axis.

To illustrate the fake rate rejection by using the stereo XFT system to confirm the tracks, the measured and expected trigger rates for the CMX trigger are shown in figure 2 as a function of instantaneous luminosity. At luminosities around $150 \cdot 10^{30} \text{ cm}^2 \text{s}^{-1}$, an additional factor of 2 in the trigger rate can be gained by using L2 stereo confirmation and pointing as opposed to the L1 stereo confirmation. At the predicted highest luminosity running, around $300 \cdot 10^{30} \text{ cm}^2 \text{s}^{-1}$, the additional gain is expected to be a




Figure 1: Schematic illustration of the axial and stereo COT layers.

factor of 4 over stereo confirmation.

The system used to obtain the XFT stereo information at L2 is described here and shown schematically in figure 3. The Stereo Finder board is used both for the L1 and L2 paths. There are 36 finder boards, 12 for each superlayer, split into two different crates (b0xft06 and b0xft07). Each board contains 2 chips (chip A and chip B) each of which look at only some of the input fibers. For axial track creation, the input fibers exit the COT XTC cards. The COT hits are combined into track segments and the ϕ position and slope are measured. The slope bits are set according to the L2 90 bit format described in section 2. These 90 bits are concatenated into the 12 Level 1 bits which are then sent to the SLAM cards. The higher resolution data is sparsified as described in section 3 then sent both to the T2FD bank, see section 4, and to one of the three PULSAR boards via the transition boards on the backplanes of the Finder crates. The data format for the Finder to PULSAR connection is described in section 5. The PULSAR boards simply merge and concatenate the stereo pixel information from several boards. The order in which the pixel information from each finder is sent out of the PULSAR is defined by a mapping described in section 6. The PULSAR sends the data over optical fibers into the L2 PC, described in section 7. The L2 PC runs a track reconstruction algorithm and L2 trigger algorithms, as described in section 9. Because of timing constraints, the stereo reconstruction algorithm is only run on those tracks that already pass all other L2 trigger requirements. Section 8 describes how this task is facilitated by the use of the TL2D bank; the stereo track information for the processed tracks is appended to the track information in that bank. The stereo track reconstruction algorithm will only be run on a subset of triggers and can be used in a number of different ways; e.g. using the track $\cot(\theta)$ to verify that the track points to a muon stub. A tool, presented in section 10 is used to simulate the L2 stereo tracking algorithm. It enables the implementation of new triggers or the improvement of existing ones.




Figure 2: Measured and predicted CMX trigger rates using no stereo confirmation, L1 stereo track confirmation and L2 track confirmation .




Figure 3: Schematic diagram of the data flow for the L2 XFT Upgrade.

2 L2 Pixel Design

Note: CDF Note 7789 contains a more detailed definition of the L1 and L2 pixel mappings [1]

To provide flexibility and simplicity, the design of the L2 pixels is similar to that of the L1 pixels ¹. A cell contains 6 pixels at L1, each with 2 slope bits: a positive and a negative slope bit. When both bits are set, the track is either a high p_T track (i.e. essentially straight, $p_T > \text{xxx GeV}/c$) or the hits are consistent with either a positive or a negative curvature. The basic unit of a cell at L1 is thus represented by $6 \times 2 = 12$ bits. At L2, a cell contains 6 sub-cells, each of which contain 3 pixels, each with 5 slope bits: positive curvature low p_T ($p_T < \text{xxx GeV}/c$), positive curvature higher p_T ($\text{xxx} < p_T < \text{yyy GeV}/c$), high p_T ($p_T > \text{yyy GeV}/c$), negative curvature higher p_T and negative curvature low p_T . The basic unit of a cell at L2 is thus represented by $6 \times 3 \times 5 = 90$ bits. The basic bit structure for L1 cells and L2 sub-cells are depicted in tables 1 and 2, respectively. It should be noted that any given pixel can have multiple tracks passing through it with different slope values, resulting in multiple slope bits being set. For SL3 at L2, the slope bits are dropped and the ϕ size is the same as the L1 pixels. This allows us to significantly reduce the total data volume in the L2 path without compromising much on the stereo reconstruction performance, as SL3 contains the least discriminating power of all stereo SL.

Bit #	11	10	9	8	7	6	5	4	3	2	1	0
Pixel #	5		4		3		2		1		0	
Slope	+	-	+	-	+	-	+	-	+	-	+	-

Table 1: L1 pixel format for a single cell. The 2 slope bits correspond to positively and negatively charged tracks or positive and negative track curvatures. Both bits are set for high p_T tracks or if the hits can be fitted with either a positive or a negative curvature track.

Bit #	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Pixel #															0
Track p_T	+	++	hpt	--	-	+	++	hpt	--	-	+	++	hpt	--	-
Curvature	++	+	0	-	--	++	+	0	-	--	++	+	0	-	--

Table 2: L2 pixel format for a single sub-cell for SL5 and SL7. The slope bit assignments are shown as a function of track p_T and of track curvature. The 3rd slope bit is set for high p_T tracks (hpt). For SL3 the internal Stereo Finder resolution is the same but only 1 bit is sent for each sub-cell.

3 Data Sparsification

The well-known idea behind sparsification is to send data only for portions of the detector that contain at least one hit. Using data sparsification results in a smaller data packet size and faster data transfers. The sparsification scheme adopted for the L2 path provides the data in two distinct sections. The first 10,

¹Note: In our naming convention a L1 pixel corresponds to the same ϕ size as a L2 sub-cell

18, and 18) words for SL3, SL5, and SL7, respectively, are used as the mask section. Each sub-cell is represented by a single bit. If the bit is set (=1) then at least 1 of the 15 slope bits (5 bits for each pixel) is set for this sub-cell. For SL3, only the header, mask section, and trailer are sent. For SL5 and SL7, the data section follows the mask section. The data section contains the 15 bits only for each sub-cell for which the mask bit was set. The information for sub-cells with no hits (for which all 15 bits are 0) are not sent. Numerous studies were carried out to investigate the the data transfer time from the Finder Boards to the L2 Pulsar boards [2]. This time was found to not exceed $4 \mu s$ when the data is sparsified.

4 T2FD Bank Format

Note: This section, included here for completeness, is a copy of the section included in CDF note 4152 [3].

The T2FD bank is designed for diagnostic readout, and to facilitate development of algorithms at level 2. The pixel information from the stereo finders for superlayers 3, 5, and 7 is made available in this bank through the format shown in table 3.

The format reflects the fact that the outermost three stereo superlayers, here called SL3, SL5, and SL7, each have 12 finder boards, but those boards are split among two crates (18 boards each). The current hardware plan is to have six boards of each superlayer type in each of the two crates. The T2FD bank definition has a flexible format for the finder pixel data, to include the possibility of sparsification. Note that the number of cells per finder board depends on the superlayer (SL3, SL5, and SL7 have 20, 28, and 36 cells/finder, respectively), so the block length can be different for each of the two blocks, depending on how the finders are distributed among the two crates, regardless of whether or not sparsification is used. For the stereo information available to level 2 for SL5 and SL7, each cell consists of 18ϕ bins, and hits in any given ϕ position can have any combination of up to five different slopes. For the purpose of sparsification, each cell is split into 6 *sub-cells* that each contain 3 ϕ positions (again with 1 to 5 of the available slopes being set). For the stereo information available to level 2 for SL3, each cell consists of 6 ϕ bins, with no slope information. This allows a fixed data volume size for SL3.

Tables 4 and 5 shows the details of the pixel data format for SL5 and SL7 as well as SL3. The information from each of the 18 successive finders within a given crate appears sequentially in 18 repeated sections, each with the format shown in table 4 for SL5 and SL7 and table 5 for SL3. The first word contains header information, including superlayer # (2 bits), buffer # for finder chip A (2 bits), buffer # for finder chip B (2 bits), the number of 32 bit words in the buffer [excluding the header word] (7 bits), bunch counter # for finder chip A (7 bits), and bunch counter # for finder chip B (7 bits). The header word is identified by the fact that bit 15 is set to 1 and bit 14 is set to 0. Also, bits 13, 30, and 31 are set to 0.

The next 5, 9, or 9 words, for SL3, SL5, or SL7, are used for the sparsification mask. In the case of SL3 this is the only data that is sent. For these 32-bit words, bits 0-5 serve as the mask for the 6 sub-cells of cell#0, bits 6-11 serve as the mask for the 6 sub-cells of cell#10 or cell#18 (depending on the SL), bits 12-14 are unused, and bit 15 has the value of 0, bits 16-21 serve as the mask for the 6 sub-cells of cell#1, bits 22-27 serve as the mask for the 6 sub cells of cell#11 or cell#13 (depending on the SL), bits 28-30 are unused, and bit 31 has the value 0. The actual pattern of cells in the mask section is somewhat complex due to the hardware implemenation, but here is an overview. The stereo finders for each of the superlayers

“T2FD”
Bank number
Bank version
Bank length
Bank type (I*4)
Number of blocks (=2)
Pointer to block 0
Pointer to block 1
Pointer to end of data
Block 0: Number of cards (=18)
Pointer to card 0
...
Pointer to card 17
Pointer to end of block
Finder pixel data
...
Finder pixel data
Block 1: Number of cards (=18)
Pointer to card 0
...
Pointer to card 17
Pointer to end of block
Finder pixel data
...
Finder pixel data

Table 3: Structure of the T2FD bank.

have two chips, called chip A and chip B. Two different versions of firmware have been written for the chips, such that they can either handle 10 cells or 18 cells (here, we ignore the added complexity of the fact that the chips also can readout neighboring cell information). For SL3 there are 20 cells (generically numbered 0 through 19 for our purposes here). Chip A reads out cells 0-9 and chip B reads out cells 10-19. The firmware that provides the cell information for L2 (for the T2FD bank and for the L2 PULSAR system), puts the SL3 cell masks into the following order: 0, 10, 1, 11, 2, 12, 3, 13, 4, 14, 5, 15, 6, 16, 7, 17, 8, 18, 9, 19. For SL5 there are 28 cells (generically numbered 0 through 27). For simplicity of firmware, chip A handles cells 0-17 and chip B handles cells 18-27. Therefore the cell mask ordering for SL5 is the following: 0, 18, 1, 19, 2, 20, 3, 21, 4, 22, 5, 23, 6, 24, 7, 25, 8, 26, 9, 27, 10, X, 11, X, 12, X, 13, X, 14, X, 15, X, 16, X, 17, X, where each X means that there are six unused bits in the mask due to the fact that chip B is “empty” while chip A is still being read out. For SL7 there are 36 cells (generically numbered 0 through 35). Here, chip A handles cells 0-17 and chip B handles cells 18-35. The cell mask ordering for SL7 is the following: 0, 18, 1, 19, 2, 20, 3, 21, 4, 22, 5, 23, 6, 24, 7, 25, 8, 26, 9, 27, 10, 28, 11, 29, 12, 30, 13, 31, 14, 32, 15, 33, 16, 34, 17, 35. Each mask word specifies the masks for 4 sub-cells, except for the last 4 mask words for SL5, which each have the masks for 2 sub-cells alternating with two

“empty” slots.

The correspondence between the above generic cell numbers within a particular finder and the global cell location in the COT is also somewhat complex. For SL3, chip A on the finder that reads out the $\phi = -90$ region of the drift chamber reads global cells 230-239 and chip B reads global cells 0-9. For SL5, chip A on the $\phi = -90$ finder reads global cells 322-335 and global cells 0-3, while chip B reads global cells 4-13. For SL7, chip A on the $\phi = -90$ finder reads global cells 414-431, while chip B reads global cells 0-17. For the finders that read out each of the other 30 degree segments of the COT, it can more simply be stated that chip A reads out the global cell positions with smaller cell numbers while chip B reads out the global cell positions with larger cell numbers.

Within the mask, a bit set to 1 corresponds to a sub-cell that contains at least one hit, while a bit set to 0 corresponds to a sub-cell with no hits. If sparsification is not desired, all of the cell masks can be set to 1’s. For SL3, the mask section is directly followed by the trailer described below. For SL5 and SL7, the next set of words provide the actual hit data for any sub-cells that were indicated by a bit set to 1 in the mask. The sub-cells appear in a *different* order than as given in the mask. They appear in numerical cell order (0-19, 0-27, or 0-35, depending on SL). Each word contains the hit information for the 3 ϕ pixels of 2 sub-cells (the last word will sometimes contain just 1 sub-cell, and the remaining bits will belong to the trailer). Within the data section, bit 15 has the value 0, and bit 31 has the value 0, except sometimes for the last word where bit 31 might be part of the trailer section. For each pixel, there are 5 bits corresponding to the 5 slope bins. If there is an odd number of sub-cells being read out, then the last word is a combination of data and trailer, in which case bits 16-29 are used for the trailer information, bit 30 is set to 1, and bit 31 is set to 1. If there is an even number of sub-cells being read out, then one additional word appears with bits 0-13 used for the trailer information, bit 14 set to 1, bit 15 set to 1, and bits 16-31 unused. The trailer contains three bits for the turn #, four bits for the finder #, four bits for the fiber #, and the bit-pattern 101 appearing in bits 11-13 or 27-29.

If sparsification were not used (and the cell masks would then all be set to 1’s), then there would be 95 and 119 words required for each finder in SL5 and SL7, respectively.

5 Data Format for the Transfer Between Finder and Pulsar Boards

A 16-bit communication protocol is used for data transfer from the L2 output chips of the Finder boards to the Pulsar mezzanine cards. The data from each Stereo Finder is packed into a packet, one per Finder per event. The size of the data packet will in general depend on how busy a particular event is, which is reflected by the COT hit occupancy, and the number of the stereo segments found. The data packet starts with the header, which has two 16-bit words, and ends with a single trailer word (see Table 6). Bit 15 in each of the 16-bit words sent from the Finder to the Pulsar is reserved to define the beginning and the end of the data envelope, which is recognized by the PULSAR firmware. It is set to 1 in the first and the last word of the data envelope, and otherwise set to 0. Bit 14 in the first and the last word is used to distinguish the beginning (set to 0) from the end (set to 1). The first header word starts with 2 bits for the superlayer number (coded as: bit 1=0, bit 0=0 for SL 3; bit 1=0, bit 0=1 for SL 5; and bit 1=1, bit 0=1 for SL 7). Bits 2-3 (4-5) give the buffer number that was read out for chip A (B) on the Finder. The second header word

Bit #	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Header	0	0	Finder A		Finder B		Bunch Counter		Finder A		Finder B		#32 bit words in buffer (excluding header)		Finder A		Finder B		Fin A		Fin B		Buffer #		SL #							
Mask	0	Unused	Unused	Cell #19	Cell #21	Cell #1	Cell #3	Unused	Unused	0	0	Unused	Unused	Cell #18	Cell #20	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2			
Mask	0	Unused	Unused	Cell #21	0	0			
Mask	0	0	0			
Data	0	Sub-cell 1, Pix 3	...	Sub-cell 1, Pix 2	...	Total of 9 words for SL5 and SL7	Sub-cell 1, Pix 1	Sub-cell 0, Pix 3	Sub-cell 0, Pix 2	0	0	Sub-cell 0, Pix 3	0	Sub-cell 0, Pix 2	0	Sub-cell 0, Pix 1	0			
Data	0	Up to 84 data words for SL5	0	0				
Data	0	Up to 108 data words for SL7	0	0			
Data	1	1	1	0	1	Fiber #	Finder #	Turn #	0	Sub-cell N, Pix 3	Sub-cell N, Pix 2	Sub-cell N, Pix 1	0	Sub-cell N, Pix 3	Sub-cell N, Pix 2	Sub-cell N, Pix 1	0		

Table 4: Format of header, mask, pixel data, and trailer for T2FD bank, corresponding to the data from one finder for SL5 and SL7. The mask section is used to sparsify based on sub-cells that are 1/6th of a cell (details are given in the text). For the data section, the notation “Sub-cell 1, Pix 1” means that the five corresponding bits indicate which combination of the 5 slope bins for ϕ pixel #1 of sub-cell #1 have hits. Note that when bit 15 or bit 31 has the value 1, then bit 14 or bit 30 is also a control bit, and indicates the beginning or end of the envelope (these control bits sometimes appear in the middle of the word since the structure of this data format was driven by the 16-bit word format used for a stereo finder to PULSAR transfer). If there is an odd number of sub-cells with data, then the last word is a combination of data and trailer word as shown above. In this case bits 16-29 are used for the trailer, bit 30 is set to 1, and bit 31 is set to 1. If there is an even number of sub-cells with data, one additional trailer word appears with bits 0-13 used for the trailer, bit 14 set to 1, bit 15 set to 1, and bits 16-31 unused.

Bit #	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Header	0	0	Finder A		Finder B		Bunch Counter		Finder A		Finder B		#32 bit words in buffer (excluding header)		Finder A		Finder B		Fin A		Fin B		Buffer #		SL #							
Mask	0	Unused	Unused	Cell #11	Cell #13	Cell #1	Cell #3	Unused	Unused	0	0	Unused	Unused	Cell #10	Cell #12	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2	Cell #0	Cell #2			
Mask	0	0	0			
Mask	0	0	0			
Data	1	1	1	0	1	Fiber #	Finder #	Turn #	0	Sub-cell N, Pix 3	Sub-cell N, Pix 2	Sub-cell N, Pix 1	0	Sub-cell N, Pix 3	Sub-cell N, Pix 2	Sub-cell N, Pix 1	0	

Table 5: Format of header, mask, pixel data, and trailer for T2FD bank, corresponding to the data from one finder for SL3. Note that the format is similar to SL5 and SL7 but without the data section.

contains the bunch counter information, ranging from 0 to 127, for Finder chip A in bits 0-6 and Finder chip B in bits 7-13. In the trailer word, bits 0-2 give the turn number, which specifies b0 markers. Bits 3-6 give the finder ID number, bits 7-10 give the fiber ID number, and bits 11-15 are set to 10111 in that order.

At L2 18×5 stereo pixel binning is employed (18ϕ -pixel bins by 5 slope bins), so that each cell requires 90 bits of information. To speed up the data transfer the information from each cell is sparsified into 6 parts (sub-cells), and only those sub-cells with at least one segment on are passed to L2 (see Section 3). The data field has the structure shown in tables 7, 8, 9.

The first several words contain the list of sub-cells masks, where one bit is reserved per sub-cell, i.e. 6 bits specify one cell. If a particular sub-cell has at least one hit, then the bit is set to 1, otherwise it remains 0. Each word specifies information for two cells.

To speed up the sparsification and mask creation in the Stereo Finder firmware, the order of cells in the mask list is the following. The first word contains masks for cell #0 followed by cell # 10, 18, and 18 for superlayers SL = 3, 5, and 7, respectively. The second word contains masks for cells #1 and # 11, 19, and 19 and so on. The bits 12-14 in each word remain unassigned. For the *SL5* Stereo Finder bits 6-14 are unassigned starting from 11-th word in the mask section. The size of the mask list is fixed and is 10 words for super-layer 3 and 18 words for super-layers 5 and 7.

The second part of the data field contains the sub-cell stereo pixel information for those sub-cells, which are not masked (corresponding to bit=1 in the mask section). The sub-cells are appended in numerical order corresponding to cells: 0,1,2, ... This is different from the mask section ordering. Masked sub-cells (bit=0) are discarded. One word is assigned to each sub-cell.

Bit #	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																				
Header	1	0	0	unused						Buffer #		SL #																								
Header	0	0	Bunch Counter																																	
		Finder B						Finder A																												
Data	0	Data																																		
...	0	...																																		
Data	0	Data																																		
Trailer	1	1	1	0	1	Fiber #		Finder #		Turn #																										

Table 6: Data envelope structure between Finder and Pulsar boards.

6 XFT Pulsar Boards

The data leaving each of the 36 Stereo Finders is sent to one of the 3 Pulsar boards. The main purpose of the Pulsar board is to concatenate the 36 Stereo Finder inputs into 3 inputs into the L2 PC. The data remains mostly unchanged except for the cases mentioned in the next paragraph and the trailer word. The

Bit #	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
10 words for SL 3	n.a. bits	Cell # 10										Cell # 0			
	n.a. bits	Cell # 11										Cell # 1			
	...										Sub-cell bit masks				
	bit=1 if has a hit										bit=0 if no hit				
				
	n.a. bits	Cell # 19										Cell # 9			
				
				
				
				

Table 7: Data field for the superlayer $SL = 3$ for Finder to Pulsar data transfer.

Bit #	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
18 words for SL 5	n.a. bits	Cell # 18										Cell # 0			
	n.a. bits	Cell # 19										Cell # 1			
	...										Sub-cell bit masks				
	bit=1 if has a hit										bit=0 if no hit				
				
	n.a. bits	Cell # 27										Cell # 9			
	...										Cell # 10				
				
	...										Cell # 17				
	Sub-cell #0	Pixel # 2										Pixel # 0			
Sub-cell #1	Pixel # 2										Pixel # 1				
...										...					
Last Sub-cell	Pixel # 2										Pixel # 1				Pixel # 0

Table 8: Data field for the superlayer $SL = 5$ for Finder to Pulsar data transfer.

mapping between the Stereo Finders and the XFT Pulsars is described in sub-section 6.1.

In order to reduce the size of the inputs arriving to the L2 PC, several additions were made to the initially envisaged configuration. The first of these has been discussed throughout and deals with using only the mask section in SL3 and not sending the higher resolution data out to the pulsars. Two further improvements were suggested, the first being of sending the data to the L2 PC only when needed. This is referred to as the L1 abort functionality and is described in sub-section 6.3. The second is the truncation of large events at the output of the Pulsar boards and is described in sub-section 6.4.

The SLINK trailer in the XFT Pulsar boards is modified from the default (0xeeee) to reflect the new L1 Abort and event truncation functionalities. Bit 2 is set (1) if the Pulsar boards are not being aborted

Bit #	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0														
18 words for SL 7	n.a. bits	Cell # 18				Cell # 0																							
	n.a. bits	Cell # 19				Cell # 1																							
	...																												
	Sub-cell bit masks																												
	bit=1 if has a hit																												
	bit=0 if no hit																												
	...																												
	n.a. bits	Cell # 35				Cell # 17																							
Sub-cell #0	Pixel # 2			Pixel # 1			Pixel # 0																						
Sub-cell #1	Pixel # 2			Pixel # 1			Pixel # 0																						
...																													
Last Sub-cell	Pixel # 2			Pixel # 1			Pixel # 0																						

Table 9: Data field for the superlayer $SL = 7$ for Finder to Pulsar data transfer.

based on L1 bits. Bit 0 is set (2) if that particular event is aborted. Bit 1 is set (1) if the event is truncated. Bits 3-15 are set to zero. Bit 16 is 1. The new trailer is shown in Table 10

Bit #	Setting
15: 3	0
2	L1 Abort disabled (1) / enabled (0)
1	Truncated Event (1)
0	Aborted Event (1)

Table 10: Bit assignment for new XFT Pulsar SLINK trailer word.

6.1 Stereo Finder to Pulsar Mapping

The Stereo Finders are arranged in 2 XFT crates (b0xft06 and b0xft07), with 18 boards in each crate (slots 3-20), in such a way that the first 6 slots of both crates correspond to SL3, the next 6 to SL5 and the last 6 to SL7.

In order to reduce the number of inputs to the L2 PC which exceed 1/4 of the total FIFO size in the FILARs, 128 words, all the inputs from SL3 are in a single XFT Pulsar board (XFT 1). The second Pulsar board (XFT 2) reads in the SL5 and SL7 outputs from crate b0xft06. The third reads those from b0xft07. This mapping is shown in Table 11. Figures 4 and 5 show the true representation in detector- ϕ of the hardware in the finder and the pulsar ordering, respectively.

Pulsar board XFT 1		Pulsar board XFT 2		Pulsar board XFT 3	
Finder Crate	Finder Slot #	Finder Crate	Finder Slot #	Finder Crate	Finder Slot #
6	3	6	9	7	9
6	4	6	10	7	10
6	5	6	11	7	11
6	6	6	12	7	12
6	7	6	13	7	13
6	8	6	14	7	14
7	3	6	15	7	15
7	4	6	16	7	16
7	5	6	17	7	17
7	6	6	18	7	18
7	7	6	19	7	19
7	8	6	20	7	20

Table 11: Mapping between Finder crate and slot number and concatenation ordering in Pulsar boards.




Figure 4: Schematic view of the 3 stereo superlayers showing the way the hardware reads out each phi slice in the finders [finder crate;slot number].




Figure 5: Schematic view of the 3 stereo superlayers showing the way the hardware reads out each phi slice in the Pulsar boards [Pulsar board; input number].

6.2 TP2D Bank Format

Note: This section, included here for completeness, is similar to the section included in CDF note 4152 [3].

The specific bank format for the XFT Pulsar banks is similar to that of all other CDF Pulsar boards with the exception of the SLINK Trailer.

The overall structure of the TP2D bank is shown in Table 12. For each Pulsar board, 6 cards are created in the TP2D bank; one each for the CONTROL1, CONTROL2, DATAIO1_1, DATAIO2_1, DATAIO2_1, and DATAIO2_2. In the case of the XFT Pulsar boards, only the CONTROL1 buffer is read out by default. Currently the 3 XFT Pulsar boards are in block 1 cards 0, 6, and 12.

The content of each XFT Pulsar card is the concatenation of the information from the 12 Stereo Finders with no words separating the boards. The overall TP2D bank structure contains 4 SLINK header and 2 SLINK trailer words. The overall TP2D bank structure for the XFT Pulsar boards is shown in table 13.

6.3 L1 Abort

The XFT L2 stereo reconstruction algorithms will only be run for a small number of triggers (such as those described in this note). Yet the data, by default, would be sent to the L2 PC for every event. A system was designed such that the data would only be sent to the L2 PC when the information might be needed. This is carried out in multiple stages. The TrackList Pulsar board [4] has the ability to read in a mask of L1 bits that define which events are potentially useful for L2 XFT. It then sends out a signal to the XFT Pulsar boards to tell them to keep the event or only send an empty data packet. the feature for the XFT Pulsar boards to listen or not to the signal from the TrackList is set in register 0x18 ² which is accessible through the CDF cardEditor variable L1ABORT. Setting the register to 1 disables this feature. It is expected that at high luminosity that up to 75% of events can be aborted using this feature. The full XFT L2 data would only be sent to the L2 PC in 25% of the events.

For aborted events, the output of the XFT Pulsars is shown in Table 14.

6.4 Event Truncation

An internal counter in the output registers of the XFT Pulsar boards was implemented that keeps track of the total number of words in a given time window. If the number of words in that time window exceeds the pre-defined threshold, the output is truncated. Only the trailer word is sent after that. The trailer word contains the information that the event is aborted.

This feature was designed as a fall-back option in case the other 2 options are not sufficient to enable the XFT L2 system to operate at the high luminosities. The L2 PC will recognize the event as being truncated and automatically accept the L2 XFT part of the relevant triggers.

It is possible to vary both the size at which events are truncated and the time interval over which the max size is accumulated. This information is stored in register 0x1c ³; bits 0-19 contain the time in ns over which the total size is integrated; bits 20-29 contain the truncation size threshold. Both of these variables can be set via the cardEditor. The default variables are a size of 1023 words and a time-integral of 16 ns; this implies that the size integration is only over a single event.

The output of the XFT Pulsars for truncated events is show in Table 15. The SLINK trailer contains in bits 16-31 the total number of words. Bit 2 is set to 1 for truncated events

7 Data Transfer from the Pulsar to the L2 PC

Each of the three XFT Pulsar boards has a single output, each containing the information from 12 finder boards. After leaving the PULSAR, the data is transferred to the L2 PC via optical links, so there are three optical inputs to the L2 PC. The PC has 2 PCI buses. One of the PCI buses has 2 FILAR cards connected to it. The other has 1 FILAR and 1 SOLAR card (used for the output to the L2TS board, that contains the information that is written into the TL2D bank). Because of the large sizes of the XFT inputs we split them over the 2 PCI buses. Two inputs are on 1 FILAR and the third is on the FILAR that shares the PCI bus with the SOLAR card.

²/vxreadreg b0l2pu02 4 0x0 0x18

³/vxreadreg b0l2pu02 4 0x0 0x1c

The specifications for the PC, FILAR and SOLAR boards can be found elsewhere.

8 TL2D Bank Format

Note: The first part of this section, included here for completeness, is a copy of the section included in CDF note 4152 [3].

The TL2D bank contains a wide range of information used in the level 2 trigger decision. For the purpose of the level 2 XFT stereo upgrade, it is used to store the stereo track parameters generated by the stereo tracking algorithm, running either offline or in the level 2 PC. The stereo track parameters, specified below, have been added to the XTRP card of the TL2D bank. We have recycled the previously defined, but unused, so-called “Match” word. The structure of the XTRP card of the TL2D bank is given in Table ???. Tables ?? and ?? show the data word structure for the XTRP track and match word, respectively. More details about the 32-bit words for XTRP track data and the End of Event word can be found in CDF note 4152 (section 31.16.4 at the time of this writing).

The structure of the XTRP card section is:

Displacement	Data Description
P3	Number of XTRP Tracks N_{TRK} (variable)
P3+1	Track 1 XTRP Data
P3+2	Track 1 Match Data
P3+3	Track 2 XTRP Data
P3+4	Track 2 Match Data
...	...
$P3+2*N_{TRK}-1$	Track N_{TRK} XTRP Data
$P3+2*N_{TRK}$	Track N_{TRK} Match Data
$P3+2*N_{TRK}+1$	Track End of Event Word

Each track has two words associated with it, the XTRP and Match words. The XTRP word is defined as:

XTRP Bit	Definition
31	XFT L2 auto-accept trigger
30:29	Unused
28:27	XFT L2 eta-gap Trigger Fired (West:East)
26:25	XFT L2 CMX Trigger Fired (West:East)
24:23	XFT L2 CMU Trigger Fired (West:East)
22:21	XFT L2 CMP Trigger Fired (West:East)
20	SL6/SL8 bit (Short track bit)
19	L1 stereo confirmation (SLAM) bit
18:12	Track curvature bin
11:0	Track phi

The Match word is defined as:

Match Bit	Definition
31:22	XFT L2 CEM Trigger Fired (Tower #)
21:12	XFT L2 Cotan theta for “best track”
11:2	XFT L2 Z5 for “best track”
1	XFT L2 Stereo track successfully reconstructed
0	XFT L2 Attempted reconstruction of stereo track

Prior to the summer 2006 shutdown (data taken through run 212133, 0i period 7), the L1 stereo confirmation (SLAM) bit was referred to as the “isolation” bit but was unused in the trigger. The stereo confirmation (SLAM) bit definition began with trigger table PHYSICS_4_01 (run 226196).

The last bit field in the XTRP word is “Track phi”, which has the 3 XFT mini- ϕ bits appended to 9 bits of linker ID. The lower edge of bin 0 is $\phi = 0$. The true ϕ in radians is given by $\phi = 2\pi * \text{bin}/2304$. If one right shifts the “Track phi” value by 3 bits (or equivalently divide by 8 and round down), then one gets the “XFT Linker ID” that identifies the XFT chip. Since an XFT chip outputs at most one track per event, this uniquely identifies an XFT track within the event. The “XFT Linker ID” runs from 0 to 287.

The information related to the L2 XFT stereo reconstruction and trigger algorithms is contained in these two words. The attempted bit is set whenever stereo reconstruction is attempted on that track. The reconstructed bit is set whenever the track is successfully reconstructed. If the track falls within a masked on XTC, or if the data was truncated in the XFT Pulsar boards, the L2 XFT part of the trigger code is auto-accepted and the auto-accept bit is set. For these cases the attempted bit is set but not the reconstructed one. The track with the lowest χ^2 is called the “best track”; the cotan theta and z5 values for this track are saved in the match word. After stereo track reconstruction, attempts are made to point the stereo track at the different detector components. If the pointing is successful the bit related to that detector component is set (CMX,CMU,CMP,CEM and etagap). For all but the CEM detector these are the West and East detectors. For the CEM detector, the bit is set for tower the track is pointing to (total of 10 trigger CEM towers).

The Track End of Event word is defined as:

Match Bit	Definition
31:23	Unused
22	End of Event bit
21	Unused
20:19	L2 Buffer number
18:9	Error flags
8	Reserved
7:0	Event Tag

The muon/electron matching data word are produced by the L2 Processor itself based on the presence of any hit in the outer detectors encountered by a track while it escapes from the interaction point.

Here is a brief explanation of the stereo track parameters and how they are encoded from floating point numbers to integers and back. The direct results of the stereo reconstruction algorithm are two floating point numbers: the z position at $r = 0$ (along the beamline) and $\cot(\theta)$. For storage in the TL2D bank, we first calculate the z coordinate at SL5 using the expression $z_{\text{SL5}} = z_0 + 94.0 \cot(\theta)$, where 94.0 cm is the radius of SL5. The z resolution at SL5 is estimated to be 1 cm, so the minimum granularity for encoding as an integer is taken to be 0.5 cm. The maximum value of z at SL5 would be 122 cm, so nine bits are used to store the magnitude as an integer. The magnitude of the z position is encoded into an integer using, schematically, $\text{int}(\text{abs}(z_{\text{SL5}})/0.5 + 0.5)$. The $\cot(\theta)$ resolution is estimated to be 0.1. The maximum value of $\cot(\theta)$ would be about 1.85, given the COT geometry. An encoding granularity of 0.005 is chosen (0.05 would be sufficient, but we have extra bits and prefer to match the number of bits used for z position encoding). The magnitude of $\cot(\theta)$ is encoded into an integer using logic of the form $\text{int}(\text{abs}(\cot(\theta))/0.005 + 0.5)$. To extract these parameters from the TL2D bank and, in the case of the z position at SL5 and $\cot(\theta)$, decode the integers back to floating point values, accessors have been written for the TL2D_StorableBank class of the TriggerObjects package. An accessor is also available to recalculate the z position at $r = 0$ from the z position at SL5.

9 Stereo Tracking Algorithm Implementation

The purpose of the stereo tracking algorithm is to quickly desparsify the L2 pixel data in the T2FD bank format and, based on these pixels, find the best stereo quantities per axial track. The algorithm is designed to quickly return the stereo quantities $\cot(\theta)$ and Z_5 for any given axial track, so these quantities can then be used in the trigger decision. Because the algorithm runs on the L2 PC and could potentially be run several times per event, it is crucial it be fast enough to execute without introducing additional dead-time into the system.

The algorithm can be broken down into three major parts, each of which will be described in greater detail below. For each axial track, the first task is to desparsify the pixel data contained in the TP2D banks of the PULSAR. We only unpack the pixel data for every relivent stereo superlayer (3,5,7) in a given window around the extrapolated axial track. The next step is then to mask and cluster these pixels.

This masking eliminates fake pixels by removing those pixels whose slope does not match expected similar slopes from the curvature of the axial track. We then cluster consecutive pixels, decreasing possible track multiplicities in the next step. The third and final part of the algorithm is to loop over combinations of clusters in the three layers, apply fiducial cuts, and choose the best quality stereo tracks.

9.1 Initialization

The initialization of the algorithm corresponds to a per-event overhead, so the code has been minimized in this section. Besides clearing memory, the only function that takes a significant amount of time is finding each of the finder cards' length in the T2FD bank for later use. This corresponds to a $\sim 1\mu\text{s}$ per event overhead. Potentially, this could be decreased in the future by finding the card length in hardware and then appending the bank format to include the length of the individual cards.

9.2 Unpacking

We unpack from the bank only those stereo pixels in a window around the given axial extrapolated phi-slice. The difference between axial extrapolated phi and the stereo hit phi is directly proportional to Z position of the track at that layer, so fiduciality requirements dictate the size of the window and vary at each of the three layers. Specifically, the track must lie within $|Z_0| \leq 90\text{ cm}$ and must be contained in the tracker $|Z_8| \leq 155\text{ cm}$. To save computation time, lookup tables were developed to map axial XFT track quantities (ϕ_6 and p_t bins) into appropriate stereo layer pixel bank locations.

Because of the format of the T2FD bank, the whole of a Finder card's mask section must be read before pixel data is extracted. This mask section parsing costs some time, so the result is stored for use in future tracks. Lookup tables were also developed to quickly determine the absolute stereo pixel value of a corresponding mask section bit. Using this machinery, we can determine the exact bank location of the needed pixels. It should also be noted that a window may possibly span two finder cards and therefore two sections in T2FD; this will increase the time needed, as the next card's mask section will need to be parsed. Also, as a time saver, the clustering and masking is done as the pixels are being unpacked.

9.3 Clustering and Masking

Clustering and masking both serve as quality cuts as well as techniques to reduce time spent looping over possible track configurations (part 3 of the algorithm).

Masking makes use of the five slope bins per pixel provided by the finder. We cut fake pixels by assuming real pixels have similar slopes to their associated axial track. Remember, only SL5 and SL7 have full pixel information at this point; SL3 pixels are just the middle of the 'mask' section without any slope information. Based on studies, we divide the extrapolated stereo pixel window up into three sections in ϕ whose selected hits would correspond to position in z , (again $z_7 \propto \Delta SL7$ where $\Delta SL7$ is the axial track extrapolation position at SL7 minus stereo hit position at SL7). We allow stereo pixel slopes to be similar or slightly different than the axial slope based on the following table:

	High ϕ Third	Mid ϕ Third	Low ϕ Third
SL7	(0,1)	(-1, 0, 1)	(-1, 0)
SL5	(-1, 0)	(-1, 0, 1)	(0, 1)
SL3	(0, 1, 2)	(-1, 0, 1)	(-2, -1, 0)

$\Delta\phi$ vs. Δ slope for different superlayers.

So for example, stereo pixels corresponding to the central part of all superlayer windows will be masked with the expected axial slope \pm one slope bin. Similarly, at SL7, pixels on the high ϕ part of the window will be masked with the axial slope bin and one slope bin less than the axial. This masking serves to reject fakes and leaves pixels most likely to be real.

We cluster stereo pixels in order to reduce the number of potential combinations of tracks from which to choose. This is motivated by studies that show real tracks often leave several concurrent pixels hit in a row at this Level 2 resolution. See the table below for a single track Monte Carlo study of real pixel clusters.

	1	2	3
SL3	27.6%	69.2%	3.2%
SL5	18.7%	72.6%	8.7%
SL7	20.2%	71.5%	8.3%

Number of adjacent pixels in ϕ of simulated single track stereo hits.

To simplify and speed up clustering, we cluster by T2FD sub-cell. One group of three pixels corresponding to a sub-cell mask bit will form a cluster whose location will be determined by the average phi location of the valid, hit pixels in that group of three pixels. Therefore, for every mask bit set in the T2FD mask header, there will be one cluster formed (assuming the pixels contain hits in the requisite slope bins). The value of this cluster corresponds to the difference between the axial extrapolated ϕ and the actual stereo pixel hits in a particular superlayer. We call this value, in superlayer seven for example, $\Delta SL7$.

9.4 Combinatorics

At this point, we have candidate pixel clusters for each of the superlayers 3, 5, 7. Any combination of these clusters, one from each layer, corresponds to a stereo track with a distinct Z_5 and $\cot(\theta)$ value. We do a χ^2 likelihood minimization to derive a linear equation for these quantities:

$$Z_0 = -4.606 \Delta SL7 + 1.032 \Delta SL5 + 6.444 \Delta SL3 \quad (2)$$

$$\cot(\theta) = 0.0589 \Delta SL7 + 0.0008 \Delta SL5 - 0.0581 \Delta SL3 \quad (3)$$

This allows us to quickly calculate these quantities. We also define the quantity $\Delta_{3,5,7} = \Delta SL3 + \Delta SL7 - 2 * \Delta SL5$ which should ideally be zero for a track that does not curve in the $r - z$ plane. This is a straight line quality value.

With several clusters at each superlayer, the combinatorics of this loop over clusters can be very substantial and therefore take substantial time. As the loop progresses, combinations from layer seven and three that lie far outside a reasonable Z value (95 cm), are skipped. From the complete loop over clusters we keep all tracks with $\Delta_{3,5,7} \leq 25$ and $-90 \text{ cm} < z_0 < 90 \text{ cm}$. The $\cot(\theta)$ fiducial cuts are enforced by appropriate choice of stereo pixel window size at each layer, limiting Z and therefore angle in θ . These are the candidate tracks to be used in the trigger logic.

10 Algorithm Performance and Stereo Track Triggers

The resolution and efficiency of the reconstruction algorithm will dictate the ultimate performance of the upgraded triggers. With this in mind, we have performance results for the baseline algorithm and maintain the flexibility to use either the single best quality track combination, or all of the highest quality tracks found for each axial track. We also will leave open the possibility of improving the reconstruction algorithm at a later date.

Two main inclusive lepton triggers will initially make use of the L2 XFT Stereo Upgrade. The inclusive CMX trigger stands to benefit most from stereo information due to the ability to point a track in either the East or West rings of the CMX detector. The CMUP trigger will also benefit from pointing to either the East or West side of the CMU.

10.1 Reconstruction Algorithm Performance

The performance of the trigger will be determined by the resolution of the stereo quantities, the efficiency of finding tracks, and the time it takes to do so. Resolution studies were performed by matching axial XFT tracks to offline tracks for tight CMX muons. SLAM confirmed tracks were required to match within six Level 2 pixels at each of the three used stereo superlayers. We then compare the best Z_0 and $\cot(\theta)$ to the matched offline values. The results are shown in figures 6 and 7. We can then use these numbers to estimate the pointing resolution at the outer lepton detectors. It is important to note that these plots are using only the best track for every axial track. Triggers designed with the 'valid track' method will use several tracks to confirm pointing to other detectors.

The efficiency of reconstructing tracks can suffer in the pixel masking, quality and fiducial cuts. When exploring the algorithms efficiencies we found an issue with simulated pixels from XFTSim. Because the L2 pixels are not read out on every event, and because L2 pixels are not completely determined by the COTD bank, XFTSim cannot make a one-to-one simulation of the exact L2 pixels. This results in an artificially low reconstruction efficiency when using simulated pixels. Instead, pixels from data must be used to study efficiency.

Reconstruction efficiency is studied for four physics objects in data. First, all axial tracks are attempted and the efficiency is found. This sample is expected to contain several fake tracks. Assuming good efficiency for offline tracks, this axial track inefficiency is really a fake rejection above and beyond Level 1 SLAM confirmation. This is further supported by the fact that a great majority of axial track reconstruction inefficiency is caused by slope masking, which was implemented to cut fake tracks. Tight muons, as defined by Joint Physics, are also studied. Their corresponding XFT track is selected by matching to the




Figure 6: Z_0 resolution for the best track, $\sigma = 11 \text{ cm}$.




Figure 7: $\cot(\theta)$ resolution for the best track, $\sigma = .12$.

muon's offline track. Finally Z events are studied by selecting two tight muons. The inefficiencies are listed table 16. For high quality physics objects, the efficiency for reconstruction is over 99.8 %, and it looks like about 3 % of SLAM confirmed XFT tracks are being eliminated as fakes.

The final figure of interest is the timing of the reconstruction. We put a lot of effort into making sure reconstruction is as fast as possible so as not to introduce latency in the L2 trigger system. We time the reconstruction with on-line tests in the backup L2 PC. Based on these tests we can tell how much time it will take, per track, to execute the reconstruction in the `reconstructTrack()` method. This time is a per track reference for trigger construction and does not include several other overheads, such as getting the data into the L2 PC. We see that the reconstruction code alone should take from 1.3 - 2.0 μS per track depending on luminosity, see figure 9. Triggers constructed with this reconstruction code should note that this does not include the encapsulating trigger code and take this into account accordingly.

10.2 ‘Best Track’ versus ‘Valid Track’

Because of drops in efficiency versus Luminosity, the original method of finding a single ‘Best Track’ choice of stereo quantities for every axial track was found to be insufficient. We now return from the func-

tion every possible combination of stereo pixels that fulfill the fiducial requirement and has $\Delta_{3,5,7} \leq 25$. This is shown to flatten out the efficiency curve at the expense of rate rejection. However, we will maintain the ability to use both methods in the future.

10.3 CMX

The existing CMX trigger at L2 uses all axial $r - \phi$ projected tracks to match to muon stubs in the detector. This means that non-CMX fiducial tracks may fake legitimate muons based on a coincidental ϕ correspondence with a muon stub. The ability to point a track to a specific region in rapidity allows a third dimension of fiducial fake rejection. The upgraded CMX trigger **CMX6.PT15.3DMATCH** uses the ‘Valid Track’ approach to extrapolate candidate stereo tracks to the radius of both the inner and outer edge of the CMX detector. The trigger then simply checks that the projected Z position of the candidate tracks at the inner CMX radius is closer to the origin than an inner radius cut. It then checks that it is farther from the origin than an outer radius cut, see figure 10. These cuts are chosen based on offline muon samples to be efficient at 99.6% for successfully reconstructed CMX axial tracks, actual efficiency may differ. The resolution of the tracks is not good enough to distinguish the slightly varied geometry of the keystone and miniskirt, so no special adjustments are made for these areas.

The efficiency of CMX reconstruction has been studied relative to Luminosity, see 11. Due to the fact that data pixels must be used for these studies, the statistics are very low. We have started taking data online with the appropriate pixel banks being read out, but it will be some time before enough statistics are gathered to make conclusive efficiency studies.

The rate reduction projections are calculated with the HTTF rate tools method and agree nicely with results from data, see figure 13.

The efficiency in simulation of L2 CMX initially was found to decrease with luminosity with the ‘best track’ approach. The new ‘valid track’ trigger is found to be flat relative to Luminosity at the expense of some rate reduction. Efficiency can be measured relative to good offline muons. When compared to offline, we look at tight muons and the associated offline track. These tracks are then matched to a XFT axial track. The offline efficiency is then the number of those reconstructed tracks that pass the trigger, over the total number of well matched, reconstructed SLAM confirmed tracks. Note this does not take into account the reconstruction inefficiency, see figure 8. The results of this approach are shown in figure 14. Note that due to the need for high statistics, these studies are done with simulated pixels. They will be checked in the future as statistics are gathered.

10.4 CMUP

The **CMUP6.PT15.3DMATCH** trigger has essentially the same functionality as the CMX trigger. Each axial track is still matched in ϕ as in the original trigger, but with the addition of stereo pointing, a specific side of the CMU detector can be used eliminating fakes from coincidental muon stub hits on the opposite side of the detector. Stereo tracking can return if the track hit the East CMU, West CMU, non-fiducial or ‘both’ if the track is indistinguishable between both sides. As in the CMX trigger method, we place cuts designed to be 99.5% efficient offline on the extrapolated Z_{CMU} . The results are shown below in figure 17. As in the CMX case, these studies are done with simulated pixels and will be checked in the future

with data L2 pixels.

References

- [1] J. Dittman *et al.*, *Run 2b XFT Stereo Mask Finding Specifications*', CDF Note 7789
- [2] R. Erbacher *et al.*, *Specification of the Data Format from Level 1 to Level 2 of Stereo XFT*', CDF Note 7772
- [3] W. Badgett *et al.*, *The CDF Run II Event Structure*, CDF Note 4152
- [4] H. Gerberich *et al.*, *Specification for Level 2 TrackList Boards*, CDF Note 8871

“TP2D”
Bank Number
Bank Version
Bank Length (variable)
Bank Type (I*4)
Number of Blocks (=2)
Pointer to Block 0
Pointer to Block 1
Pointer to End of Data
Block 0: Number of Cards ($= i \times 6$)
Pointer to Pulsar Board 1 (CONTROL1) Card
Pointer to Pulsar Board 1 (CONTROL2) Card
Pointer to Pulsar Board 1 (DATAIO1_1) Card
Pointer to Pulsar Board 1 (DATAIO1_2) Card
Pointer to Pulsar Board 1 (DATAIO2_1) Card
Pointer to Pulsar Board 1 (DATAIO2_2) Card
...
Pointer to Pulsar Board i (CONTROL1) Card
Pointer to Pulsar Board i (CONTROL2) Card
Pointer to Pulsar Board i (DATAIO1_1) Card
Pointer to Pulsar Board i (DATAIO1_2) Card
Pointer to Pulsar Board i (DATAIO2_1) Card
Pointer to Pulsar Board i (DATAIO2_2) Card
Block 1: Number of Cards ($= j \times 6$)
Pointer to Pulsar Board $i+1$ (CONTROL1) Card
Pointer to Pulsar Board $i+1$ (CONTROL2) Card
Pointer to Pulsar Board $i+1$ (DATAIO1_1) Card
Pointer to Pulsar Board $i+1$ (DATAIO1_2) Card
Pointer to Pulsar Board $i+1$ (DATAIO2_1) Card
Pointer to Pulsar Board $i+1$ (DATAIO2_2) Card
...
Pointer to Pulsar Board $i+j$ (CONTROL1) Card
Pointer to Pulsar Board $i+j$ (CONTROL2) Card
Pointer to Pulsar Board $i+j$ (DATAIO1_1) Card
Pointer to Pulsar Board $i+j$ (DATAIO1_2) Card
Pointer to Pulsar Board $i+j$ (DATAIO2_1) Card
Pointer to Pulsar Board $i+j$ (DATAIO2_2) Card

Table 12: Structure of the TP2D bank. This example shows the bank format when there are two crates reading out Pulsar boards. Block 0 contains the data from the i number of Pulsar boards to be readout in the first crate. Block 1 contains the data from the j number of Pulsar boards to be readout in the second crate. For each Pulsar board in the readout chain, 6 card pointers are created, which corresponds to the 6 onboard DAQ buffers.

Word	Data Description
0	Header Word
1	SLINK BOF
2	SLINK Header 1
3	SLINK Header 2
4	Header from Finder 0
5	Mask word 0 from Finder 0
...	...
i	Trailer from Finder 0
i+1	Header from Finder 1
...	...
n	Trailer from Finder 11
n+1	SLINK Trailer: 0x10000
6	SLINK EOF

Table 13: Structure of the Pulsar XFT boards. For events that are not L1 aborted nor truncated.

Word	Data Description
0	Header Word
1	SLINK BOF
2	SLINK Header 1
3	SLINK Header 2
4	data word: 0xc000c000
5	SLINK Trailer: 0x10001
6	SLINK EOF

Table 14: Structure of the Pulsar XFT boards for events aborted based on L1 mask.

Word	Data Description
0	Header Word
1	SLINK BOF
2	SLINK Header 1
3	SLINK Header 2
4	data word
5	data word
...	...
N	data word
5	SLINK Trailer: Bit # 2 set to 1
6	SLINK EOF

Table 15: Structure of the Pulsar XFT boards for truncated events. N represents the word at which the total size equals the truncation threshold.

XFT Tracks	3.1 %
Tight Muons	.10 %
CDF Offline tracks	.18 %
Z Events	0 % \pm 2.2 %

Table 16: L2 track reconstruction inefficiency of different objects.




Figure 8: Reconstruction efficiency for SLAM confirmed Axial tracks.




Figure 9: Time per XFT Axial Track Stereo Reconstruction, in units of $.1 \mu\text{S}$.




Figure 10: CMX stereo track reconstruction.

CMX Trigger Efficiency vs Luminosity, $\Delta 357 < 25$, reduced SL3




Figure 11: L2 Efficiency of the CMX trigger versus Luminosity, not including reconstruction inefficiency $\approx .1\%$




Figure 12: L2_CMX_3DMATCH rates for different possible versions of the trigger. Existing CMX and CMX_3D are shown for consistency. The ‘Best track’ method is shown in black, as well as the all or ‘Valid Track’ methods for different cuts on $\Delta_{3,5,7} \leq 15, 25$.




Figure 13: The rates in data agree with simulation. Shown here is the final CMX_3DMATCH trigger corresponding to CMX_VALIDTRACK_25 in 12, prescaled by 10.




Figure 14: L2_CMX_3DMATCH efficiency when compared to good, well matched offline muon tracks.




Figure 15: CMUP stereo track reconstruction.




Figure 16: L2_CMUP_3DMATCH ‘valid track’ rates from data. They match projections from the HTTF rate tools.




Figure 17: Offline method of finding efficiency for good matched CMUP offline muons.