

# Chapter 12

## Trigger

### 12.1 Overview

The trigger plays an important role in hadron collider experiments because the collision rate is much higher than the rate at which data can be stored on tape. In run II the collision rate will be effectively equal to the crossing rate of 7.6 MHz while the tape writing speed will be less than 50 Hz. The role of the trigger is to efficiently extract the most interesting physics events from the large number of minimum bias events. For example, the total  $t\bar{t}$  cross section is approximately nine orders of magnitude smaller than the minimum bias cross section.

Due to changes in the detector and the accelerator the entire trigger system used in run 1 must be replaced for run II. The primary reason for replacing the trigger electronics along with all CDF front-end electronics, is the reduction in the accelerator bunch spacing from  $3.5\mu\text{sec}$  to  $132 - 396\text{nsec}$ . In the past, trigger signals from the calorimeters were sent to the control room, where they were processed, with the trigger decision sent back to the detector before the next beam crossing. As a result, the data from only one crossing needed to be stored on the detector. In run II there will not be enough time to send detector signals to the control room between bunch crossings, let alone make a trigger decision and distribute it back to the detector. In addition most of the old trigger is incompatible with new or upgraded detector elements.

In run 1b, the trigger had to reduce the raw collision rate by a factor of  $10^5$  to reach the tape writing speed of  $< 10$  Hz. In run II, the conditions will be more challenging as the luminosity will increase by an order of magnitude, but the rate of data written to tape will only increase by a factor of 3 to 5. This will require the trigger to have a larger rejection factor while maintaining high efficiency for the broad range

of physics topics we study.

The CDF trigger system has a three level architecture with each level providing a rate reduction sufficient to allow for processing in the next level with minimal deadtime. Level-1 uses custom designed hardware to find physics objects based on a subset of the detector information and makes a decision based on simple counting of these objects (e.g. one 12 GeV electron or two 1.5 GeV muons). The Level-2 trigger uses custom hardware to do a limited event reconstruction which can be processed in programmable processors. The Level-3 trigger uses the full detector resolutions to fully reconstruct events in a processor farm.

Figure 11.1 is a functional block diagram of the three level pipelined and buffered trigger system. To allow time for transmission and processing of the trigger signals to make the trigger decision, we have selected a  $5.5\mu\text{sec}$  Level-1 latency. This requires each detector element to have local data buffering for the 42 beam crossings (at 132 nsec separation) that occur during the latency period.

If an event is accepted by the Level-1 trigger, the front-end electronics move the data to one of four on-board Level-2 buffers. This is sufficient to average out the rate fluctuations and allow a 40 kHz Level-1 accept rate with  $\leq 10\%$  deadtime for the anticipated  $20\mu\text{sec}$  Level-2 processing time. The data acquisition system will allow the Level-2 trigger to accept as many as 300 events per second. These are transferred to the Level-3 trigger processor farm where the events are reconstructed and filtered using the complete event data, with  $\leq 50$  Hz written to permanent storage. The custom Level-1 and Level-2 trigger hardware is described in this chapter while the Level-3 trigger is described in Chapter 11.

Both the Level-1 and Level-2 trigger systems will be provided with detector information not available

in the respective run 1 systems. These enhanced capabilities are required to provide the rejection power needed for run II while significantly expanding the triggers physics potential. The most significant change for Level 1 is the addition of track finding. Previously available only at Level 2, tracks in the outer tracking chamber will be reconstructed within  $2.7 \mu\text{sec}$  after a  $\bar{p}p$  collision. This allows a track to be matched to an electromagnetic-calorimeter cluster for improved electron identification, a track to be matched to a stub in the muon system for better muon identification and momentum resolution, and tracks to be used alone for triggers such as  $B^0 \rightarrow \pi^+\pi^-$ .

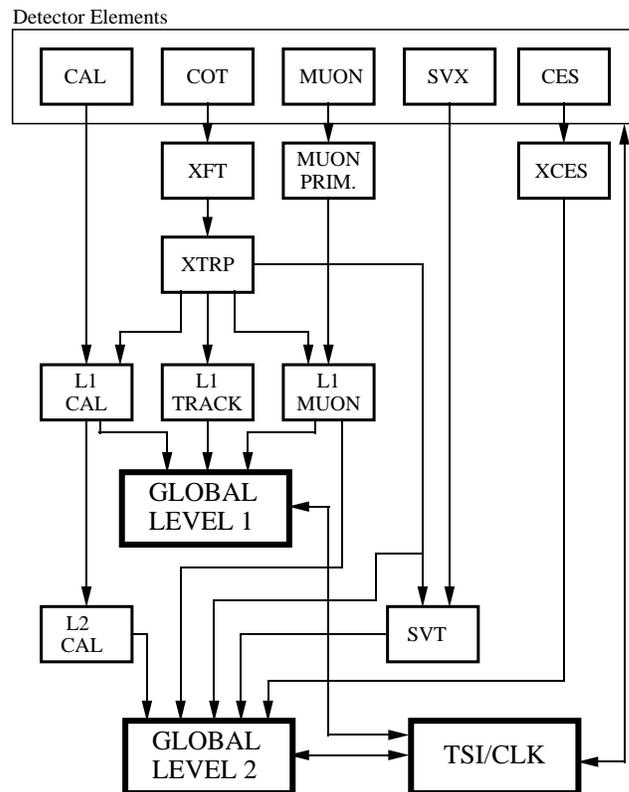
The most significant addition to the Level-2 trigger is the Silicon Vertex Tracker (SVT) which will more fully exploit the physics potential of the high precision silicon vertex detector. A major effort has been made to build this system that, for the first time in a hadron-collider experiment, can trigger on secondary vertices. This will make accessible a large number of important processes involving hadronic decay of  $b$ -quarks outlined in section 2.6. In addition, the SVT provides improved momentum resolution for tracks and finer angular matching between muon stubs and central tracks.

The block diagram for the run-II trigger system is presented in Fig. 12.1. The input to the Level-1 hardware comes from the calorimeters, tracking chamber, and muon detectors. The decision to retain an event for further processing is based on the number and energies of electron, muon, and jet candidates, as well as the  $\cancel{E}_T$  in the event. A Level-1 accept can also be generated based on the kinematic properties of observed track pairs.

Events accepted by the Level-1 system are processed by the Level-2 hardware. All of the information used in the Level-1 decision is available to the Level-2 system, but with higher precision. In addition, data from the central calorimeter shower-max detector allows improved identification of electrons and photons. Jet reconstruction is provided by the Level-2 cluster finder; secondary-vertex information is produced by the SVT. A Level-2 accept initiates full detector readout for the event. An extension of the Level-2 system to include tracking in the  $1 < |\eta| < 2$  region using the Intermediate Silicon Layers is under consideration.

The Trigger System Interface (TSI) and Clock systems which synchronize the trigger and DAQ sys-

## RUN II TRIGGER SYSTEM



PJW 9/23/96

Figure 12.1: The run-II trigger-system block diagram.

tems are described in sections 11.6.3 and 11.6.4. All subsystems of the Level-1 and Level-2 triggers use the VME hardware and protocols described in section 11.6.

### 12.2 Level 1 Trigger hardware

The Level-1 hardware consists of three parallel synchronous processing streams which feed inputs of the single Global Level-1 decision unit. One stream finds calorimeter based objects (L1CAL), another finds muons (MUOM PRIM-L1MUON) while the third finds tracks in the central tracking chamber (XFT-XTRP-L1 TRACK). Since the muon and electron triggers require the presence of a track pointing at the corresponding outer detector element, the tracks must be sent to the calorimeter and muon streams as well as the track only stream. Up to 64 different

triggers can be formed using simple ANDs and ORs of objects from these streams.

All elements of the Level-1 trigger are synchronized to the same 132nsec clock with a decision made every 132nsec by Global Level-1. During abort gaps events are automatically rejected. When the accelerator is operating in 36 bunch mode (396nsec), the trigger will be clocked every 132nsec with the two intermediate clock cycles rejected automatically.

### 12.2.1 Level 1 calorimeter hardware (L1CAL)

The goal of the Level-1 (L1) calorimeter trigger is to trigger on electrons, photons, jets, total event transverse energy ( $\Sigma E_T$ ) and missing transverse energy ( $\cancel{E}_T$ ). The calorimeter triggers are divided into two types: object triggers (electrons, photons and jets) and global triggers ( $\Sigma E_T$  and  $\cancel{E}_T$ ). The object triggers are formed by applying thresholds to individual calorimeter trigger towers, while thresholds for the global triggers are applied after summing energies from all towers. In turn, the object triggers are of two types: single object triggers (e.g. electrons from  $W \rightarrow e\nu$ ) where finding one in the event is sufficient to generate a L1 accept, and di-object triggers where the rate is too high to trigger on only one (e.g. electrons from  $J/\psi \rightarrow ee$ ).

Through each board of the trigger, there are two parallel paths: one for the object triggers and one for the global triggers. In the object trigger path, thresholds are applied to the individual tower energies, then the number of towers passing each threshold is counted. For the single-object triggers, this is a 1-bit sum (0,  $\geq 1$  towers) and for the di-object triggers it is a 2-bit sum (0, 1, 2,  $\geq 3$  towers). Electron and photon triggers are formed by applying energy thresholds to the electromagnetic (EM) energy in a tower, while jet triggers are formed using the total (EM + HAD) energy in a tower. To reduce the low- $E_T$  electron trigger rate, tracks from the Level-1 track processor (XFT) are available for matching to the tower, and towers with significant HAD energy can be rejected. The global path sums the total energy (HAD+EM) in all the towers (with appropriate weighting for the  $\cancel{E}_T$  calculation) and applies thresholds to these sums.

A total of 16 object-trigger bits are available for use by all single and di-object triggers. The default system configuration parses the 16 bits into 10 single-

object triggers and 3 di-object triggers. There is a limited capability to reconfigure the bit usage to allow for more or fewer di-object triggers. There is also flexibility in how many can be electron/photon and how many jet, with a limitation of no more than 8 of either kind. A total of 4 thresholds are available for the global triggers: 2 for  $\Sigma E_T$  triggers and 2 for  $\cancel{E}_T$  triggers. The result is a 20-bit word presented to the final Level-1 decision card to summarize the calorimeter triggers.

#### 12.2.1.1 Data Flow

The data flow in the calorimeter trigger is shown in Fig. 12.2. The digitized calorimeter data are summed into trigger-tower energies and weighted by  $\sin\theta$  on the front-end (ADMEM) cards to produce  $E_T$ . Studies of Level-1 trigger rates for Run II[1, 2] indicate that the tower segmentation of the original CDF trigger provides acceptable rates. The existing segmentation[3] of the central and wall calorimeters, with towers of approximately  $\Delta\eta = 0.2 \times \Delta\phi = 15^\circ$ , is carried into the new plug calorimeter, resulting in a  $24 \times 24$  map in  $\eta - \phi$  space. The tower at highest  $\eta$  on either side of the detector covers  $2.6 < |\eta| < 3.6$ . The same segmentation is used for the hadronic and electromagnetic calorimeters.

The data are transmitted upstairs to the trigger as 10-bit parallel words, with a least count of 125 MeV and a full scale  $E_T$  of 128 GeV[4], on SCSI-type cable consisting of 10 twisted pairs.

The data are processed in 6 L1CAL crates with final global sums and trigger summaries made in the Global L1 crate. The data are forwarded from the L1CAL system to the Level-2 cluster finder (L2CAL) for use after a L1 accept.

A more detailed view of the L1CAL pipelines is presented in Fig. 12.3. The whole L1 system is pipelined with each stage taking less than 132 ns and processing a new event every 132 ns. The bulk of the L1 processing takes place on 16 Digital Information Receive And Compare(DIRAC) boards in each of 6 L1CAL VME crates. The DIRAC boards are arranged in sets of four, handling the data from a full-rapidity,  $\Delta\phi = 15^\circ$  wedge of the calorimeter: West Plug, West Central, East Plug and East Central. Each DIRAC card has 12 channels covering 6 EM towers and the 6 corresponding HAD towers. The central cards each cover 5 towers in the central and the first tower in the plug, reaching  $\eta = \pm 1.3$ . Each



track isolation, and 1 bit to flag  $|\eta| > 1$ . Track  $P_T$  thresholds are applied, and 3 bits are fanned-out to the 6 trigger towers.

The tower threshold cuts and track  $P_T$  cuts are performed in two 32k-byte SRAMs per tower, one for jet triggers and one for electromagnetic triggers. The inputs to the jet-trigger SRAM are 8 bits of Total energy and 3 bits of track information. The SRAM for electrons and photons receives 8 bits of EM energy and 3 bits of track data and the 3 least significant bits of HAD energy for HAD/EM. Each of these SRAMs has an 8-bit output corresponding to the 8 possible triggers (combination of threshold(s) and track condition). For example, a low- $E_T$  electron trigger might consist of a tower over 3 GeV of energy, a track with  $P_T$  greater than 3.0 GeV/c, and no more than 250 MeV of HAD energy. Note that the cuts for a particular trigger could be different for different regions of  $\eta$  since each tower has its own SRAM.

The next DIRAC stage summarizes how many towers passed each of the thresholds on the board. The single-object triggers are ORed together and the di-object triggers are summed. The 16 bit summary is transmitted to the CRATESUM card using 16 dedicated lines on the J3 backplane. The allocation of these bits among EM and Jet is programmable, as is the number of di-object triggers.

In parallel with the tower trigger logic, there is a series of adders which calculates the total  $\Sigma E_T$  for the board. The input to these adders is the 8-bit sum of HAD+EM energy for each of the 6 towers. A tower is only included in the sum if its energy is above a programmable threshold, to reduce sensitivity to multiple interactions. The output is an 8-bit word that is transmitted on the J3 backplane to the CRATESUM card.

The tower-trigger and  $\Sigma E_T$  outputs are not time aligned; the data for the  $\Sigma E_T$  path leaves DIRAC several crossings before the triggered tower data. The data is treated this way since the processing of the  $\Sigma E_T$  data requires more pipeline stages in the CRATESUM and PreFRED cards. As shown in Fig. 12.2 mismatches in timing between paths are corrected using FIFOs at the end of processing on the PreFRED modules. In fact all PreFRED modules have FIFOs to ensure that the data presented to FRED is correctly aligned.

### 12.2.1.3 Crate Summary Card

The Level-1 Crate Summary card (CRATESUM) receives the 16 bits of trigger data and 8 bits of  $\Sigma E_T$  data from each of 8 DIRAC cards. The trigger data are summarized into one 16-bit word in the same fashion as on DIRAC. This summary from each CRATESUM is transmitted differentially over twisted pair cable to the TOWTRG PreFRED module in the GLOBAL Level-1 Crate. This PreFRED card combines trigger bits from all 12 CRATESUM cards into the final 16-bit summary of the tower triggers which is put into a FIFO to align it with the bits from other parts of the trigger system. These 16 bits are sent via the J3 backplane to the L1 Decision card (FRED).

In an independent pipeline, CRATESUM forms the  $\Sigma E_T$  for each of the 2 wedges covered by its 8 DIRAC cards. The 10 bit summary for each wedge is transmitted over twisted-pair cable to the SUMET PreFRED module in the GLOBAL Level-1 Crate. The SUMET PreFRED receives the  $\Sigma E_T$  word from all 24 wedges (12 CRATESUMS) and calculates the total  $\Sigma E_T$  and  $\cancel{E}_T$ . For the  $\cancel{E}_T$  calculation, each of the  $\Sigma E_T$  words received from the CRATESUMS is weighted by  $\cos \phi$  and  $\sin \phi$  for the corresponding wedge to form  $\Sigma E_{Tx}$  and  $\Sigma E_{Ty}$  from which  $\cancel{E}_T^2$  is calculated. Thresholds are then applied, two for  $\Sigma E_T$  and two for  $\cancel{E}_T$ . The four  $\Sigma E_T/\cancel{E}_T$  bits are aligned using a FIFO and sent via the J3 backplane to FRED. The calculated  $\Sigma E_T$  and  $\cancel{E}_T$  are stored in L1 FIFOs and are transmitted to the L2 processors after a L1 accept for use in the L2 decision.

The DIRAC, CRATESUM and PreFRED cards are all 9U x 400mm VME cards [5]. These cards are addressable through VME to download thresholds and to read out trigger information [5] into the event data stream. There are also registers that can be loaded or read back through VME for diagnostic purposes.

## 12.3 The eXtremely Fast Tracker (XFT)

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

in the Level-2 trigger decision.

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

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

- The track-finding efficiency will be greater than 96% when the single-hit efficiency of the central tracker is greater than 92%.
- The momentum resolution reported to the trigger will be  $\Delta P_T / P_T^2 < 2\%$ .
- The resolution on  $\phi_0$  will be better than 6 mrad.
- The fake-track rejection will be at least twice as good as the CFT.
- The minimum track  $P_T$  will be 1.5 GeV/c.

The goal of high track efficiency is set by the desire to be efficient for high- $P_T$  physics. Good momentum resolution is a requirement for the high- $P_T$  muon triggers where the primary handle on rate, after matching with muon chamber stubs, is the track momentum in the central tracker. The improved fake-track rejection is also needed to reduce the high- $P_T$  muon trigger rate at high luminosity. The lower  $P_T$  cutoff is set by the desire to maximize the acceptance for  $B$  decays that can be used, for example, to measure CP violation. The resolution on  $\phi_0$  is set by the requirement that XFT tracks will be used as the seed for the SVT at Level 2.

### 12.3.1 Algorithm Description

A block diagram of the XFT is shown in Fig. 12.4. The processor works off of hit data from the 4 axial layers of the COT. There are a total of 16,128 axial wires, and the data on each wire is classified as prompt and/or delayed, for a total of 32,356 bits of information. The definition of prompt or delayed will depend upon the maximum drift in the COT. Assuming that this maximum is  $\sim 100$  nsec, a prompt hit occurs whenever there is a hit in the time window

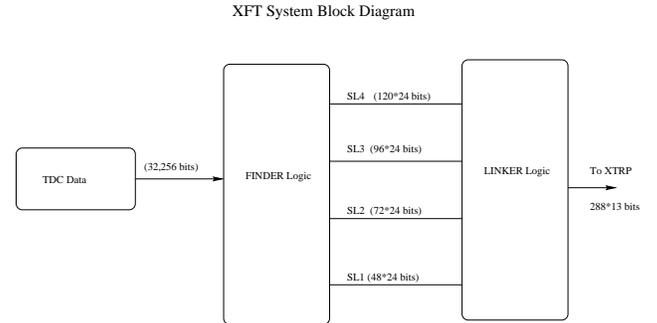


Figure 12.4: Block diagram of the XFT trigger.

0-33 nsec, and a delayed hit is defined as a hit falling in the window 33-100 nsec.

Track identification is accomplished in two processes, the Finder and the Linker. The Finder searches for high- $P_T$  track segments in each of the outer-four axial superlayers of the Central Tracker. Each found segment is characterized by a mean- $\phi$  position in the axial superlayer. The Linker searches for a four-out-of-four match among segments in the 4 layers, consistent with a prompt high- $P_T$  track. An example of a track traversing the COT is shown in Fig. 12.5.

A block diagram of the Finder logic is shown in Fig. 12.6. TDC information is brought from the detector to the XFT on Ansley cables. These cables can carry up to 24 signal pairs. The TDC data is multiplexed 4:1, such that prompt and delayed hit information for 4 neighboring COT cells comes in on a single cable. The Finder is designed to look for valid track segments in each of these 4 cells. A track segment is then defined by the cell it traverses for each of the 12 wire planes in a superlayer, and whether it generated a prompt or delayed hit. A collection of the cell numbers and hit types for the 12 wires in an axial superlayer is called a mask. The mask will change depending on the  $\phi$  of the track, and its angle through the cell (or  $P_T$ ). The finder works by storing

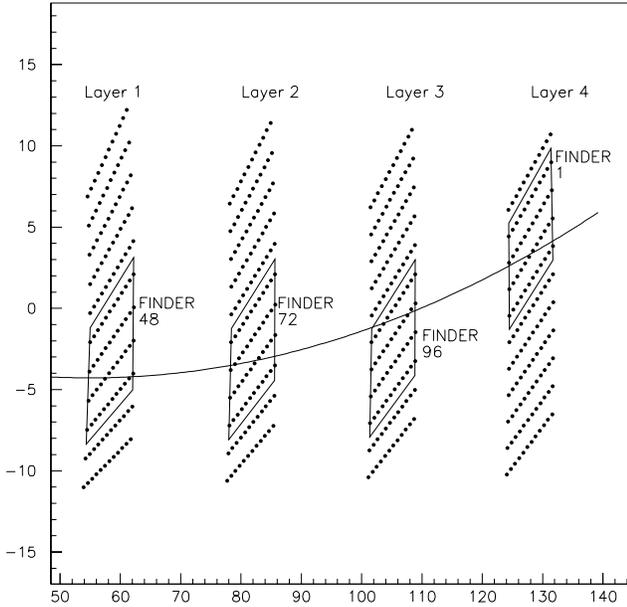


Figure 12.5: A close up view of a track in the COT. All cells in all four axial layers are shown. The relevant Finder in each layer is highlighted.

all possible masks for tracks with  $P_T \geq 1.5$  GeV/c in a database (or equivalently, hard wired on a chip). To keep the total number of masks relatively low, the prompt and delayed bits are “or”ed together on a subset of the wires. The finder compares incoming TDC information with all masks for the given axial superlayer, allowing a programmable (up to 3) number of missed wire planes. Every 132nsec, the finder outputs 24 bits per cell (6 pixels for each of the 4 possible cells) to the linker.

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

A block diagram of the Linker logic is shown in Fig. 12.7. Each Linker is given all of the pixel information from the Finders needed to find the tracks in a  $\Delta\phi = 1.25^\circ$  phi-slice of the tracking chamber. The Linker begins by searching (in parallel) a list of about 1300 roads, where a road is a group of 4 pixels, one from

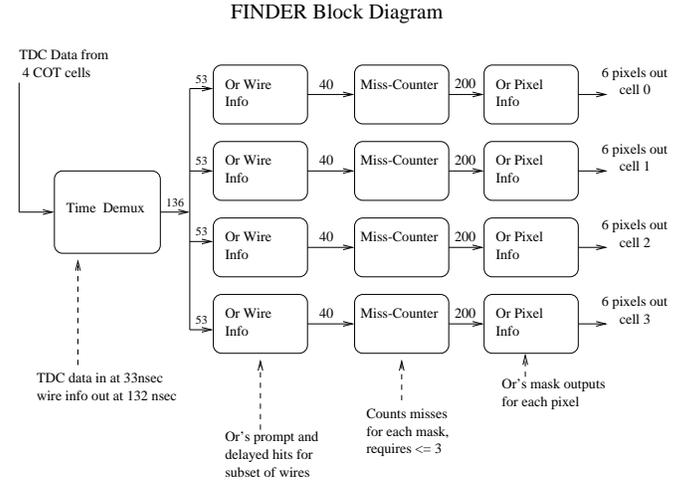


Figure 12.6: Logic flow for the Finder part of the XFT trigger.

each axial superlayer, corresponding to a valid track with  $P_T \geq 1.5$  GeV/c. The roads are defined by their  $P_T$  and the pixel position in layer 3. The roads are then “or”ed down to 128  $P_T$  bins and 8 pixel locations. Found tracks are passed through a priority encoder to find the highest- $P_T$  track in each  $\phi$  bin. Finally, all  $\phi$  bins in a Linker are combined and the highest- $P_T$  track in the  $1.25^\circ$  region covered by the Linker is identified.

### 12.3.2 Algorithm Simulation

A software simulation of the XFT algorithm has been developed using FORTRAN. This simulation models the COT geometry, the expected pulse width ( $\sim 40$ nsec) of hits in the COT, and the occupancy expected under Run II conditions. We assume 132 nsec bunch spacing. The simulation was run on a data sample from Run I. This data sample was collected requiring a muon stub as defined at the trigger level. To simulate the performance under Run II conditions, additional hit data from minimum bias events from Monte Carlo are added, so that the effective luminosity of this sample can be dialed from  $2 \times 10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$  to  $10 \times 10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$  (at 132nsec bunch spacing). The number of Finder masks used for the simulation is approximately 200 per super-

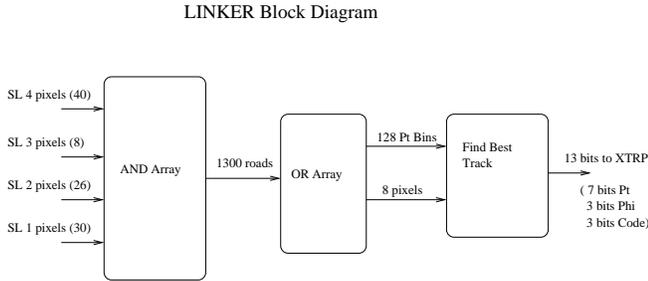


Figure 12.7: Logic flow for the Linker part of the XFT trigger.

layer, and both the Finder masks and the Linker roads are specified using a database.

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

- Track finding efficiency,
- Momentum resolution,
- $\phi_0$  resolution,
- Muon trigger rate vs  $P_T$ ,
- Fake track rate vs  $P_T$ .

These are studied at a luminosity of  $2 \times 10^{32} \text{ cm}^{-2}\text{sec}^{-1}$  at 132nsec bunch spacing. The performance is dependent upon the wire efficiency of the COT. The lower the true chamber efficiency, the more misses the Finder needs to allow to attain high segment finding efficiency. Unfortunately, the performance degrades the more misses are allowed. We study the performance at a wire efficiency of 92% and 99%. For these two cases, the number of misses allowed is 3 and 1 respectively.

We begin by studying the track finding efficiency. First we determine which of the muon triggers are associate with a real track with  $P_T^{offline} \geq 1.5 \text{ GeV}/c$ , as determined by our standard offline algorithm. Of this sample, the percentage found by the XFT is  $(96.5 \pm 0.7)\%$ .

The momentum resolution is defined as  $(P_T^{Offline} - P_T^{XFT}) / (P_T^{Offline})^2$  where  $P_T^{Offline}$  is the momentum using our offline tracking algorithm and  $P_T^{XFT}$  is the momentum returned by the XFT. To calculate the  $\phi_0$  resolution, we use the  $P_T$  and pixel position at the third layer, and extrapolate the track to the origin (assuming a beam constraint). This gives an estimate of the  $\phi_0$  of the track and can be compared with  $\phi_0^{true}$  of the track. In Fig. 12.8 we show the momentum and  $\phi_0$  resolution obtained in the simulation, for the two different wire efficiencies. We observe a momentum resolution of 1.0-1.2%/GeV/c and a  $\phi_0$  resolution of 2.5-3.5mRadians, which are well within our specifications.

We next examine muon trigger rates. A muon trigger at Level 1 (for Run II) is defined as a muon stub match with an XFT track. A muon stub is defined as a pattern of hits in the muon chambers consistent with the passage of a charged particle. Matching is done by extrapolating XFT tracks to the muon chambers and looking for the presence of a muon stub within  $5^\circ$ . This will overestimate the true rate since in run II the matching requirement will be  $2.5^\circ$ . To calculate the muon trigger rates, we assume a muon stub cross section of  $10 \mu\text{b}$  (measured in from run 1 data). We use the simulation to determine what fraction of XFT tracks match with muon triggers. We then increase expected rate by a factor of 1.3 for increased coverage expected for the muon chambers for Run II. In Fig. 12.9 we show the muon trigger rate vs the XFT  $P_T$  threshold imposed. This is the rate expected at a luminosity of  $2 \times 10^{32} \text{ cm}^{-2}\text{sec}^{-1}$ , 132nsec bunch spacing. For a  $P_T$  threshold of 10 GeV/c the rate is approximately 30 Hz. This is clearly well within the Level 1 budget of 50 Khz, but more importantly is also well under the Level 2 budget of 300 Hz. Recall that there is very little to reduce the muon rate at Level 2 besides increasing the threshold.

Finally, each XFT track is classified as either fake or real, by comparing with the list of tracks found by our standard offline algorithm. The XFT track is required to match an offline track to within 1 pixel at 3 of the 4 layers to be considered as a real track, otherwise it is classified as fake. In the bottom plot of Fig. 12.9 we show the fraction of muon triggers which are caused by fake XFT tracks. This is quite low, and there is little evidence of any growth vs  $P_T$  threshold. A comparison of XFT specifications vs a full simulation of the conceptual design is shown in Table 12.1.

	Specification	Simulation
Track finding efficiency	$\geq 96\%$	96%
Momentum resolution	$\leq 2.0\%/GeV/c$	$1.1\%/GeV/c$
$\phi_0$ resolution	$\leq 8\text{mR}$	3.0mR
$\mu$ Trigger rate (@10 GeV/c)	$\leq 50\text{ Hz}$	30 Hz
Fake Fraction (@10 GeV/c)	$\leq 50\%$	$\leq 10\%$

Table 12.1: Comparison of XFT specifications vs results obtained from simulation. The luminosity used in the simulation is  $2 \times 10^{32} \text{ cm}^{-2}\text{sec}^{-1}$  at 132nsec bunch spacing.

### 12.3.3 Algorithm Implementation

As mentioned above, we plan to use the existing cable plant to drive trigger information from the TDC's mounted on the endwall of the detector to the XFT located in the first floor counting room. Since the Ansley cables can carry 24 signal wires, and we plan on 4:1 multiplexing in 132nsec, each Ansley covers 4 adjacent COT cells.

The Finders will also cover 4 adjacent COT cells. This reduces the total chip count for the Finders, but is dependent on how fast we can run the Finder chips, and how many masks can be stored on a single chip. This setup matches the inputs of one Finder to one Ansley cable, However, since the COT cells are tilted with respect to the radial, each Finder will also need neighbor wires, from either side of the set of 4 cells.

We are pursuing implementation of the Finder using field programmable gate arrays (FPGAs). The great advantage of these devices is the ability to re-program them in-situ. The EPF10K50 series from Altera Corp meets our needs in terms of the available logic gates, I/O pins, and speed. We have implemented about 90% of the design shown in Fig. 12.6 using Altera design software. This preliminary work indicates that we can fit about 250 masks, running 4 cells serially in a single Finder, in 132 nsec. This indicates that we will need to “or” some subset of wires in the two outer axial layers to fit them in the 10K50 chips. Another possibility is increasing the total number of Finders used for the outer two layers. This approach is under consideration. Note that the performance of the system using a maximum of 200 masks per layer is excellent. The total number of Finders needed is 336.

We also plan to implement the Linker using Altera FPGAs. As mentioned above, the Linker contains approximately 1300 roads. The automatic place and

route software indicates that this design will fit into a much smaller device than the Finder. The number of Linkers required for the full system is 288. The track information output by the Linker is in the following format:

- 7 bits of  $P_T$  (includes sign),
- 3 bits of  $\phi$ ,
- 1 bit indicating if the track did *not* go through axial layer 4 (i.e. a possible high-eta track),
- 1 bit indicating that the track is isolated,
- 1 bit undefined, reserved for future use.

We plan to place Finder and Linkers on separate boards. The Finders will be arranged on boards in phi slices of  $7.5^\circ$ . Each board would contain 2, 3, 4, and 5 Finders from layers 1, 2, 3, and 4, respectively. We plan on arranging Linker boards such that 6 or 12 Linkers are on a board, so that the information can be brought off of the board in  $7.5^\circ$  increments. The reason for this is that the rest of the detector is segmented into  $15^\circ$  slices. Each board sends its list of tracks to the extrapolation unit (XTRP). For each of these tracks, the above 13 bits of information are sent to XTRP. The XTRP is responsible for mapping the XFT tracks onto muon and electron primitives found by other Level 1 trigger processors, and will be described in the following section.

One problem not highlighted earlier is that very often neighboring Linkers finder the same physical track. If both tracks are kept, this could present a problem for the trigger. On a single board, we can eliminate this problem by only keeping one track. This algorithm will be implemented by passing track information from neighboring LINKERs into another FPGA. However, this procedure is difficult to implement when we reach Linkers on separate boards.

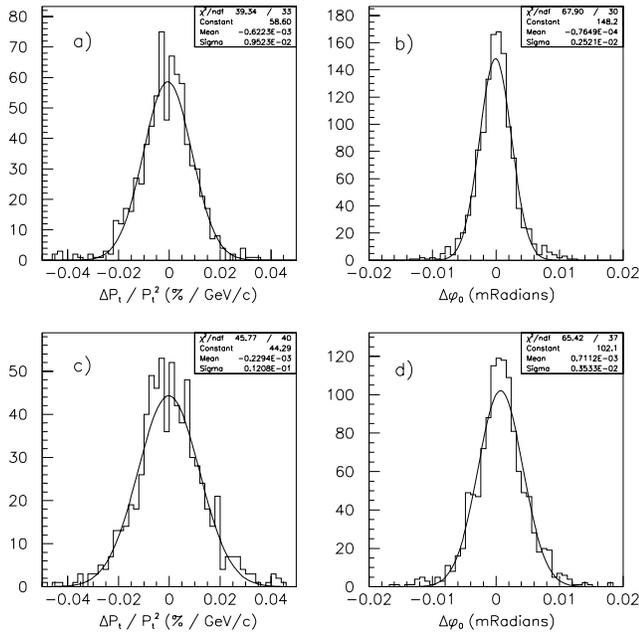


Figure 12.8: Expected performance of the XFT. Figures a) and b) shows the momentum resolution and  $\phi_0$  resolution expected when the chamber wire efficiency is 99%. Figures c) and d) shows the momentum resolution and  $\phi_0$  resolution expected when the chamber wire efficiency is 92%.

These will instead be handled by the XTRP, using the same algorithm.

### 12.3.4 The Extrapolation Unit (XTRP)

The purpose of the XTRP is to receive tracks from the XFT and distribute the tracks or information derived from the tracks to the Level 1 and Level 2 trigger subsystems. After receiving the tracks from the XFT, signals are sent to the Level-1 muon system (L1MUON), the Level 1 Calorimeter trigger (L1CAL), and the Level 1 Track Trigger (L1TRACK) as shown in Fig. 12.1. The tracks are also put into a pipeline and upon receiving a Level 1 accept are stored in Level 2 buffers. The tracks then sent to the Level 2 processor and the silicon vertex trigger (SVT).

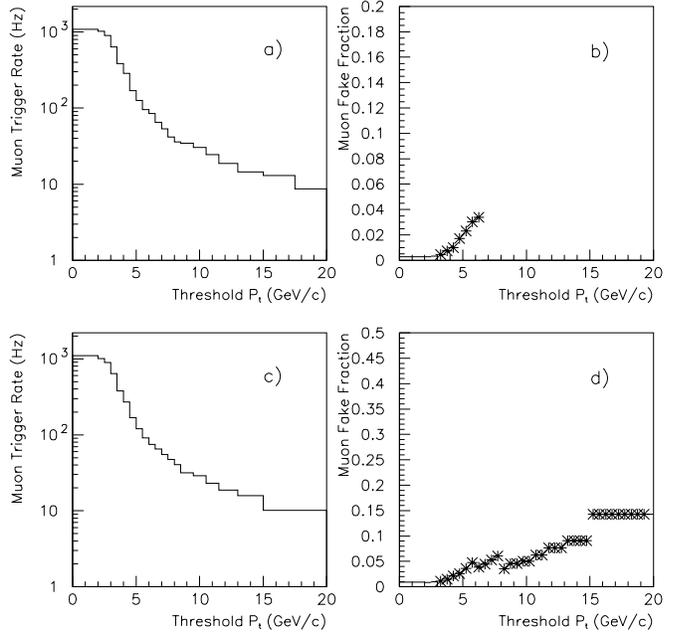


Figure 12.9: Expected muon stub plus XFT track rate. The top plot shows the total rate (in Hz) vs minimum  $P_T$  threshold. The bottom plot is the same, except for the subset of XFT tracks determined to be fakes.

To remind the reader, the XFT logically divides the COT into 288 segments, each covering 1.25 degrees. One track is allowed per segment. Each track has 6  $P_T$  bits, a sign bit, 3 mini-phi bits, an isolation bit, a superlayer 6/8 bit, and one spare bit. Every 132 ns, data from all 288 segments are sent from the XFT to the XTRP even if no track is found. See section 12.3 for more details.

The following information is sent to the Level 1 trigger subsystems:

- **CMU and CMX muon systems (L1 MUON).** Each XFT track is extrapolated to the radii of the CMU and CMX chambers by means of look-up tables. One or more bits, corresponding to 2.5 degree segmentation, are set according to  $P_T$ ,  $\phi$ , and amount of multiple scattering. These bits are sent to the Level 1 Muon Trigger system. Three separate  $P_T$  thresholds are available.

- **Central Calorimetry (L1 CAL).** A set of eight bits for each 15 degree wedge segment is sent to the Central Calorimetry Level 1 trigger. These bits correspond to eight separate momentum thresholds. Track extrapolation is done using look-up tables so that tracks crossing wedge boundaries are handled correctly.
- **Level 1 Track Trigger (L1 TRACK).** The Level 1 Track Trigger is an adjunct to the XTRP. It resides in the same VME crate and provides Level 1 triggers based on XFT track information only. The XTRP modules select tracks above a given  $P_T$  threshold and passes them on a bus to the Track Trigger. The total number of tracks is counted. If 5 or more tracks are found an automatic Level 1 accept is generated. If there are 2, 3, or 4 tracks, the  $P_T$  and  $\phi$  information is used to interrogate look-up tables to generate various Level 1 triggers.

After sending out the Level 1 trigger information all tracks are put into a pipeline and stored pending the Level-1 trigger decision. If a Level 1 accept is received the tracks are latched into Level 2 buffers. All non-trivial tracks are then extracted and put into two separate FIFO's for delivery to the Level 2 processor and to the SVT respectively.

### 12.3.5 Level-1 Muon Trigger (L1 MUON)

The purpose of the Level-1 muon trigger is to provide single and dimuon objects for the Level-1 trigger decision. The run II Level-1 muon trigger includes three major upgrades over the run 1 hardware. First, it is designed to provide deadtimeless Level-1 operation for 132 nsec bunch separation. Second, it uses central-tracking information in the pipelined Level-1 decision to reduce the Level-1 trigger rate. Third, to reduce accidental triggers, it takes advantage of the full detector granularity in matching muon stubs to tracks.

The muon system is composed of wire chambers and scintillators stacked radially from the interaction point. For the purposes of muon identification, the list of detectors traversed by a muon in traveling from the interaction region defines the components of the muon system. The central region is represented in the  $\eta - \phi$  map of Fig. 12.10. In the central-most region,  $|\eta| < 0.6$ , a muon from the interaction region encounters the central tracker, the hadron calorimeter, the

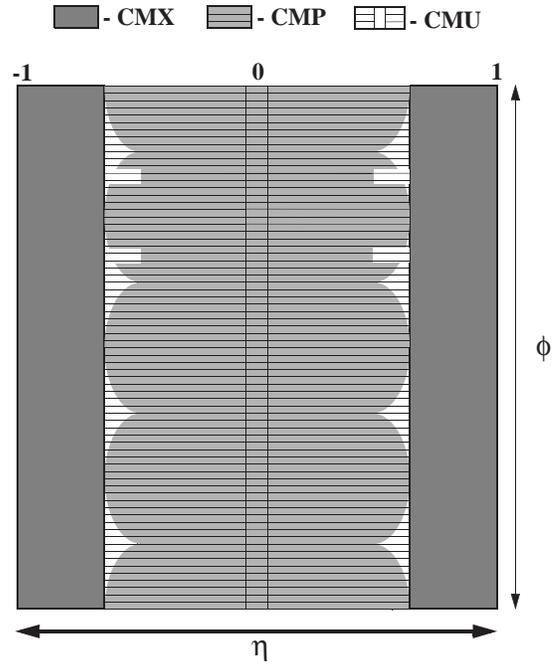


Figure 12.10: CDF  $\eta - \phi$  map for central muons.

projective wires of the CMU, the half-cell-staggered chambers of the CMP, and the CSP scintillators. Beyond the central-most region in  $\eta$ , a muon traverses the central tracker, the hadron calorimeter, and the wire chambers and scintillators of the central muon extension, CMX/CSX.

In contrast to Run I where muon-track matching was performed in  $5^\circ$   $\phi$  segments, the Run-II muon trigger makes use of the finest granularity available from the detector elements for Level-2 processing. For speed and simplicity, this data is combined into fixed  $2.5^\circ$  azimuthal intervals and four  $\eta$  intervals for Level-1 processing. The utility of the fine match has been demonstrated in CDFNOTE-3416. For Level 2, the match is  $1.25^\circ$  in azimuth for CMU and CMX, and for the CMP with its rectangular geometry, the segment is a one-tube stack. The angular interval subtended by a tube stack varies from  $\sim 1.25^\circ$  for the nearest chambers to about half that at the corners of the rectangle.

#### 12.3.5.1 Muon Trigger Primitives

The signal flow begins with the generation of candidate track “primitives” separately in each detector element. These primitives are developed in each 132 nsec crossing and are synchronously pipelined by the

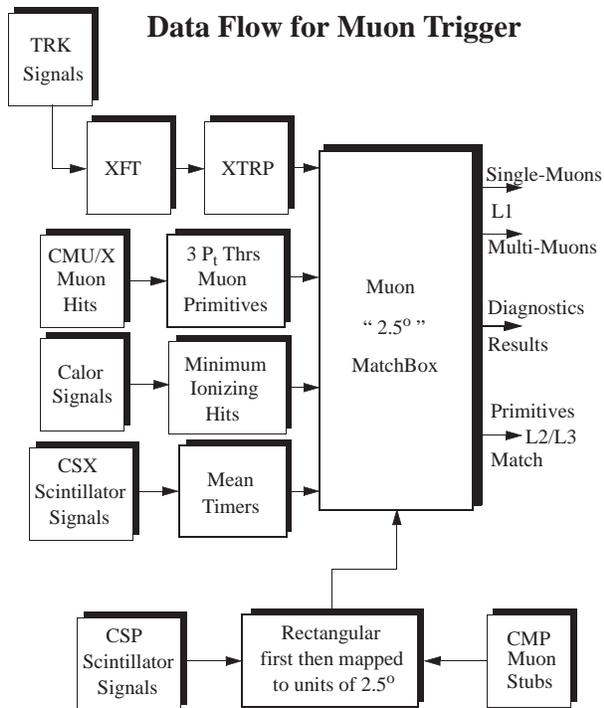


Figure 12.11: Data flow for the muon trigger.

beam-crossing clock. The primitives are derived from single hits or coincidences of hits for the scintillators of the muon system. For the wire chambers the primitives are derived from patterns of hits or from pairs of hits on projective wires with the requirement that the difference in the arrival times of signals be less than a preset threshold. This maximum allowed time difference imposes a minimum  $P_T$  requirement for hits from a single track. The wire chamber primitives are also referred to as muon “stubs”. CMU and CMX primitives represent a pair of hits on projective wires with a time difference less than some downloaded value. The value is set to select tracks above a  $P_T$  threshold. CMP primitives represent the presence of a pattern of hits in a tube stack consistent with the passage of a track of interest. The information flow is shown in Fig. 12.11 which illustrates the primitive sources and the interconnection of units that process the data into trigger signals.

Trigger primitives are generated from the various muon detector elements according to Table 12.2. The name of the trigger card originating the data flow is given in the table.

### 12.3.5.2 Muon and Track Match

The Level-1 trigger match has a  $2.5^\circ$  granularity. Since the geometry of the CMP and CSP is rectangular, data from these detectors are combined in rectangular format, saved for Level 2, and then converted to the nearest  $2.5^\circ$  cylindrical bin for correlation with other detector data at Level 1. The detectors that have  $1.25^\circ$  segmentation are also saved for Level 2 in this finer unit and merged in pairs to  $2.5^\circ$  for matching in the Level-1  $\phi$ -correlation unit. The cards that perform the rectangular match for the CMP/CSP are the Pre-Match cards. The cards that form the  $2.5^\circ$  Level-1 correlation, the Match cards, are designed to handle  $30^\circ$  each, or twelve  $2.5^\circ$  segments. The Match cards form the Level-1 and Level-2 trigger signals. There are many outputs from each 30 degree Match card representing the presence of high, medium, and low  $P_T$  muons. These signals are combined in a Muon summary card and presented to FRED as 8 independent bits. The summary card is programable for flexibility. For example, one might program individual bits for  $\geq 1$  high- $P_T$  muons found,  $\geq 1$  medium- $P_T$  muons found,  $\geq 1$  low- $P_T$  muons seen,  $\geq 2$  medium- $P_T$  muons seen,  $\geq 2$  low- $P_T$  muons seen, and have 3 spare bits for expansion. The 3  $P_T$  bins are derived from the XTRP unit which has higher precision than that available from the differential timing of the CMU/CMX projective wires. The  $P_T$  thresholds are separately programable for the various regions, CMU and CMX, both in the XTRP and muon primitives cards. To avoid double counting when computing  $\geq 2$  low or medium  $P_T$  muons, an intervening empty  $2.5^\circ$  segment is required for separate muons. These signals are transmitted to FRED as Level-1 trigger imperatives. If a Level-1 accept is received by the  $30^\circ$  card, it transfers the pipelined data (both input to and output from the  $30^\circ$  card) to one of four Level-2 buffers. The information in the buffer is also output to the Level-2 trigger. The Pre-Match cards provide similar data for Level 2.

### 12.3.6 Global Level-1

The Level-1 Decision card or “FRED” module is located in the Global Level-1 crate and is responsible for issuing the Level-1 trigger decision. Single-bit trigger signals from the various calorimeter, muon, and tracking PreFRED modules, which are also located in the Global L1 crate, are combined by FRED to form the final Level-1 triggers. The Run-II FRED

Detector	Basic Unit	Unit	CARD	Algorithm Description	# Outputs
CMU	Wire pair	1.25°	MUIT	Hi, low $P_t$ , plus a “lefover to 384ns” determined from differential timing.	288 x 2 x 2 $\phi \times \eta \times P_t$
CMP	4 tube stack	0.6° 1.2°	MPIT	2 or 3 out of 4 hits for patterns from radial tracks	336 $\phi$
CMX	Wire pair	1.25°	MXIT	Hi, low $P_t$ , plus a “lefover to 384ns” determined from differential timing	288 x 2 x 2 $\phi \times \eta \times P_t$
CSX	Coincidence	15°/8	MS1X	Gated Mean Time from 1/2 overlapped scintillators	192 x 2 $\phi \times \eta$
CSP	Scintillator	1.2° 2.4°	MS1P	Gated Scintillator hit	168 x 2 $\phi \times \eta$
HAD	Calorimeter	15°	MHIT	Signal in calorimeter PMT for $\eta$ intervals 0/4, 4/5 and 6/9	24 x 6 $\phi \times \eta$

Table 12.2: The muon primitives produced for each of the central muon detectors.

module (Level-1 Decision card) has the following features:

- receives up to 64 Level-1 inputs from a possible 8 PreFRED modules
- generates up to 64 Level-1 triggers, each based on a subset of 8 of the 64 Level-1 inputs
- has the ability to prescale or rate limit each of the 64 Level-1 triggers separately
- provides the Level-1 trigger decision to the Trigger Supervisor (TS) running the physics partition
- provides the 64 Level-1 trigger bits to the 7 Level-2 processors
- provides signals to scalers to count the Level-1 input rates, the raw Level-1 trigger rates, and the prescaled/rate-limited Level-1 trigger rates
- contains 4 Level-2 buffers that can be read out over VME

### 12.3.6.1 The Global Level-1 Crate

The Global L1 crate is a VME crate with  $9U \times 400$  mm slots. A schematic diagram of the crate is shown in Fig. 12.12. The 64 separate Level-1 inputs to the FRED module arrive from the various calorimeter,

muon, and tracking PreFRED modules over a custom J3 backplane. Each PreFRED module receives information via cables from the appropriate section of the trigger and summarizes it for presentation to FRED. Before being transmitted on the J3 backplane, the output bits pass through a FIFO which phases all 64 FRED inputs to ensure that they come from the same event. The TOWTRG and SUMET PreFRED modules, which summarize the calorimeter trigger, were described in section 12.2.1.3. Since the Muon and Track bits need no further processing in the PreFRED module, these PreFREDs contain only differential receivers and the time-alignment FIFO.

Event data is read out from the FRED and PreFRED modules using the standard VRC and TRACER modules discussed in section 11.6 in reference [6]. The VRC is also used to configure FRED at the start of a run. Custom-built scalers, which are also located in the Global Level-1 crate, keep track of the “raw” input signal rates and the final L1 trigger rates.

### 12.3.6.2 The Global Decision (FRED) Module

Since the FRED module is part of the run-II pipelined architecture, trigger decisions are made on a cyclic basis, with a cycle time of 132 nsec to accommodate the maximum crossing rate expected for run

## Global Level 1 Crate

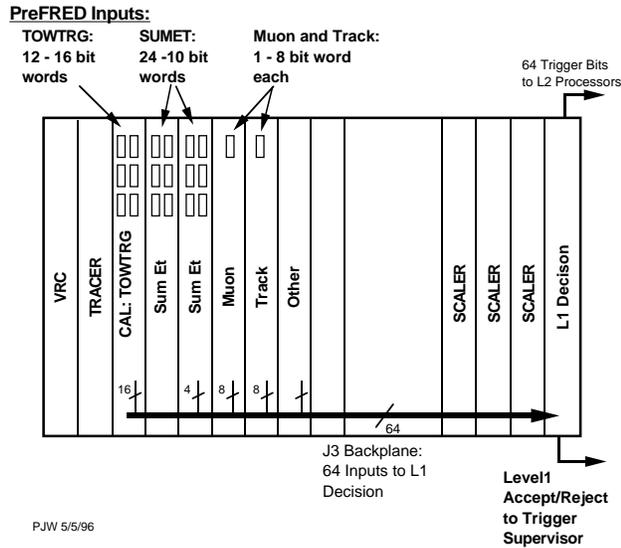


Figure 12.12: The Global L1 Crate.

II.

The FRED module expects the Level-1 inputs from the PreFRED modules to be valid at the rising edge of the 132 nsec system clock. The 64 inputs arrive in parallel, pass through buffers, and are latched. They are then multiplexed down to subsets of 8 inputs using AMCC S2024 “Crossbow” chips. The subset of inputs to be used is part of the information downloaded at the beginning of a run via the VME interface. Each subset of 8 inputs addresses a static RAM to determine if this pattern of 8 inputs constitutes a valid trigger. The contents of the SRAM are also downloaded at the beginning of the run. Each SRAM has 8 outputs. Thus a single “L1 trigger block” allows for 8 separate L1 triggers using a selection of 8 inputs (or some subset of these 8 inputs). The current design uses 8 such L1 trigger blocks. Thus for any given run there would be a total of 64 separate L1 triggers.

Each output line from the SRAMs is latched after a fixed delay in order to allow the SRAM outputs to become valid and stabilize. The rate-limiting and prescaling logic is then enabled. The presence of any trigger bits after this stage causes a Level-1 accept signal to be sent to the Trigger Supervisor.

The FRED module has the ability to independently prescale each Level-1 trigger. The prescale logic is implemented with a divide-by-N counter im-

plemented in Xilinx FPGAs. The values for the prescaling are downloaded at the beginning of the run.

The Level-1 Accept/Reject signal is sent to the Trigger Supervisor over a dedicated cable on the front of the VME card. All other communication between the FRED module and the Trigger Supervisor, such as the transmission of the final Level-1 decision back to FRED, is done via the TRACER module and the VME backplane.

The event data to be read from the FRED module include the latched Level-1 inputs and the Level-1 RAM outputs. The event data are stored in the standard L1 pipeline and L2 buffer system described in chapter 11. Upon receipt of a Level-1 accept from the Trigger Supervisor, the 64 Level-1 trigger bits will also be sent to the Level-2 processors through a dedicated cable to aid in the L2 trigger decision.

## 12.4 Level 2 Trigger Hardware

The Level-2 trigger consists of several asynchronous subsystems which provide input data to programmable Level 2 Processors in the Global Level-2 crate which evaluate if any of the Level-2 triggers are satisfied. In run II, we expect about 100 different Level-2 triggers to check compared 64 different triggers in run 1.

Processing for a Level-2 trigger decision starts after the event is written into one of the four Level-2 buffers on all front-end and trigger modules by a Level-1 accept. While Level-2 is analyzing the event that buffer cannot be used for additional Level-1 accepts. If all four Level-2 buffers are full then the experiment starts to incur deadtime. The time required for a Level-2 decision needs to be less than about 80% of the average time between Level 1 accepts in order to keep the deadtime at an acceptable level. For a 50 kHz Level-1 rate, this would mean that the average decision time needs to be 16  $\mu$ sec or less. The best estimate of the Level-2 processing time is 20  $\mu$ sec. The Level 1 rate would be limited to 40 kHz to keep the deadtime below 20%.

In order to keep the Level-1 rate at 50 kHz, we have decided to pipeline Level-2 in two stages each taking approximately 10  $\mu$ sec. Although the latency remains 20  $\mu$ sec, the time between Level-2 decisions will be 10  $\mu$ sec resulting in minimal deadtime even with a Level-1 accept rate of 50 kHz.

The first stage of the Level-2 pipeline is an event

building stage: data is collected from the Level-2 buffers of the Level-1 trigger systems (XFT and L1 MUON) and from the calorimeter shower maximum detector (XCES). Simultaneously a hardware cluster finder (L2CAL) processes the calorimeter data and a track processor finds tracks in the silicon vertex detector (SVT). All the Level-2 event data is collected into the memory of the Level 2 processors. We estimate that this stage will take 10 microseconds. After all of the data is stored in the processors the event is examined to find out if the criteria for any of the Level 2 triggers is satisfied. Analysis of an event in the L2 processors can occur at the same time as the data for the next event is loaded into memory.

#### 12.4.1 Level-2 Cluster Finder (L2CAL)

Since jets are not fully contained by trigger towers in the Level-1 trigger, the Level-1 thresholds must be set much lower than the jet energy to provide an efficient trigger. This results in rates that are too high for readout into Level-3. To provide a reduction in the jet trigger rates at Level-2 these triggers use the  $E_T$  from clusters of trigger towers. The Level 2 cluster finding algorithm is essentially the same as was used in run 1. In this algorithm, contiguous regions of calorimeter towers with non-trivial energy are combined to form clusters. Each cluster starts with a tower above a “seed” threshold (typically a few GeV) and all towers above a second somewhat lower “shoulder” threshold that form a contiguous region with the seed tower are added to the cluster. The cluster finder design effort has focused on implementing the algorithm in the multi-buffered deadtimeless architecture, with emphasis on improving performance with a parallel implementation of high speed Xilinx FPGA’s.

The data from the calorimeter system is collected and processed by the Level-1 trigger as described in section 12.2.1 and shown in Fig. 12.2. The towers are summed on the detector into trigger towers of  $0.2$  by  $15^\circ$ , resulting in a  $24 \times 24$  array for a total of 1152 towers (576 EM and 576 HAD). The tower energies are weighted by  $\sin\theta$  and are gain and offset corrected. The data from each crossing is output from the DIRAC boards, received by the Digital Cluster And Sum (DCAS) boards and stored in FIFOs. Upon a Level-1 accept the data is latched into one of four Level-2 buffers in the DCAS board.

The cluster finding process starts as soon the

DCAS boards are given the address of the next Level 2 buffer to process. The cluster finder algorithm is explained in Fig. 12.13. The size of each cluster expands until no towers adjacent to the cluster have energy over the shoulder threshold. After the complete cluster is found the tower energies are sent to an adder tree and the next seed tower is selected. The seed finding, clustering and summing processes are pipelined such that while the summing is carried out for one cluster the seed finding and clustering can start for the next cluster. For each cluster found the total EM and Had energies are calculated and recorded along with the number of towers, and the  $\eta$  and  $\phi$  coordinates of the seed tower. The energy weighted  $\eta$  and  $\phi$  positions which were available in run 1 are not calculated since they were found to be unnecessary.

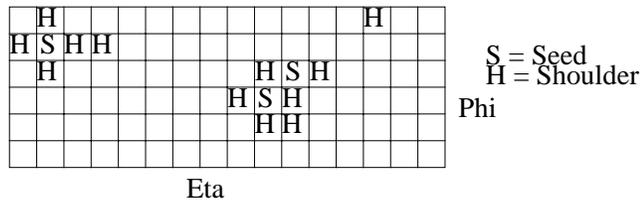
Each DCAS card performs the local clustering and summing for 16 trigger towers (8 EM + 8 HAD). For each  $15^\circ$  calorimeter wedge there are 3 DCAS boards which receive data from 4 DIRAC boards. The 12 DCAS cards are contained in each of 6 VME crates mounted directly above each of the L1CAL crates. A single local controller (LOCOS) provides control over clustering in each crate and forms crate-wide sums. The local controllers are in turn controlled by a single Cluster Queen (CLIQUE) in the L2 Processor crate. The CLIQUE forms the final cluster sum and forwards it to the L2 Processors.

After cluster finding on the event in buffer N is finished, cluster finding on another buffer can start. This allows the data collection and data analysis phases of Level 2 to be pipelined.

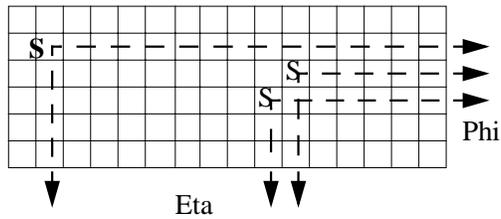
A “hook” for expansion is built into the cluster finder. The  $\eta$  and  $\phi$  address of each cluster is sent to a slot which can contain a Neural-Net Processor (NN). This board can read the individual energies from the seed tower and the surrounding towers. There are separate data paths for cluster finding and NN reading. This will allow the implementation of the NN photon trigger and NN tau trigger used in Run I.

#### 12.4.2 Level-2 Shower Maximum Data Used in the Trigger

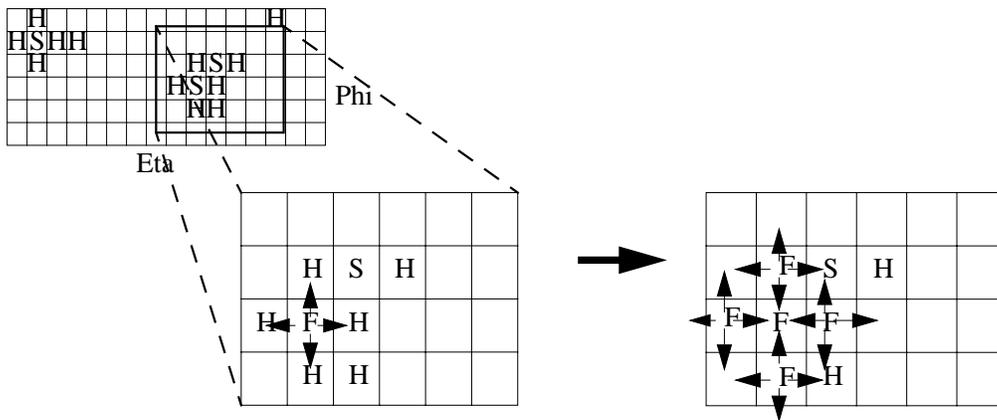
The shower maximum detectors of the EM calorimeters can be used to reduce trigger rates for electron and photon triggers. For both electrons and photons, requiring a cluster above threshold in the shower maximum detector eliminates the background



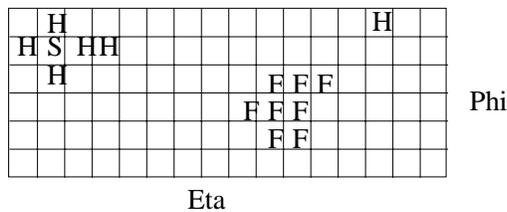
1. The inputs from DIRAC are checked to see if they are over the seed or shoulder threshold.



2. The Cluster Control board tells the DCAS's to enable their Phi and Eta address lines for all seeds. The Control board then selects the seed with the lowest Eta address, and then within that Eta, the lowest Phi address to begin cluster finding.



3. Cluster finding then begins with the selected seed sending a “found” bit to its four orthogonal neighbors. If the neighbors are over the shoulder threshold and receive a “found” bit, they in turn send out the signal to their four orthogonal neighbors, and so on, and so on.



4. When a tower is “found” it is disabled from being found in another cluster. i.e. if the tower is a seed and is found in a cluster initiated by another seed, it will not enable its address lines when the control board prompts for seed addresses. After one cluster is found, the tower energy values are enabled into an adder tree. Steps two through four are repeated until all clusters are found.

Figure 12.13: Calorimeter trigger cluster finding

from single-phototube discharge. Since the spatial resolution is much smaller than a calorimeter wedge, matching tracks from the central tracker to the shower maximum provides a significant reduction in combinatoric background for electron triggers.

CDF upgraded the electron and photon trigger hardware for run 1b to make use of shower maximum information from the central shower maximum detector (CES) in the Level-2 trigger decision[8]. For electrons, this upgrade resulted in a 50% reduction in backgrounds, while retaining approximately 90% of the signal. Topics in b physics such as the search for rare processes demand a low trigger threshold. By using the XCES, we were able to lower the inclusive electron threshold to 8 GeV, which is 1 GeV lower than the threshold for run 1a, while maintaining the same bandwidth of events accepted by the CDF hardware Level-2 trigger.

#### 12.4.2.1 Central Shower Maximum (XCES)

While the frontend electronics are completely different, the basic scheme for the XCES trigger is the same for run II as for run 1b. The signals from four adjacent CES wires are summed together and compared to a threshold to form an XCES bit with  $\sim 2^\circ$  azimuthal segmentation. This bit can then be matched in azimuth and momentum to a XFT track to generate a Level-2 trigger. There are 16 XCES bits per detector wedge (8 bits times 2 thresholds) yielding a total of 768 XCES bits.

The QIE for shower max digitizes each crossing, and stores the data in a level 1 pipeline, awaiting a Level-1 trigger decision. The data is compressed into a floating point format, consisting of 3 bits representing the measurement range, and 5 bits of ADC data, for a total of 8 bits. Once a Level-1 accept has been issued, 4 time slices for each wire of digitized data are received by the VME readout boards. The CES VME readout board will have an extra trigger circuit to produce the trigger output bits (XCES) to be used in the Level-2 decision. This circuit converts the floating point output to a linear format and sums the 4 time slices for each wire. The signals from four adjacent wires at a time are summed and compared to a threshold to define an XCES bit. Two ADC thresholds will be implemented in the trigger. One, on the order of a few GeV, will be used for the  $J/\psi \rightarrow ee$  trigger. The other, on the order of 8 GeV, will be used for the inclusive electron trigger. The XCES

bits are received by the Level-2 Processor within 10 microseconds of a Level-1 accept where the matching of the XCES cluster in momentum and azimuth to a XFT track is performed to generate a Level-2 trigger.

#### 12.4.2.2 Plug Shower Maximum (XPES)

Because this is a new detector, with a substantially different geometry from the CES, the options for using information from it in the level 2 trigger are numerous. The plug shower maximum interface is planned to match the signals of the strips with a signal over some threshold to a listing of the plug calorimeter towers that should also have a signal. There are two alternatives being investigated at present. The first sums the signals from 4 adjacent strips and compares the result to a threshold value, much as is done for the XCES. The 50 resulting bits per 45 degree phi sector are then used to match calorimeter towers to the strips. The other possibility is to use the last dynode output from each multi-anode PMT, which is equivalent to summing the signals from 16 adjacent strips. In either case some logic will have to be performed to correlate the strips over threshold in each view to the appropriate calorimeter tower using a series of ands and ors. The performance and possible implementation of each scheme is still being investigated at this point, and we plan to have the design finalized in the late fall of 1996.

#### 12.4.3 Silicon Vertex Tracker (SVT)

The ability to use impact-parameter information in the trigger to detect secondary vertices can substantially increase the physics reach of a hadron-collider experiment. For example, the resolution and mass-scale systematics for the top-quark mass determination could be measured and controlled if there were a large-statistics dijet mass peak from W or Z decay. Unfortunately the large QCD jet cross section makes the signal-to-background ratio poor and the trigger rate prohibitively high. Impact-parameter triggering, however, can provide a  $Z \rightarrow b\bar{b}$  signal, for which the continuum background is substantially reduced. With the Silicon Vertex Tracker (SVT), we expect a large sample of  $Z \rightarrow b\bar{b}$  events sitting well above the continuum background.

Studies of B decay will also be greatly enhanced by the SVT. Some of the decay channels that are important in the study of CP violation will be virtually undetectable at the Tevatron without an impact

parameter trigger (*e.g.*  $B^0 \rightarrow \pi^+\pi^-$ ) [9]. Other interesting decay modes (*e.g.*  $B^0 \rightarrow J/\Psi K_s$ ), although accessible by requiring one or two leptons in the final state, will benefit from an impact parameter trigger in terms of better rejection against background and higher statistics on tape.

Some physics processes not involving  $b$  quarks will also be aided by the SVT. The high- $P_T$  inclusive-muon trigger, needed for the W-mass measurement, has a high Level-2 rate. The SVT can both reduce the accidental rate by demanding an SVX track pointing at the primary vertex and remove lower- $P_T$  muon background by using the improved momentum resolution to tighten the  $P_T$  threshold.

In order to obtain impact-parameter information in the Level-2 trigger, the SVX II is read out after each Level-1 trigger. The SVT combines this data with the Level-1 tracking information from the central tracking chamber and computes track parameters ( $\phi$ ,  $P_T$ , and impact parameter  $d$ ) with resolution and efficiency comparable to full offline analysis. Figure 12.14 displays the resolution of the impact parameter in Monte Carlo for a full SVT simulation, including detector misalignments within tolerances, and for the offline reconstruction.

We give a brief summary of the SVT below. A full description of the device can be found in Ref.[10].

### 12.4.3.1 Implementation

The strategy we follow combines the use of the Associative Memory (AM) technique with a processor farm. The overall architecture of SVT is shown in Fig. 12.15.

SVX-II data digitization and sparsification occur in the front end. From each of the SVX-II readout chips, there is an 8-bit data stream consisting of the chip ID and status words and a series of word pairs containing a channel number and the digitized pulse-height in that channel. The data is transmitted over fiber optic lines to a splitter where one output goes to the SVX-II readout and the other output goes to the SVT. There are 12 fibers for each of the 12 SVX-II  $\phi$  sectors. These are fed into G-link Fiber-Receiver transition cards where the data is demultiplexed and transmitted through the backplane to the Hit Finder boards, three for each  $\phi$  sector. The Hit Finder performs pedestal subtraction and bad channel suppression. It then scans the data stream for clusters of hit strips. When acceptable clusters are found, the Hit

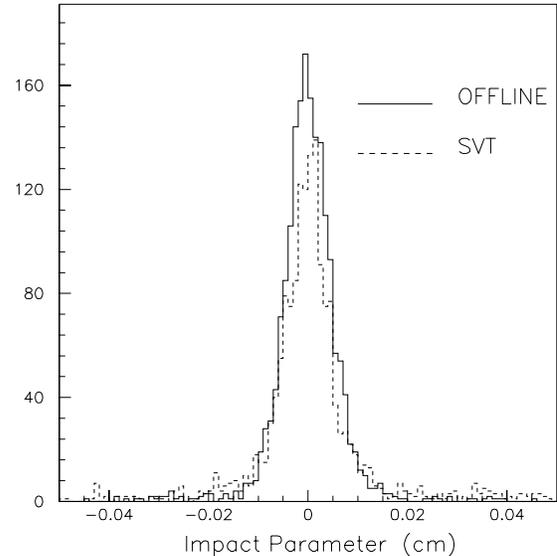


Figure 12.14: Comparison of the impact-parameter distribution of SVX tracks from the SVT trigger simulation and offline procedures.

Finder computes and outputs the centroid of each cluster, the most probable track intersection point at that SVX layer.

The output from the three Hit Finders for each  $\phi$  sector are merged into one stream with a Merger board. Communication between SVT boards is standardized. The transmission is over 50-conductor flat-ribbon twisted-pair cable. The data is pushed by the transmitting board at 30 MHz into a FIFO on the receiving board.

While the Hit Finders process SVX data, the list of tracks found in the outer tracking chambers during Level-1 processing is transmitted from the XTRP to the Associative Memory Sequencer (AMS). As soon as the clusters from the Hit Finder are available from the output of the Merger, they are also transmitted to the Associative Memory Sequencer. The AMS uses a lookup table to convert the 14-bit cluster centroid into a coarser Superstrip in preparation for the pattern-recognition stage. The size of the Superstrip is programmable. It is a compromise between a small size that would provide more precise pattern recognition and produce fewer fake-track candidates but require a larger memory, and a large size which would output more fakes but require a small memory. (The fakes would be rejected at the next stage, but they

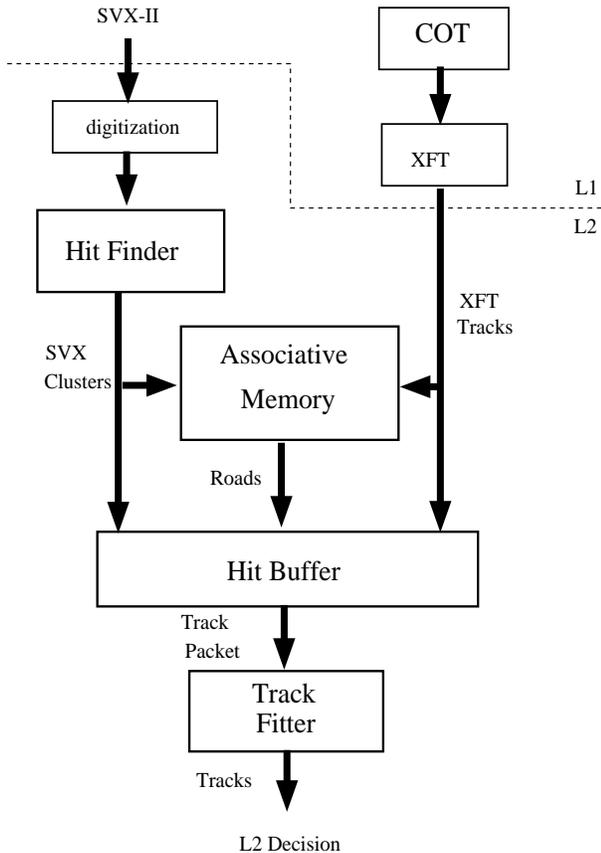


Figure 12.15: Architecture of the SVT trigger.

waste time.) A study of the compromise results in the choice of a  $250 \mu\text{m}$  Superstrip.

Each AMS broadcasts the outer tracks and coarse SVX clusters on the backplane where they are received by two Associative Memory (AM) boards. Each AM board fans this data stream out to 128 AM chips mounted on the board. The AM chip is a custom CMOS device that contains 128 roads. It is implemented in  $7\mu$  technology and runs at 30 MHz. Each chip stores 128 combinations of outer tracks and coarse SVX hits. Each combination represents a legitimate particle trajectory and is called a *road*. As the outer tracks and clusters are received, each chip checks to see if all of the components of one of its roads (outer track and four SVX hits) are present in the data stream. If so, the chip records the track candidate's road number. The roads are downloaded, so they can be modified, for example to correct for geometric misalignment or to select which layers to use.

At the same time the clusters and outer tracks are sent to the AMS, they are also sent to the Hit Buffer where they are stored by Superstrip number. Once the AM is done, the road numbers of candidates are also sent to the Hit Buffer. The hit buffer uses

another lookup table to find the Superstrips which correspond to each road. It then goes to those Superstrip bins and retrieves the full outer-track and SVX-cluster information.

This set of one outer-track and four SVX hits is called a track packet. The Hit Buffers sends these packets to the Track Fitters. The collection of Track Fitters, the Track-Fitter Farm, is designed so that tracks are distributed equally to all Fitters, and enough Fitters are included so that each one fits on average one track per event. Studies indicate that ten Fitters will meet these requirements. The fit is a linear approximation, consisting of a set of scalar products. It takes 70 cycles on a currently available commercial CPU, approximately  $3.5\mu\text{s}$ . Other methods currently under study may be significantly faster, even though the processing time with the current technology is not a concern.

#### 12.4.3.2 Expected Performance

The resolution of the fit is nearly as good as offline reconstruction. The expected resolution is  $\sigma_d = 35 \mu\text{m}$  (at  $P_T = 2 \text{ GeV}$ ),  $\sigma_\phi = 1 \text{ mrad}$  and  $\sigma_{P_T} = .3\%P_T^2$  (where  $P_T$  is in GeV). The track parameters are available for the Level-2 trigger decision.

To see the power of the SVT, consider  $B^0 \rightarrow \pi^+\pi^-$ . This is one of our most difficult trigger problems because the rate for two low- $P_T$  tracks is so high. At Level 1, only the outer-track information is available. We will require one track with  $P_T > 2 \text{ GeV}$ , and a second, oppositely-charged track with  $P_T > 3 \text{ GeV}$  and  $30^\circ < \Delta\phi < 135^\circ$ . The rate for this at  $\mathcal{L} = 10^{32}$  is 16 kHz, well within the Level 1 design limit of 50 kHz. At Level 2, the SVT provides impact parameter information. Applying full SVT simulation to SVX' data taken during run Ib, we find that a cut of  $|d| > 100 \mu\text{m}$  provides three orders of magnitude in rejection; such a cut is 50% efficient for the signal. The resulting Level-2 rate is  $< 10 \text{ Hz}$ , much smaller than the design limit of 300 Hz.

There are two potential problems that must be controlled. The first is the precise mechanical positioning of the SVX II. Since the SVT does not have  $z$  information for tracks, an angular misalignment between the SVX and the beam would produce incorrect impact-parameter results. Therefore SVX-II detectors must be aligned collinear to the beam to within  $100 \mu\text{rad}$ . The ladders themselves must be aligned to each other within this specification during

## Level 2 Trigger Block Diagram

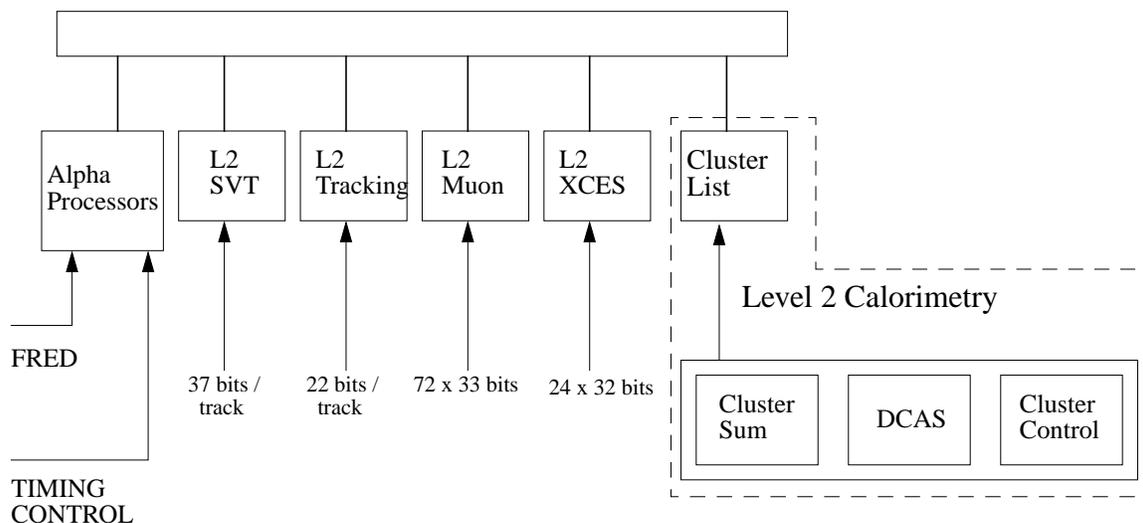


Figure 12.16: The Level-2 decision crate.

construction. It will also be necessary to align the beam to the detector during the first few minutes of each run. Studies have been performed showing that the SVT itself can provide the needed feedback to the accelerator control room.

The other potential problem is the occupancy of the SVX II (Section 5.8.1), which directly impacts the ability of the SVT to complete its calculation in time for the Level-2 trigger decision. To study this, a Monte Carlo simulation was tuned until it could accurately reproduce the occupancy of the current SVX detector over a wide range of luminosities. The occupancy is observed to be proportional to the luminosity, as expected. The tuned Monte Carlo was run with an SVX-II simulation and run-II luminosity to produce a data sample for SVT timing simulations. The result shows that the SVT will complete its calculations well before the end of its allotted  $17\ \mu\text{s}$ .

### 12.4.4 Level-2 Decision Hardware (Global Level 2)

The final Level-2 decision is made in a set of four Level-2 Processors in the Level 2 crate. There are five major sources of data for the Level 2 trigger. The data from each source is collected into a single board in the Level 2 crate as shown in Fig. 12.16. The data is stored in a device dependent format in FIFO's. The data from the output of the FIFO is formatted into a 128 bit word and transmitted to the

Level 2 processors. The processors store the data directly into memory locations (common blocks in Fortran parlance). The storing of data in all processors occurs simultaneously. Each data source maps the type of data (tracks, calorimeter clusters, SVT, muon, XCES) and the Level 2 buffer number into the distinction address. This allows data for one buffer to be assembled while the processors are examining an event in a different buffer.

The cluster list board collects information about each calorimeter cluster. The  $\eta$  and  $\sigma$  address, number of towers, and EM and hadronic energy for each cluster is encoded into a 64 bit word. A 2 bit field indicating which one of the four sets of seed and shoulder thresholds was used is stored in each cluster list element. The L2 buffer address is also sent for diagnostic purposes. The cluster list card takes two clusters and transmits them in a 128 bit transfer to the Level 2 processors.

The trigger data from the central strip chambers (XCES) consist of 32 bits per wedge, for a total of  $24 \times 32$  bits. This data is packed into six 128 bit words and transmitted to the Level 2 processors. The XCES information is used to require hits in the central strip chambers for low energy electrons and photons.

The tracking data comes from two sources, the XFT tracks from the XFT via the XTRP, and the tracks from the SVX processed by the SVT to give impact parameter. The tracks from XTRP are trans-

mitted immediately after a Level 1 accept. They are received by the L2 tracking module, packed 4 tracks per 128 bits and transmitted to the Level 2 processors.

The data from the SVT arrives later than the data from the other four systems. The data from the SVX must first be transmitted from the detector, then processed in the SVT and combined with XFT tracks. This will take a minimum of  $7\mu\text{sec}$  and most likely  $10\mu\text{sec}$ . Since this is the total time allocated to collect Level 2 data, the processors will start analyzing the event before the SVT data is complete. Only if the impact parameter required to make the Level 2 decision will the SVT data be used. If all triggers which require SVT impact parameters are rejected by other cuts (e.g. kinematics), then the impact parameter cut won't be tested.

The data from the muon match board is zero suppressed and then transmitted to the Level 2 system. The muon information is packed into 72 bits representation each  $2.5^\circ$  wedge for each muon system, and then transmitted to Level 2 processors.

#### 12.4.4.1 Level-2 Processors

The Level 2 processors are custom VME modules based on the Dec 21164 Alpha chip. Figure 12.17 shows a block diagram of the processor board design which is an extension of the processors used in run 1b. The new processors have a 433 MHz clock instead of a 150 MHz clock. The data path into the board is 128 bits wide instead of 64 bits. The memory and I/O interface use the 21172 chip set instead of discrete logic. The boards use VME format and reatout instead of Fastbus used in run 1b.

The data from the data collection boards is transmitted over a 128 bit wide bus ("Magic Bus") using a custom J3 backplane. The address is used as a 'DMA channel' and determines which of the 4 L2 buffers are being processed and which of the 5 data sources are sending the data. There will be four Level 2 processors in the system.

All four Level 2 processors latch the data from all the data collection boards. The data is initially stored in a FIFO and then transmitted to the main memory. Care will be taken that successive Level 2 buffers do not map to the same cache lines. This will allow the processor to analyze one event out of the third level cache while data from another Level 2 buffer is being written to main memory and invalidat-

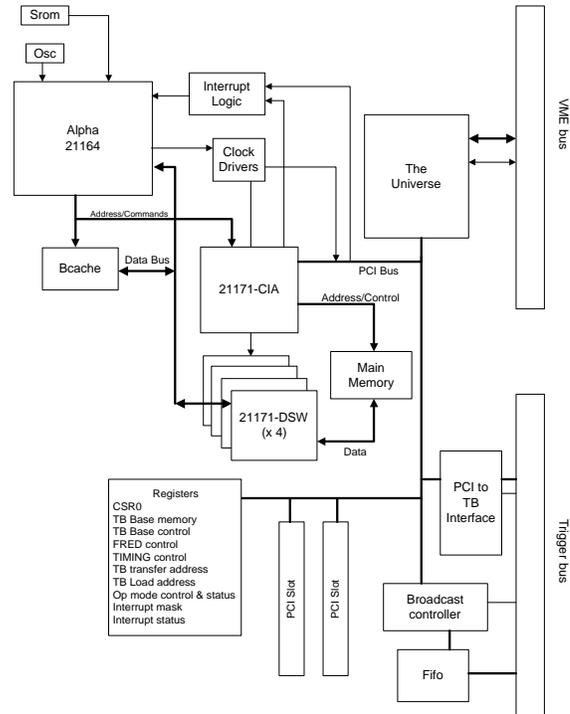


Figure 12.17: Block diagram of the Level-2 processor board.

ing the cache lines holding the event the previously came in on this buffer. The second level cache is on chip and is 3-way set associative.

Each of the four Level 2 processors will be assigned a set of triggers to examine. The triggers will be distributed across the processors by hand. All muon triggers will be in one processor and all electron triggers will be in another. The processor forming the electron plus muon trigger will have to find the low  $P_T$  electrons and muons in parallel with the electron processor and the muon processor. This allows the triggers to be ordered such that most triggers will not have to be evaluated. For example, if the 6 GeV central muon trigger fails, then any trigger requiring a 6 GeV or greater muon need not be tested. This kind of optimization was used successfully in run 1. Each of the four processors will independently forward a Level-2 accept or reject for the set of triggers it is examining to the Trigger Supervisor.

Negotiations are underway with the Michigan State DØ group to use the same processor in DØ.

## 12.5 Trigger Rates

The CDF trigger system is very flexible. It can pick event topologies that are simple, like a single high- $P_T$  lepton, or complex, like two charged particles in a narrow mass window coming from a secondary vertex, or a multiple-jet system with large  $\cancel{E}_T$ . In run I, we normally took data with over 50 separate trigger paths. This will also be the case in run II. However, as in the past, the trigger rate, and thus potential deadtime problems, will be dominated by a few triggers. Here we focus on the three highest-rate triggers; the single-electron and single-muon triggers give the largest Level-2 rates, and the two-track trigger has the highest Level-1 rate.

The high- $P_T$  inclusive-electron trigger is used for the W and top-quark studies as well as many exotic-particle searches. In the central rapidity region, current trigger rates can be reliably extrapolated to run II. For a calorimeter tower with EM  $E_T > 16$  GeV matched to a track with  $P_T > 12$  GeV, the Level-2 trigger rate will be 20 Hz for  $2 \times 10^{32}$  luminosity and 132 nsec bunch separation. Since both calorimeter and tracking information are available to the Level-1 system, that rate won't be much higher. In the plug rapidity region, studies of the trigger rate with IFT/SVT tracks are in progress. We do however have an upper limit based on use of calorimeter information only. The Level-2 rate is  $< 100$  Hz. A track requirement should reduce this by at least a factor of two.

A lower-threshold central-electron trigger is important for B physics and calibration samples for the W mass. The Level-2 trigger rate for  $E_T > 8$  GeV will be 85 Hz. For calibration samples, prescaling is an acceptable way to reduce the rate. For B physics, a large impact-parameter requirement in the SVT reduces the rate by more than an order of magnitude.

For the single-muon trigger, again there are two rapidity regions to be considered. In the central region covered by the CMUP detector, the Level-1 trigger rate for  $P_T > 12$  GeV is 45 Hz, again for  $2 \times 10^{32}$  luminosity and 132 nsec bunch separation. At Level 2 with its tighter match between the track and muon stub, the rate drops to 15 Hz. For  $0.6 < |\eta| < 1.0$ , the region covered by the CMX system, the Level-2 rate will also be 15 Hz for  $P_T > 12$  GeV. This rate increases to 25 Hz for  $P_T > 6$  GeV, so lower thresholds for B physics can be accommodated.

The greatest challenge for the Level-1 system is

the  $B^0 \rightarrow \pi^+\pi^-$  trigger. As presented in the SVT section, the Level-1 rate is 16 kHz at  $10^{32}$  luminosity when one track is required with  $P_T > 3$  GeV and a second track is found with  $P_T > 2$  GeV and  $30^\circ < \Delta\phi < 135^\circ$  relative to the first track. At Level 2, with an impact parameter  $> 100 \mu\text{m}$  required, the rate drops to  $< 10$  Hz. This will provide 6000  $B^0 \rightarrow \pi^+\pi^-$  events on tape per  $\text{fb}^{-1}$  for a branching ratio of  $1 \times 10^{-5}$ .

The trigger system will allow us to collect data for our broad range of physics topics even at the highest planned luminosity.

# Bibliography

- [1] C.Campagnari, M.Miller, M.Shochet, and G.Sullivan, *Estimates of trigger rates and efficiencies for the 1993 CDF calorimeter level 0 trigger.*, CDF/DOC/TRIGGER/CDFR/1180 (1990).
- [2] F.Ukegawa, *Low Momentum Electron Trigger Rates*, CDF/ANAL/BOTTOM/CDFR/2480 (1994).
- [3] H. Frisch and P. Wilson, *Trigger Tower organization and summing in  $\eta - \phi$  space for Run II and beyond*, CDF/DOC/TRIGGER/CDFR/-2045 (1993).
- [4] C. Campagnari *et al.*, *Specifications for the Calorimeter Front End Electronics*, CDF/PIG/-DOC/2060, (1994).
- [5] T. Shaw and G.Sullivan, *A CDF standard VME readout crate for the Upgrade*, CDF/DOC/TRIGGER/CDFR/2388, (1993).
- [6] T. Shaw and G. Sullivan, "A Standard Front-End and Trigger VME Readout Crate for the CDF Upgrade," CDF note #2388; T. Shaw and W. Stuermer, "Designer's Guide for Modules to go into the CDF Custom Backplane."
- [7] "Trigger Supervisor Protocols for Run II", CDF Group/Yale University, October 10, 1995.
- [8] K. Byrum *et al.*, NIM in Physics Research A364 (1995) 144-149.
- [9] S.Donati and G.Punzi, " $B^0 \rightarrow \pi\pi$  Trigger Studies with Run 1a Data", CDF Note #3167, 1995.
- [10] S.Belforte *et al.*, "Silicon Vertex Tracker Technical Design Report", CDF Note #3108, 1994.