DØnote # ????

A Silicon Track Trigger for the DØ Experiment in Run II

— Proposal to the Fermilab PAC —

H. Evans[1], U. Heintz[2,3], T. Heuring[5], J. Hobbs[4], M. Johnson[2], S. Mani[6], M. Narain[2,3], F. Stichelbaut[4], H. Wahl[5]

[1] *Columbia University*
[2] *Fermi National Accelerator Laboratory*
[3] *Boston University*
[4] *State University of New York at Stony Brook*
[5] *Florida State University*
[6] *University of California, Davis*

August 20, 1998

Contents

1 Introduction

1.1 Purpose of this Document

This document describes the design of a Silicon Track Trigger preprocessor (STTpp) for the DØ detector in Run II. The STTpp interfaces with the approved trigger system of the DØ detector as an additional Level-2 preprocessor and thus requires no modification of the approved system. The STTpp performs a precise reconstruction of charged particle tracks in the Central Fiber Tracker (CFT) and the Silicon Microstrip Tracker (SMT). It allows the selection of events that contain decays of long-lived particles, *e.g.* b quarks, at the trigger level and enhances the capabilities of the DØ detector to trigger on charged particles in general. The STTpp improves the capabilities of the DØ detector to exploit the large range of high-p_T physics that will be accessible in Run II.

This introductory section outlines the motivation for proposing such a device and summarizes the components of the DØ detector that it will interface with. Section 2 describes the conceptual design of the processor hardware and section 3 addresses various performance issues. Section 4 gives a preliminary schedule and cost estimate and funding outline. The document concludes with a summary in section 5.

1.2 Physics Motivation

1.2.1 The Physics Program for Run II

During Run II, the Fermilab Tevatron will run at a center-of-mass energy of 2 TeV and deliver an integrated luminosity of 2–4 fb^{-1}, 20–40 times more than during Run I. This large data set will allow sensitive studies of the standard model of the electroweak interactions [56]. Among these are intermediate vector boson physics, a detailed study of the top quark, and the search for the Higgs boson and for supersymmetry or other extensions of the standard model. These issues address some of the most fundamental problems in particle physics research today: the question of electroweak symmetry breaking and the generation of masses.

1.2.2 Studies of the Top Quark

The dominant production process for top quarks in $p\overline{p}$-collisions is $t\overline{t}$-pair production via the strong interaction. In a data set of 4 fb^{-1}, we expect to identify several thousand $t\overline{t} \to \ell\nu bq\overline{q}\overline{b}$ events. This sample of $t\overline{t}$ decays will allow us to study the production dynamics and measure the decay branching ratios of the top quark. If the top quark, for example, couples to an unobserved heavy state, as some theories suggest, it will manifest itself in the shape of the $t\overline{t}$ mass distribution.

The top quark mass is of particular interest. It plays an important role in radiative corrections to many processes. The Standard Model prediction for the mass of the W boson includes radiative corrections which depend on the masses of the top quark and the Higgs boson. Figure 1 shows the predicted W boson mass as a function of the top quark mass for several assumed values of the Higgs boson mass. The data point indicates the world average measurements of the W boson mass [3, 4] and the top quark mass [5, 6]. The measurement of the W boson mass is one of the most sensitive tests of the Standard Model since its value follows directly from the structure of the model. Assuming the Standard Model as correct, the measurements can be translated into a constraint on the mass of the unobserved Higgs boson. The top quark mass is also interesting because of its sheer size. With a mass of about 175 GeV, the top quark is the heaviest known fundamental particle and it may well hold the key to understanding electroweak symmetry breaking and the generation of the fermion masses.

The most precise top quark mass measurements from Run I use the $t\overline{t}$ decay channel in which one top decays to $q\overline{q}b$ and the other to $e\nu b$ or $\mu\nu b$, the "lepton+jets" channel. The precision of the DØ measurement (about 7 GeV) is limited by the number of events, the jet energy scale calibration, and the combinatorics due to the large jet multiplicity. Each of the four final-state quarks fragments into a jet of hadrons. If light quark and b jets cannot be distinguished, there are 12 different ways to assign the four jets to the decay of the t and $\overline{t}$ quarks. If both b jets can be identified, this reduces to a mere two-fold ambiguity, significantly reducing the associated systematic uncertainties. The jet energy scale uncertainties can be reduced by reconstructing W or Z boson decays to jets. A measurement of the position of the invariant mass peak would help calibrate the jet energy scale, which is the largest systematic uncertainty in the top quark mass measurement. The

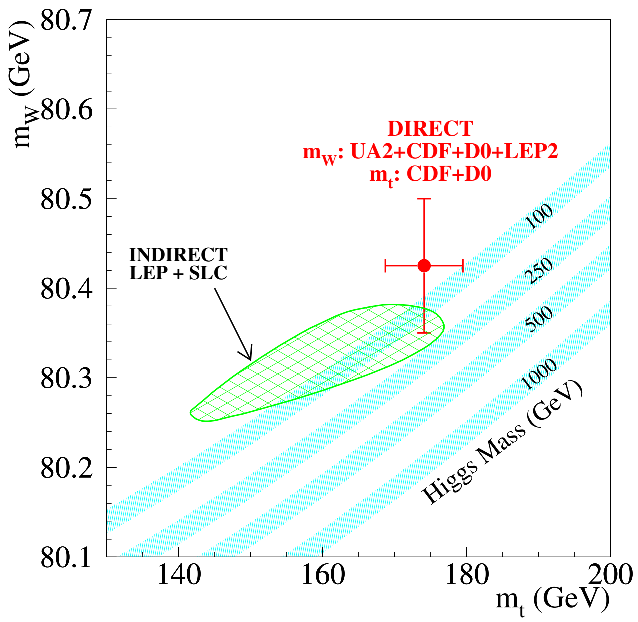

Figure 1: Predicted W boson mass versus the top quark mass (bands).

most promising channel is the decay $Z \rightarrow b\bar{b}$, in which b-jet tagging can provide rejection of the dominant gluon-jet background. This rejection is already needed at the trigger level, to keep the event rate low enough for the experiment to be able to run an unprescaled trigger. With these additional handles on systematic uncertainties we should be able to measure the top quark mass to about 3 GeV precision in Run II.

Production of single top quarks via the electroweak interaction is also an interesting process to study. The cross section is smaller than for $t\bar{t}$-pair production so that this process could not be observed in the Run I data. For $m_t = 170$ GeV, Run II with an integrated luminosity of 4 fb^{-1} will produce approximately 8000 single top quarks [49]. Single top quark production proceeds via a Wtb-vertex and its cross section is thus proportional to $\Gamma(t \rightarrow Wb)$ which, in turn, is proportional to $|V_{tb}|^2$. A measurement of the production cross section for single top quarks will probably provide the most precise measurement of the width of the top quark.

1.2.3 Search for the Higgs Boson

With a large data set (> 5 fb^{-1}), a signal for $H \rightarrow b\bar{b}$, produced in association with a W boson decaying to $e\nu$ or $\mu\nu$, can be observed at the Tevatron for $80 < m_H < 120$ GeV [56]. This mass window is extremely interesting from experimental and theoretical points of view. The current lower bound on the Higgs mass is 89 GeV[54] and is expected to reach approximately 105 GeV by the end of the LEP2 running [55]. The Tevatron can thus extend the accessible Higgs mass range beyond the reach of LEP2 before LHC starts running. Fits to the present electroweak data, assuming the completeness of the standard model, give $m_H = 115^{+116}_{-66}$ GeV [4]. In supersymmetric extensions to the Standard Model the lightest Higgs boson is expected to have a mass below about 130 GeV. It is therefore imperative that this window be covered reliably.

To see a Higgs signal during Run II, the inclusion of additional channels, like Higgs bosons produced along with hadronically decaying W or Z bosons, is required to increase the sensitivity beyond the study in reference [56]. For WH production, 67% of the decays are to $q\bar{q}b\bar{b}$. For ZH production 70% of the events decay to $q\bar{q}b\bar{b}$. In order to exploit these channels we have to be able to acquire all events with $q\bar{q}b\bar{b}$ final states for analysis. This requires rejecting backgrounds without b quarks at the trigger level.

To interpret the results of a Higgs search, the observation of a $Z \rightarrow b\bar{b}$ signal is almost mandatory.

It would prove that we can see a narrow resonance that decays to $b\bar{b}$ and it would also provide essential measurements of b-tagging efficiency and $b\bar{b}$-mass resolution.

1.2.4 Technicolor and Topcolor Theories

Technicolor breaks the electroweak symmetry dynamically by introducing a new strong interaction. It predicts the existence of particles in the TeV mass range, some of which couple to mass and therefore like to decay to final states containing b quarks. In $p\bar{p}$ collisions the techni-ρ (ρ_T) and the techni-ω (ω_T) may be produced. The ρ_T likes to decay to technipions which decay to heavy fermions, *e.g.* $\rho_T^+ \rightarrow \pi_T^+\pi_T^0 \rightarrow b\bar{c}b\bar{b}$. If the mass of the ρ_T is less than twice the mass of the π_T, the decay $\rho_T^+ \rightarrow W^+\pi_T^0$ dominates. This decay has a signature identical to that of Higgs production in association with W bosons, but a significantly higher cross section because of the ρ_T-resonance enhancement, so that it should be visible in Run II [7].

Topcolor explains the high mass of the top quark through new strong dynamics. It predicts Z' bosons and top gluons that preferentially couple to the third fermion generation. Their decays to $b\bar{b}$ would give rise to resonance structures in the $b\bar{b}$-mass spectrum.

1.2.5 The Importance of b-Tagging

In all the processes described in this section, b quarks appear in the final state. The identification (tagging) of b-quark jets will therefore play a central role in the high-p_T physics program during Run II. It can be used to reduce backgrounds from light-quark and gluon jets and to reduce combinatoric effects. If b-quark jets can be tagged already at the trigger level, the rate of triggers for these processes can be reduced without loss of efficiency. This is especially important for processes with all-hadronic final states like $WH \rightarrow q\bar{q}b\bar{b}$ or $Z \rightarrow b\bar{b}$. Triggers for these processes are swamped by gluon-jet background. As a result, triggers with jet-p_T thresholds low enough to see a mass resonance in the 100 GeV range cannot be operated without significant prescale. Additional background rejection, like that afforded by tagging b jets, is required to exploit the full integrated luminosity delivered by the Tevatron.

This is the motivation for proposing a trigger processor for the DØ experiment, that processes the data from the SMT at trigger level 2. This processor reconstructs tracks from charged particles precisely enough that tracks from the decay of long-lived b quarks can be identified. The b quarks produced in $p\bar{p}$ collisions live long enough (about 1.5 ps) to travel several millimeters. Tracks of charged particles from b quark decays can therefore be identified by their large impact parameter with respect to the beam trajectory.

1.3 The DØ Trigger System

During Run II, the proton and anti-proton beams will cross in the center of the DØ detector every 132 ns, resulting in a crossing frequency of 7 MHz. On the other hand the data acquisition system can only read out the detector and write the information to tape at a rate of 50 Hz. The function of the trigger system is the recognition of beam crossings during which an interesting event occurred that should be stored for off-line analysis. In order to perform this function efficiently, the DØ experiment will have a three-level trigger system in Run II.

Level-1 is entirely implemented in hardware. A logic network tests for calorimeter towers with energy above preprogrammed thresholds and hit patterns in the muon system and the CFT that are consistent with tracks above preprogrammed transverse momentum thresholds. This test is complete 4.2 μs after the beam crossing. The front ends are buffered by a pipeline 32 levels deep, which makes the Level-1 trigger deadtimeless. If the signals from an event satisfy one of the preprogrammed Level-1 trigger conditions the digitization of the detector signals is initiated and the event is passed on to the second trigger level. The maximum rate at which Level-1 can accept events is 5–10 kHz. As soon as the digitization is initiated, the detector stops accepting new events.

Level-2 consists of dedicated preprocessors, each of which reduces the data from one detector system (calorimeter, muon system, CFT, preshower detectors), and a global Level-2 processor which collects the data from all preprocessors and takes the trigger decision based on the combined information of all detector systems. This decision has to be taken in a very short time to limit the deadtime incurred while the Level-2 is processing the event. The mean decision time for Level-2 is 100 μs, resulting in no more than 5% deadtime. Level-2 can accept events and pass them on to the third trigger level at a rate of 1 kHz.

Level-3 consists of a farm of fast hight-level computers and performs a simplified reconstruction of the entire event. Events that satisfy the desired characteristics can be written to a permanent storage medium at a sustained rate of up to 50 Hz.

1.4 The DØ Silicon Microstrip Tracker

Figure 2 shows a view of the SMT which consists of cylindrical barrel sections and disks between the barrel sections and at the two ends of the detector. The barrel sections provide precise measurement of tracks in the central region. The luminous region of the Tevatron has an rms length of 25 cm in beam direction, which requires the barrel section to be long (72 cm). The disks extend the acceptance of the detector to forward tracks.

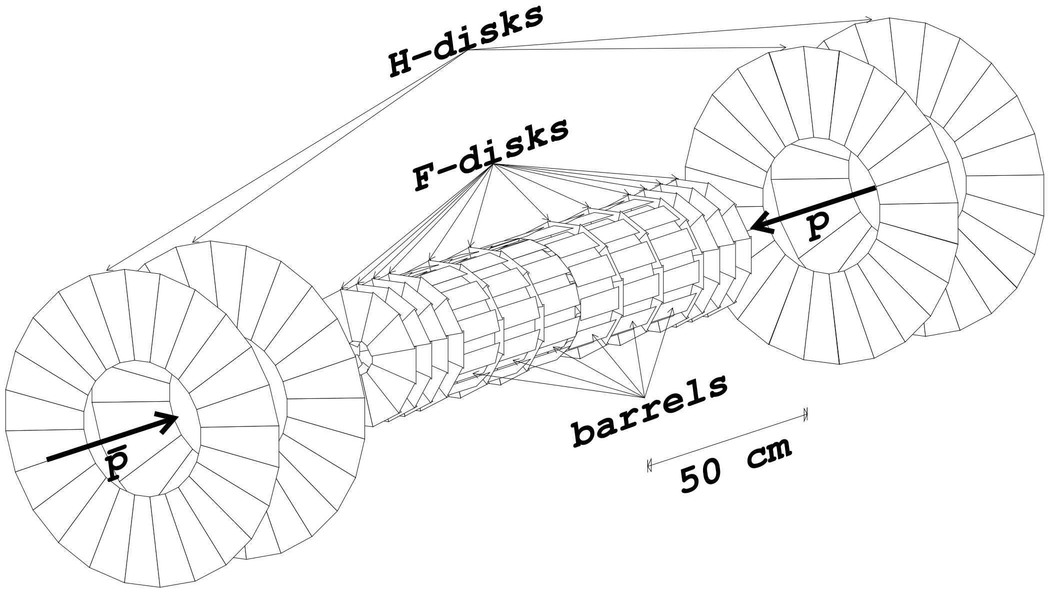

Figure 2: View of silicon microstrip tracker.

The barrel sections consist of rectangular silicon detectors arranged in four layers. Figure 3 shows the arrangement of the detectors in the plane transverse to the beam direction. All detectors have axial strips (parallel to the beam axis).

Figure 4 shows three 30° sectors in the silicon detector. The lines indicate the active areas of the ladders and the shaded regions the intervals in azimuth for which tracks with infinite momentum, coming from the origin, are contained in one 30° sector. There is only a tiny interval (1/4°), for which tracks hit two sectors. The acceptance loss from tracks that cross sectors depends in general on their p_T and impact parameter b. Figure 5 plots the acceptance loss versus curvature ($=1/p_T$) of the track for several impact parameter values. The plot of acceptance loss reflects the slight asymmetry of the detector between tracks with positive and negative curvatures and impact parameters. For $|p_T| > 1.5$ GeV the acceptance loss is never more than 1.2% for $b = 0$. For $b = 1$ mm it is on average 1.0% and for $b = 2$ mm it is 2.2%. This loss in acceptance is driven by the small overlap between the active regions of the ladders in layer 3. If tracks with hits in only the other three layers (1, 2, and 4) are accepted, there is no acceptance loss. We can thus safely ignore tracks that cross between sectors and divide the SMT into 12 independent sectors, each covering 30° in azimuth as shown by the shaded regions in figure 3. This substantially simplifies the design. The STTpp can then be broken down into 12 sectors that do not communicate with each other.

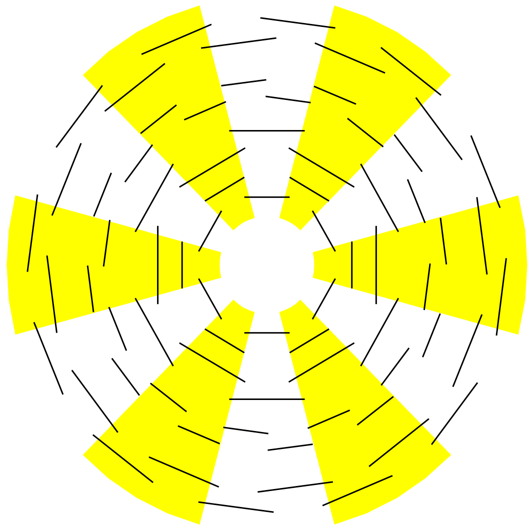

Figure 3: Arrangement of barrel detectors in transverse plane.

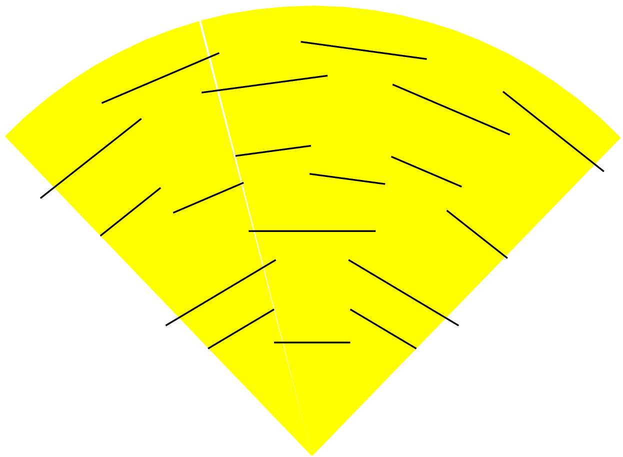

Figure 4: Three sectors in the silicon detector.

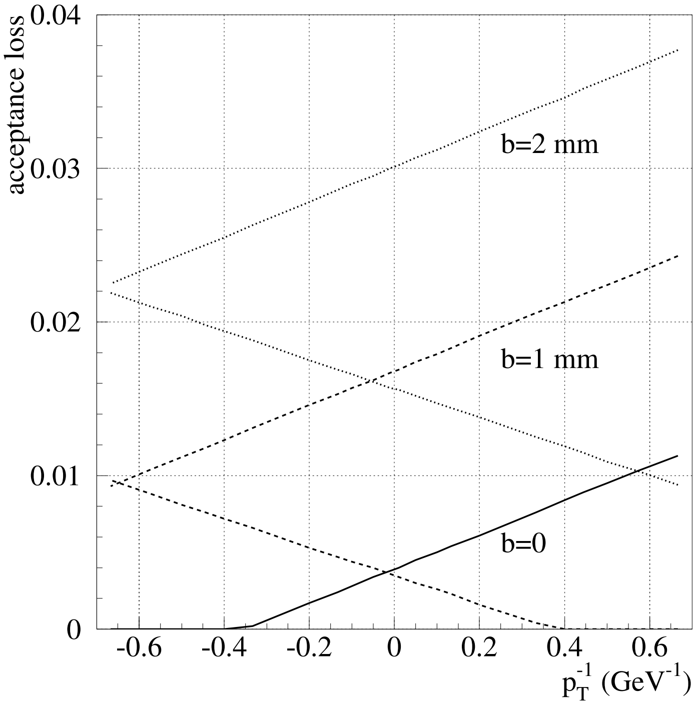

Figure 5: Acceptance loss vs curvature.

2 Conceptual Design

2.1 Overview

This section describes the conceptual design of the STTpp. The STTpp consists of fast digital electronics that process input data from two detector components: the CFT and the SMT. The CFT Level-1 trigger communicates a list of track candidates to the STTpp. Each track candidate defines a road in the SMT. The STTpp receives all data from the SMT barrel detector portcards. It forms hit clusters from the raw data and filters these clusters, retaining only the clusters that fall into one of the roads defined by the CFT data. Of all the hit clusters in a road, it finds the hits most consistent with the corresponding CFT track and fits a trajectory to the selected SMT hits and the CFT track. Then it communicates a list of fitted track parameters to the Level-2 CFT preprocessor (CFTpp), which merges the information with the Level-1 CFT information and formats the data for transmission to the global Level-2 processor.

As discussed in section 1.4, the SMT detectors can be divided into 12 sectors, each covering 30° in azimuth. In the STTpp we ignore tracks that cross sectors, which substantially simplifies the design. The STTpp can then be broken down into 12 subunits, each covering one sector, that do not communicate with each other. We have to custom-design the electronics for these subunits in order to optimize them for the required speed. Two such subunits share a VME crate. There are six such crates. In each crate there is one fiber road card, which receives the information from the CFT Level-1 trigger and broadcasts it across the backplane to up to 12 trigger cards. The trigger cards convert the CFT information into roads in the SMT. They receive the SMT data from the portcards, find hit clusters and sort the clusters by roads. The data processing on the trigger cards takes place in field programmable gate arrays (FPGA). The track fitting is performed by digital signal processors (DSP), located on one or more processor cards.

The electronics in the CFTpp, consisting of processors and communication links are copies of components that are already being designed for other parts of the DØ trigger system. In this way the STTpp will fit seamlessly into the existing Level-2 design. Figure 6 shows the crate structure of the STTpp.

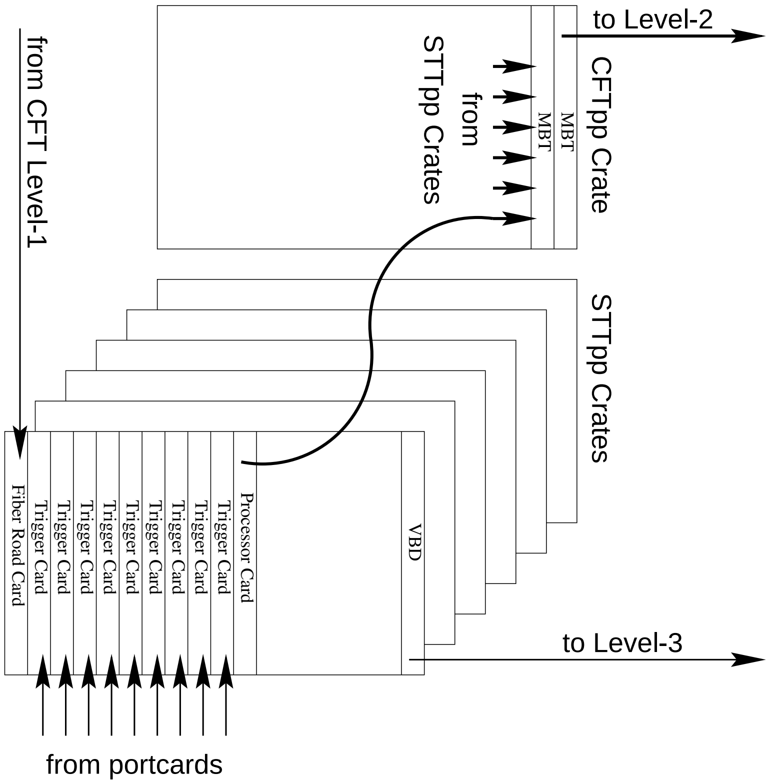

Figure 6: Crate structure of STTpp.

2.2 Fiber Road Card

2.2.1 Overview

The fiber road card serves three functions in the overall STTpp design.

- It receives roads from the CFT Level-1 trigger.
- It sends the roads out over a special purpose bus to the trigger cards.
- It is used as the controller and buffer manager for the STTpp crate.

One of these cards will be necessary for each STTpp VME crate, corresponding to one STTpp sector. A block diagram of the fiber road card is given in Figure 7.

2.2.2 Input from CFT Level-1 Trigger

The CFT Level-1 trigger transmits up to 32 CFT track candidates over a single fiberoptic cable to each STTpp sector 60° wide. For each track candidate it sends the identities of the fibers hit in the inner and outermost CFT layers. Tracks that pass through a given STTpp sector can hit any one of 773 fibers in the outermost CFT layer, spanning a region of 60°+19° in azimuth [9]. Tracks with $p_T > 1.5$ GeV that pass through a given outer fiber can hit any one of 28 fibers in the inner CFT layer. We thus need at least 15 bits to completely specify a CFT track candidate, 10 to identify the outer fiber and another 5 for the inner fiber.

The exact scheme to map the 80 azimuthal CFT sectors onto the six 60° STTpp sectors, that looses as few as possible tracks which cross sector boundaries, is currently under study.

The CFT Level-1 track lists will be obtained by passively splitting the fibers that transmit the CFT track lists from the CFT Level-1 concentrator cards to the CFTpp. The track lists will have one 16-bit word for each track.

A fiber optic receiver will be necessary to convert the CFT data to electrical signals. The data will then be put into a FIFO which decouples the data-sending clock from the local crate clock.

2.2.3 Buffer and Crate Control Functions

In addition to dealing with the standard DØ crate controls from the trigger framework, the fiber road card must also control buffering and readout of the data from the processor cards to Level-3.

2.3 Trigger Card

2.3.1 Overview

The trigger card has to fulfill the following tasks:

- receive CFT data from the fiber road card over the backplane and convert it to roads in SMT,
- receive SMT data from the portcards over optical fibers and find hit clusters,
- select SMT hit clusters that are inside the roads defined by CFT data,
- buffer SMT data,
- transmit the hit clusters in the roads to the processor card(s),
- buffer hit clusters for readout to Level-3 over VME bus.

Figure 8 shows a block diagram of the trigger card.

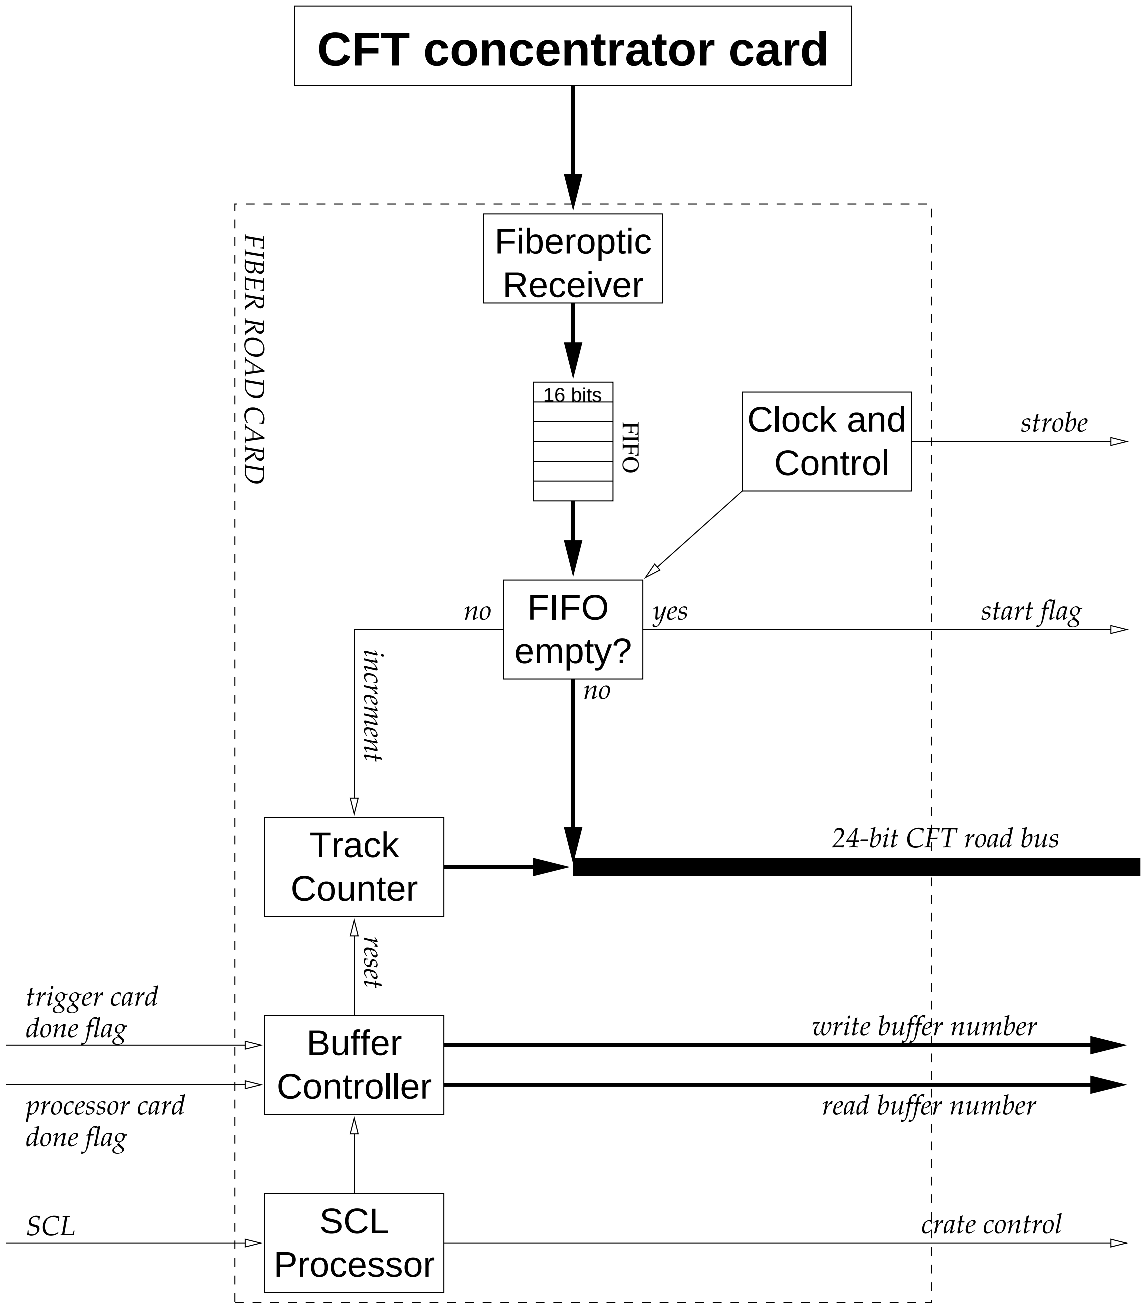

Figure 7: Block diagram of fiber road card.

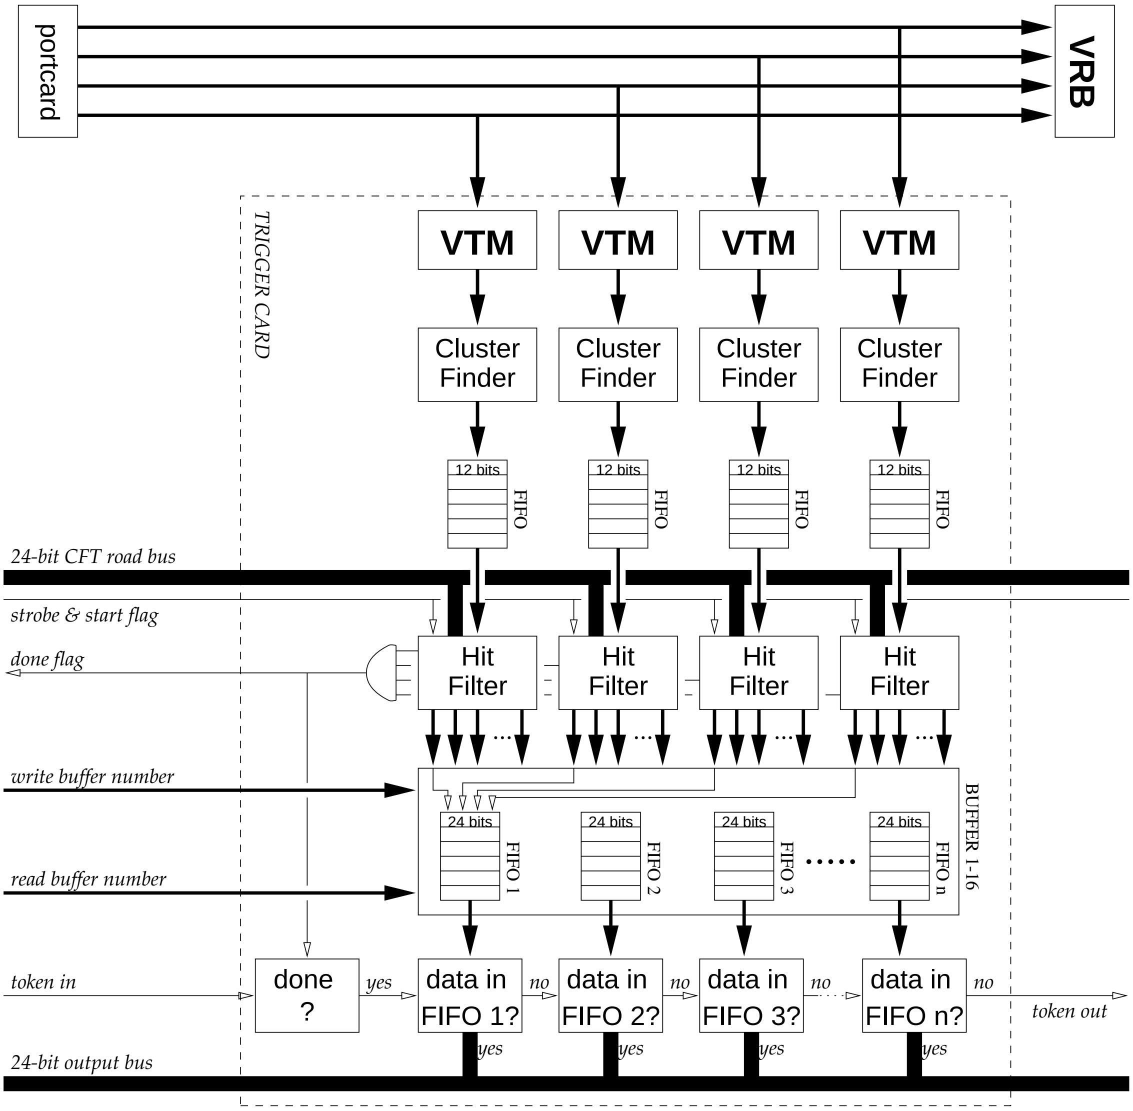

Figure 8: Block diagram of trigger card.

2.3.2 Input from SMT Front-Ends

Optical fibers transmit the SMT data from the portcards on the detector platform to the VRBs [10] in the counting house. The data transmitted over a fiber consist of alternating 8-bit address and signal words. These words are transmitted at a rate of 53 MHz. On every fiber an end-of-record code indicates that all data for an event have been sent. Passive splitters in the fibers create a data path into the STTpp, where the fibers plug into the same VTM receiver module [11] as used for the VRBs. The data are loaded into a FIFO as they are received.

Each barrel detector is read out over a single HDI. Each portcard accepts data from eight HDIs and outputs it over four optical fibers such that two HDIs map into one fiber. The four fibers from one portcard plug into the same VTM at the VRB end and at the STTpp end. The mapping of the barrel detectors into portcards must therefore arrange the SMT data in the way that the STTpp requires it.

The STTpp has to provide seamless coverage of all barrel segments that belong to one SMT sector, because many tracks cross from one barrel segment into another. Very few tracks are not contained in a single SMT sector (section ??) so that each sector can be treated as an independent unit. This means that the data from all barrel segments of any one SMT sector must be routed to the same STTpp crate. For a six-barrel configuration there are 36 detectors per sector which populate 4.5 trigger cards. We can therefore accomodate two SMT sectors in each STTpp crate.

2.3.3 Hit Finding

An FPGA processes the data further. It corrects for pedestals and gains and supresses noisy strips based on a lookup table. It supresses all strips with pulseheights below a threshold. It has to be determined whether these operations need to be performed on a strip-to-strip, chip-to-chip, or global scale. The FPGA then scans the corrected pulseheights for local maxima, *i.e.* strips with larger pulseheights than their two neighbors. Each such local maximum forms the seed for a hit cluster. The position of the hit cluster is determined from a pulseheight averaged center of gravity of up to five adjacent strips, centered on a seed strip [?]. The cluster outputs 13 bits: the seven bits used to identify the strip number plus an additional two added by the cluster finding algorithm, three bits to label the eight chips that read out the axial strips or the seven chips that read out the stereo strips, and one bit to indicate whether the cluster is on the axial or the stereo side (see figure 9).

13	12 11 10	9 8 7 6 5 4 3 2 1
axial stereo	chip	strip

Figure 9: Bit allocation in output of cluster finder.

Based on bit 13 the clusters are stored in two FIFOs. Only clusters from the axial strips participate in the track finding. The axial strips are read out by the first (last?) three chips for layers 1,2,5, and 6, and by the first (last?) five chips for layers 3,4,7, and 8. The clusters from the axial strips are stored in one FIFO, the clusters from the stereo strips are stored in the other FIFO. All clusters are buffered for readout to Level-3 over the VME bus (?).

2.3.4 Input from Fiber Road Card

The data from the CFT Level-1 trigger are transmitted from the fiber road card over a bus on the backplane. For every CFT track 24 bits have to be transmitted: 16 bits to identify the inner- and outermost hit fibers and an 8-bit track number. The bus requires 24 signal lines and one strobe line. It runs at 25 ns per cycle. According to Monte Carlo simulations [12], we expect on average less than two tracks per STTpp crate. The maximum number of tracks expected is 25, the maximum number allowed per STTpp crate is 32. It therefore takes 50–800 ns to transmit the CFT data across the backplane to the trigger cards.

2.3.5 Conversion of CFT Tracks to SMT Roads

The 16 lines that identify the CFT track address a lookup memory. The output of the lookup memory is a range of strip numbers, consisting of 10 bits each (corresponding to bits 3–12 in figure 9) for the first and last strips. If a track does not go through a detector the corresponding lookup memory will return zero for both strip numbers. Since every input fiber carries the data from two barrel detectors, up to eight strip number ranges have to be determined for every CFT track. Tracks that go through an SMT sector of 60° in azimuth can hit any one of 773 outer fibers. For each outer fiber there are 28 inner fibers that can be hit. So there are 21644 possible CFT tracks for a 60° SMT sector. We therefore need either one lookup memory of 433 kB with 16 address lines and 160 output lines or eight lookup memories of 54 kB each with 16 address lines and 20 output lines. The coordinate conversion can then be performed "on the fly" just like the cluster finding and no time is lost.

2.3.6 Road Definition

The fiber hits in the inner and outer CFT layers define a family of tracks, which all go through these fibers. The envelope of all these tracks defines the road in the SMT. For radial positions smaller than that of the inner hit fiber, the tracks with the smallest and largest curvature define this envelope. Figure 10 shows the road in the SMT spanned by two hit fibers in the CFT. The solid lines indicate the road for tracks that pass through the origin. If we also want to accept all hits from tracks with finite impact parameters b the road widens as indicated by the dashed lines.

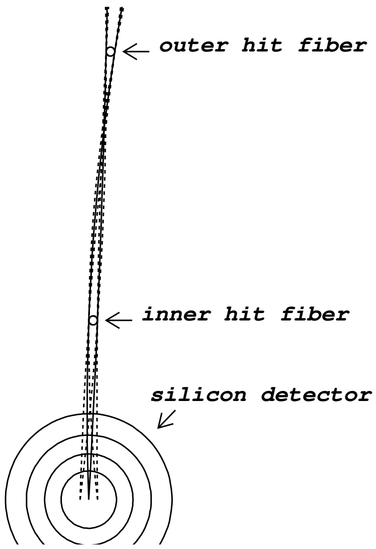

Figure 10: Road definition in the silicon detector.

Figure 11 shows the width of the road in azimuth around the track, defined by the centers of the inner and outer fibers and the origin, versus radius. The black region is the road defined by the envelope of all tracks passing through the origin, the dark shaded region is the road defined by tracks with impact parameter $|b| < 1$ mm, and the light shaded region for $|b| < 2$ mm. Figure 12 shows the road widths in arclength. Table 1 lists the road widths at the mean radii of the silicon layers. In the table "a" and "b" refer to the two overlapping sublayers that make up one silicon layer.

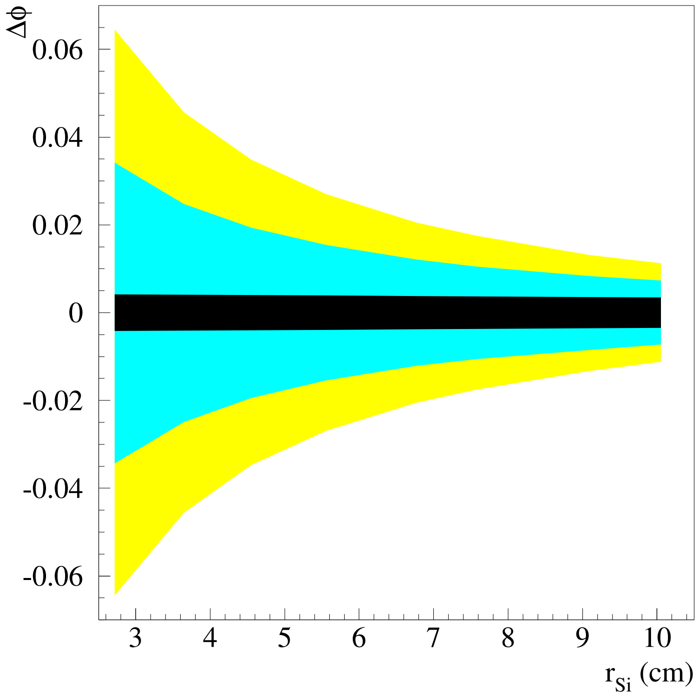

Figure 11: Width of silicon roads versus radius.

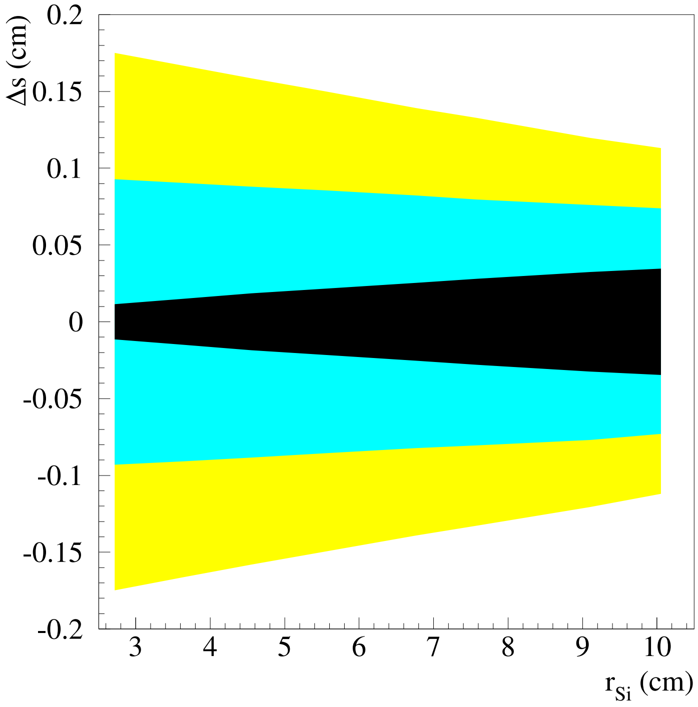

Figure 12: Width of silicon roads versus radius

Table 1: Road sizes in silicon detector.

layer	radius	road for $b = 0$		road for $\|b\| < 1$ mm		road for $\|b\| < 2$ mm	
		mrad	mm	mrad	mm	mrad	mm
1a	2.72 cm	$\pm$ 4.2	$\pm$0.12	$\pm$34.3	$\pm$0.93	$\pm$64.4	$\pm$1.75
1b	3.65 cm	$\pm$ 4.1	$\pm$0.15	$\pm$24.8	$\pm$0.91	$\pm$45.7	$\pm$1.66
2a	4.55 cm	$\pm$ 4.0	$\pm$0.18	$\pm$19.4	$\pm$0.88	$\pm$34.8	$\pm$1.58
2b	5.55 cm	$\pm$ 3.9	$\pm$0.22	$\pm$15.4	$\pm$0.86	$\pm$27.0	$\pm$1.50
3a	6.77 cm	$\pm$ 3.8	$\pm$0.26	$\pm$12.1	$\pm$0.82	$\pm$20.6	$\pm$1.39
3b	7.58 cm	$\pm$ 3.7	$\pm$0.28	$\pm$10.5	$\pm$0.80	$\pm$17.5	$\pm$1.33
4a	9.10 cm	$\pm$ 3.6	$\pm$0.32	$\pm$ 8.4	$\pm$0.76	$\pm$13.2	$\pm$1.20
4b	10.05 cm	$\pm$ 3.4	$\pm$0.35	$\pm$ 7.3	$\pm$0.73	$\pm$11.2	$\pm$1.13

2.3.7 Hit Filter

As the cluster data appear at the output of the cluster-position FIFO, they are processed by a second FPGA. This FPGA simultaneously compares the cluster positions to the upper and lower limits of the all address ranges for all roads. Bits 3–12 in figure 9 are compared to the road limits. Since there can be up to 32 roads and two address ranges per road per fiber, a total of 128 10-bit comparators are required. Figure 13 shows a block diagram of the hit filter FPGA.

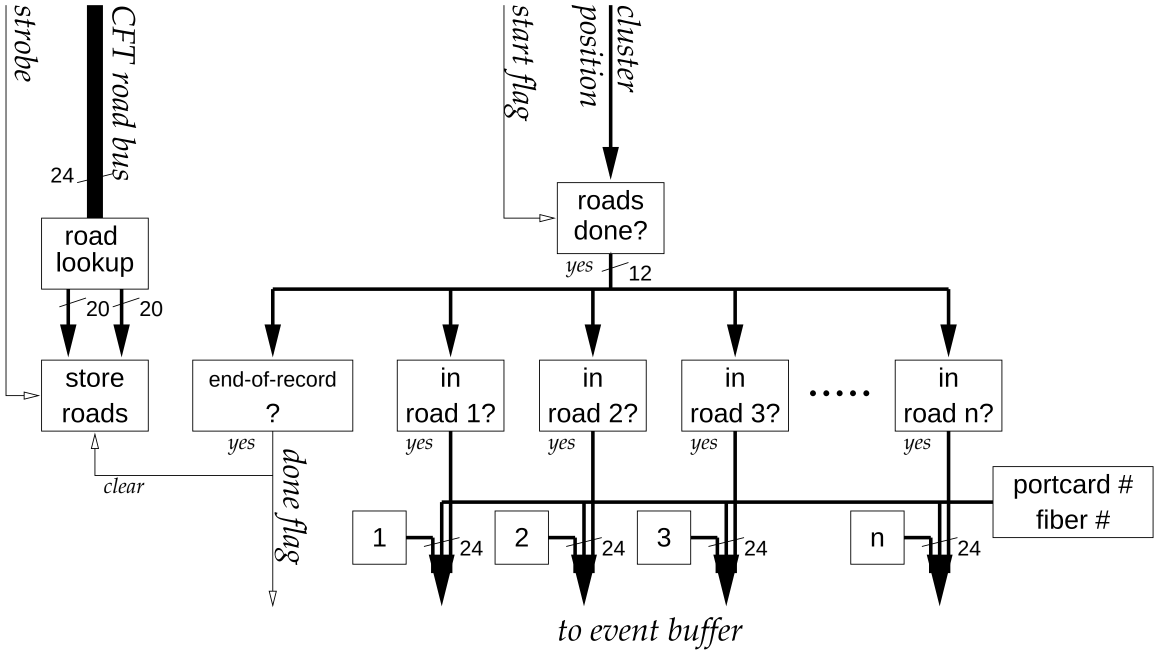

Figure 13: Block diagram of hit filter FPGA.

2.3.8 Event Buffering

The output of the hit finder FPGA is stored in a buffer 16 events deep. There is a FIFO for every road that stores all (12-bit) cluster positions that are inside an address range for this road. An end-of-record code is passed through the cluster finder unaltered. When it appears at the hit filter, it sets a "done" flag to indicate that all data from the fiber has been processed through the cluster finder and the hit filter and has been loaded into a buffer. The logical "and" of the flags from the four hit filters on a trigger card indicates that the trigger card has finished loading an event. The logical "and" of the "done" flags from all trigger cards is formed on the backplane (?) and signals to the fiber road card that an event has been completely loaded into the buffer. The fiber road card then updates the buffer pointer for the next event to point to an empty buffer. At this time the STTpp is ready to accept another event.

2.3.9 Timing

It takes about 2.4 μs after a Level-1 trigger accept for the CFT data to arrive at the fiber road cards. Bus transfer to the trigger cards takes another 0.8 μs. That means 3.2 μs after a Level-1 trigger accept the road addresses are latched on the trigger cards.

Digitizing the SVX-II data takes a little under 3 μs and there is an additional small latency of about 100 ns from the cluster finding. Thus the SMT data arrive in the FIFO at the output of the cluster finder roughly at the time the roads are set up.

Every fiber carries data from about 2000 SMT channels. Monte Carlo simulations predict an average of 40 hits per fiber, corresponding to an occupancy of 2% [12] and a maximum of 180 hits per fiber. We add 1% occupancy from noise (see analysis of test beam data [13]). We thus have 60–200 channels to read out. At 53 MHz this will take 2.3–7.6 μs. Only half the hits (from the axial strips) form clusters and the reduction from number of hits to number of clusters should be about a factor two. Thus there will be 15–50 clusters which have to be compared to the strip ranges that define the roads. The comparisons for all roads are performed in parallel. For each road there is a FIFO and a cluster that is in the strip range for a road is stored in the corresponding FIFO. This should not take longer than 100 ns (?) per cluster, *i.e.* 1.5–5.0 μs. This is less time than required for the readout of the SMT. The cluster filtering can therefore effectively be performed without latency.

All hit clusters in the roads for an event are loaded into a buffer about 5. 5–10.8 μs after the Level-1 trigger accepted the event, but not later than completion of the readout of the SMT data. At this time the STTpp is ready to accept another event from Level-1.

2.3.10 Transfer of Hits to Processor Card

For each hit 26 bits need to be transferred: 7+2 bits for the fractional strip number, 3 bits to identify the HDI/chip within a fiber, 2 bits to specify the fiber within a portcard, 4 bits to identify the portcard, and a 8-bit track number (figure 14). The lower 12 bits are the same as at the output of the cluster finder (figure 9).

26 ... 19	18 ... 15	14 13	12 ... 10	9 1
track	portcard	fiber	HDI/chip	strip

Figure 14: Bit allocation in output to processor card.

Hit transfer to the processor card is initiated by the presence of data in the event buffer. The first trigger card transmits the lists for all roads over a 26-bit bus to a processor card. For each road the first record consists of the CFT data and the road number. These are the 24 bits sent from the fiber road card to the trigger cards. Then all hits in the road are transferred. This sequence is repeated for all roads. When the card has finished transmitting the list it passes a token to the next trigger card and so on until all trigger cards that serve the same azimuthal sector in the SMT have transmitted their data to the processor card.

There are on average 42 and at most about 500 clusters in all roads of an STTpp crate [12]. If the bus operates at 25 ns/cycle, the transmission of these clusters and the CFT data to the processor card will take 1.1–13.3 μs.

2.4 Track Fit Card

2.4.1 Hit Conversion and Track Fitting

The match between CFT trigger tracks and clustered silicon information is made in the STC. The next stage of processing – silicon coordinate conversion and track fitting – is described in this section. It contains an overview of the fitting system followed by a description of each of the three functional blocks.

Figure ?? shows a block diagram of the track–fitting card. There are three basic functional blocks. They are:

- an **input stage** which receives combined SVT cluster and CFT track information and provides conversion from cluster strip coordinates to global coordinates needed for track fitting,
- a **fitting stage** consisting of an array of dedicated processors to perform the track fit dedicating one processor to each track, and
- an **output stage** which communicates the fit results to the rest of the trigger system.

Because many details remain to be optimized, the design presented here should be seen as an initial, but workable, solution. The most important pieces of the track fitting process have been considered, and the resulting design is believed to be sufficient. In particular, the processor design presented here is based on Texas Instruments TI320C62 digitial signal processors(DSP). These may be replaced in the final design by programmable logic devices, most probably the Altera 10k100 device. However, this would be driven by price and performance considerations, not, we believe, by problems intrinsic to the DSP–based design presented here.

2.4.2 Input Stage

The input stage receives the matched SVT cluster and CFT track information from the silicon trigger card over a dedicated bus. The bus provides a uni–directional, data–driven broadcast from the trigger cards to the track fitting card. When the cluster finding and road matching in the STC is complete, each cluster and the ID of the associated CFT track are driven onto the bus. Each transfer is 30 bits wide, and the cycle time is 25 ns. The cluster information includes the barrel number (3 bits), ladder number (7 bits) and cluster position in strip coordinates (12 bits)[1]. The track ID is taken to be a global serial number uniquely identifying each track for a given event. For a system capable of handling 200 tracks, eight bits are needed for the track ID. In addition to the cluster information, start– and end–of–data transfers will provide synchronization and error checking, and CFT track parameters will be sent as needed by the fitting procedure.

The input stage consists of bus receivers, simple control logic, a table to provide track ID to processor ID translation, and a look up table to provide conversion from hits positions in ladder coordinates (barrel, HDI, SVX chip ID and centroid position) to positions in global (r, ϕ) coordinates.

There will be one processor for each track to be fit, with an upper bound of 32 tracks for each 60^o sector. Each silicon hit will be broadcast with a global track ID indicating which CFT track the hit is associated to. The track ID is a global ID, and a look up table is used to associate a given track to the appropriate processor. A global track ID is input to the look up table, and the ouput is the ID of the processor handling that track. The table is reset at the start of each event and built on–the–fly as data for each track is received. The needs are very modest, 32 8-bit comparators, 32 8-bit registers and minimal control logic.

In addition to the look up table for track ID to processor ID conversion, a second table is used to convert the cluster into global coordinates for the track fit. We have adopted the simplest possible approach to perform this conversion. A look up table is built containing one entry for every possible cluster position in the geometrical sector. This implies between 128k and 256k entries in the look up table for one 30^o sector. The address is formed from the barrel, HDI and SVX chip ID's and the cluster position in strip coordinates. Each entry contains the aligned hit position in (radius,phi) space. If the radius is stored in units of 10 microns and the phi angle granularity is derived from the smallest angular resolution resulting from the 1/4

[1] It is likely that the cluster information will be in the format channel number, SVX chip ID and HDI side. In this case the number of bits can be slightly reduced.

strip cluster precision, the table needs at most 32 bits for each slot. This corresponds to 14 bits for radius and 18 bits for the angle.[2]

The cost of such a table implemented in 15 ns static RAM will be less than $70 per 30^o sector. Processor clock cycles are typically 5 ns, considerably faster than the external bus cycle, so converted hits can be copied directly into processor memory.[3]

Although the coordinate conversion table is quite large, the download is not a significant problem because it changes only when the detector alignment changes. Battery backed up RAM could be used to eliminate the need to download after power failures.

2.4.3 Processing Stage

During the track/hit loading phase, the input stage is used to fill processor internal memories with the initial track parameters and positions of associated silicon hits. When this is completed the processors are started, and the tracks are fit using the silicon information. The fit algorithm studied thus far is a linearized fit to the function

$$\phi(r) = b/r + \kappa r + \phi_0$$

This is the same algorithm used in the physics studies. The results from the fit are the three track parameters b, ϕ_0 and curvature and the χ^2/dof. Here b is the vertex-signed impact parameter and ϕ_0 is the angle of the track at the point of closest approach. The algorithm used in most of the physics studies constrains the hits such that the barrels must be in a mechanically sensible order, and if the χ^2 is too large, the track is refit, dropping the hit dominating the χ^2.

The processors will be commercial units, either digital signal processors or programmable logic. The simplest approach would likely be the use of DSP's, but the final choice will be driven by overall cost and performance. At this point in the design, it is important to ensure that the fit can complete within the overall L2 timing constraints. We are conservatively allowing 10 μs for the fit. The time taken has been estimated for one possible processor, the Texas Instruments TMS320C62, by counting instructions needed by the fit. Simulations are underway using silicon Monte Carlo and DSP compiler/simulation packages from TI.

To perform the timing estimates, the fortran code used in the physics studies was reordered to give an apparently minimal number of operations. The number of operations was tallied under the assumption that in DSP processors, additions following multiplications take the same amount of time as multiplications. It was further assumed that the initial track parameters were available in memory, and that hit resolutions were also available as look up. Under these assumptions, the number of operations per track N was found to be

$$N = 27n \text{ multiplies} + (3+n) \text{ divides}$$

in which n is the number of hits on the track.

A simple estimate of the dynamic range needed during the fit gives 35 integer bits, beyond the 32 bit integer capacity of the processor. This implies that floating point formats will be needed[4]. In estimating the fit times, it is assumed that the emulation of 32-bit floating point operations in 32-bit integer format requires 10 instructions for each floating point operation. Roughly, this corresponds to two loads, two shifts, an integer multiply, an integer sum, a shift and a store. This gives eight operations, which is taken as 10 here.

The DSP under consideration is rated as a 1600 MIPS processor, but local tests indicate the processing capacity is 400 MIPS. The external clock frequency is 200 Mhz. The processing time is given as

$$\begin{aligned} T &= (\mathrm{Numberofinstructions})(\mathrm{timeperinstruction})(\mathrm{emulationfactor}) \\ &= (27n + 3 + n) * 2.5 \text{ ns} * 10 \end{aligned}$$

[2] It will be important to insure that the mapping of discrete phi values in the table meshes with the discrete precision of the input clusters. One way to avoid possible mismatches would be to store a ladder ϕ and a cluster $\delta\phi$ in the look up table. The true position would be the sum of the two numbers. This, however, does introduce the need for an addition operation in the conversion.

[3] If synchronizing the loading of the table outputs into the assigned processor is difficult to complete within the bus cycle, the output could be latched for a single cycle. The processor load would then occur during the following cycle. This is essentially creates a two-stage pipeline, and has no real impact on performance.

[4] Other possibilities such as dual range calculations will be studied. If possible, integer execution will be used.

For a fit using four silicon hits, the corresponding processing time is 2.8 μs. Thus, it appears that single–pass fitting can be completed well within the alloted 10 μs. Two–pass fitting in which the worst hit is discarding for tracks with a large χ^2 is also apparently feasible but given the precision of this result two–pass fitting clearly needs more study. Furthermore, this result assumes that only one hit per layer is available. In practice this filtering is best performed after hit conversion because misalignment effects can be considered. The feasibility of performing both hit filtering and a multiple pass fit remains an open question.

2.4.4 Output Stage

The ouput stage is used to transfer the fit results to the global DAQ. The transmitted information consists of the three track parameters and the χ^2 resulting from the fit. In addition, one may want to read out the hit positions for a subset of events. The output will either be to a dedicated local card via a custom bus or directly to a concentrator crate via a serial link.

An upper bound on the data transfer can be derived under the assumption of 32 bits of precision for each of the four quantities from the track fit. At 50 tracks per event, and 16 bytes per track, 800 bytes would be transferred. However, the output can probably be packed to reduce size as follows:

- **Impact Parameter** storing with 1 μm precision and a range from -1.0 cm $\leq b \geq$ 1.0 cm requires 15 bits,
- Momentum storing with 0.1 GeV precision and a range from -100.0 to 100.0 GeV requires 11 bits,
- ϕ_0 storing with 0.025 rad precision requires 8 bits
- χ^2/dof storing with a precision of 0.2 and a range from 0 to 10.0 requires 6 bits.

For a 50 track event, this results in 250 bytes transferred.

2.5 z-Vertex Processor

2.5.1 Primary Vertex Determination

We present here the possibility to measure the z position of the primary event vertex using hit information from the SMT at Level-2. The four inner barrel sections of the SMT, which cover the range $-25 \leq z \leq 25$ cm around the center of the detector, are equipped with four layers of double-sided detector elements. Layers 1 and 3 have strips that are parallel to the beam direction on one side and strips that form a 90° angle with the beam direction on the other side. The hit information from the 90° strips only can be used for a fast determination of the z position of the primary event vertex because these strips are projective in the r-z plane and no hit information from the axial strips is necessary. Using these detector elements we can measure the z vertex position in the range $-30 \leq z \leq 30$ cm with a resolution of 525 μm.

Knowledge of the z position of the primary event vertex would allow a more precise measurement of transverse momentum (p_T) at Level-3 from the calorimeter, which measures energy and needs the z vertex position to convert this into a momentum vector. A better transverse momentum resolution translates into reduced trigger rates because the background p_T spectra, which dominate the trigger rate, typically fall rapidly with increasing p_T and because a sharper threshold allows higher threshold values without loss of efficiency. Moreover the knowledge of the vertex multiplicity can be used to veto on events with many interactions, as an intelligent prescale for high rate processes.

In the following we describe the algorithm used, and present a conceptual design of a possible implementation of the algorithm in the trigger hardware. A simulation of the performance and physics benefits from this device is described in reference [15].

2.5.2 Algorithm

Given the hits in layers 1 and 3 of the SMT we first considered connecting all hits in layer 1 with all hits in layer 3 and histogramming the intercepts of all these lines with the beam axis. Peaks in this histogram would indicate event vertices. However the number of combinations may be prohibitively high (over 4000

Table 2: Occupancies (in clustered hits)

	Number of hits in 30° ϕ sectors		
	1 interaction	3 interactions	6 interactions
hline layer 1	6	13	26
layer 3	4	10	17

for events with two additional interactions) and therefore such an algorithm, which requires many floating point operations, would far exceed the available processing time at Level-2.

We therefore designed an algorithm which is in principle equivalent to the above but utilizes massive parallelism to compute the results within the prescribed time-frame. The loop over the hit combinations is replaced by a loop over candidate vertex locations. The algorithm sequentially assumes discrete values of the vertex z position and for each value it counts the number of hit pairs in layers 1 and 3 that line up. This can be easily achieved by scaling the z positions of all hits in the first layer as follows,

$$z'_{1,i} = \frac{R_3}{R_1} z_{1,i} + \left(1 - \frac{R_3}{R_1}\right) z \tag{1}$$

where $z_{1,i}$ is the i^{th} hit in layer 1, z is the assumed vertex position and R_j is the radial position of layer j. If multiple coulomb scattering is negligible, then $z_{3,i} = z'_{1,i}$ for tracks that point to the assumed z vertex position. Therefore, we can simply bin the hit positions in layer 3 and the scaled hit positions in layer 1 with a resolution that is similar to or coarser than the hit resolution of the silicon devices in z direction (about 30 μm) and count how many bins have entries both from layer 1 and layer 3. Equation 1 can be evaluated for $z = 0$ and histograms for other values of z can be obtained simply by shifting this histogram.

Thus, this algorithm has the advantage that the computation is independent of the hit multiplicity in the event and it only involves one multiplication per hit. Figure 15 illustrates this concept.

A simple peak finding algorithm is used to determine the vertex location. The peak finder scans the track multiplicity for each z-position and identifies the one with largest track multiplicity as the event vertex. To compute the z-vertex a weighted average of the two neighbouring bins is used. The peak finder sets these bins to zero and goes on to find the next highest bin and so on to determine the location of other vertices in the event until the track multiplicity of the highest bin falls below a threshold. This threshold is set to 5 times the average track multiplicity of the event.

2.5.3 Implementation

We implemented the above algorithm using 1 mm bin sizes for layer 3. We then scaled the hits in layer 1 to the average radius of layer 3 at 7.5 cm. As the even and odd ϕ ladders in layer 3 are at different radii, for simplicity we also scaled the layer 3 hits to 7.5 cm, before binning them into 1 mm bins. A finer granularity (about a factor of 5–10 more) could be used, but is more expensive and time consuming. The granularity of 1 mm gives us acceptable vertex resolution, as will be shown below, and is cost effective. The step size chosen for the vertex scan is also set to 1mm.

We have studied the clustered hit occupancy in these ladders. The hit occupancy (without multiplexing in the 90° ladders and 156 μm strip pitch) in 30° ϕ sectors in layers 1 and 3 for events with one, three, and six interactions is given in table 2.

2.5.4 Design of the z-vertex finder

For the two-layer system in D0, the tracking problem can be reduced to a simple coincidence logic as described in section 1. The hits in the outer layer are projected to a common radius of 7.5 cm and binned. An adequate bin size, determined from simulation, is 1 mm but the design is flexible enough to accommodate any optimized value. The theoretical lower limit of 35 microns is probably excessively precise but an intermediate value may prove to be more cost-effective.

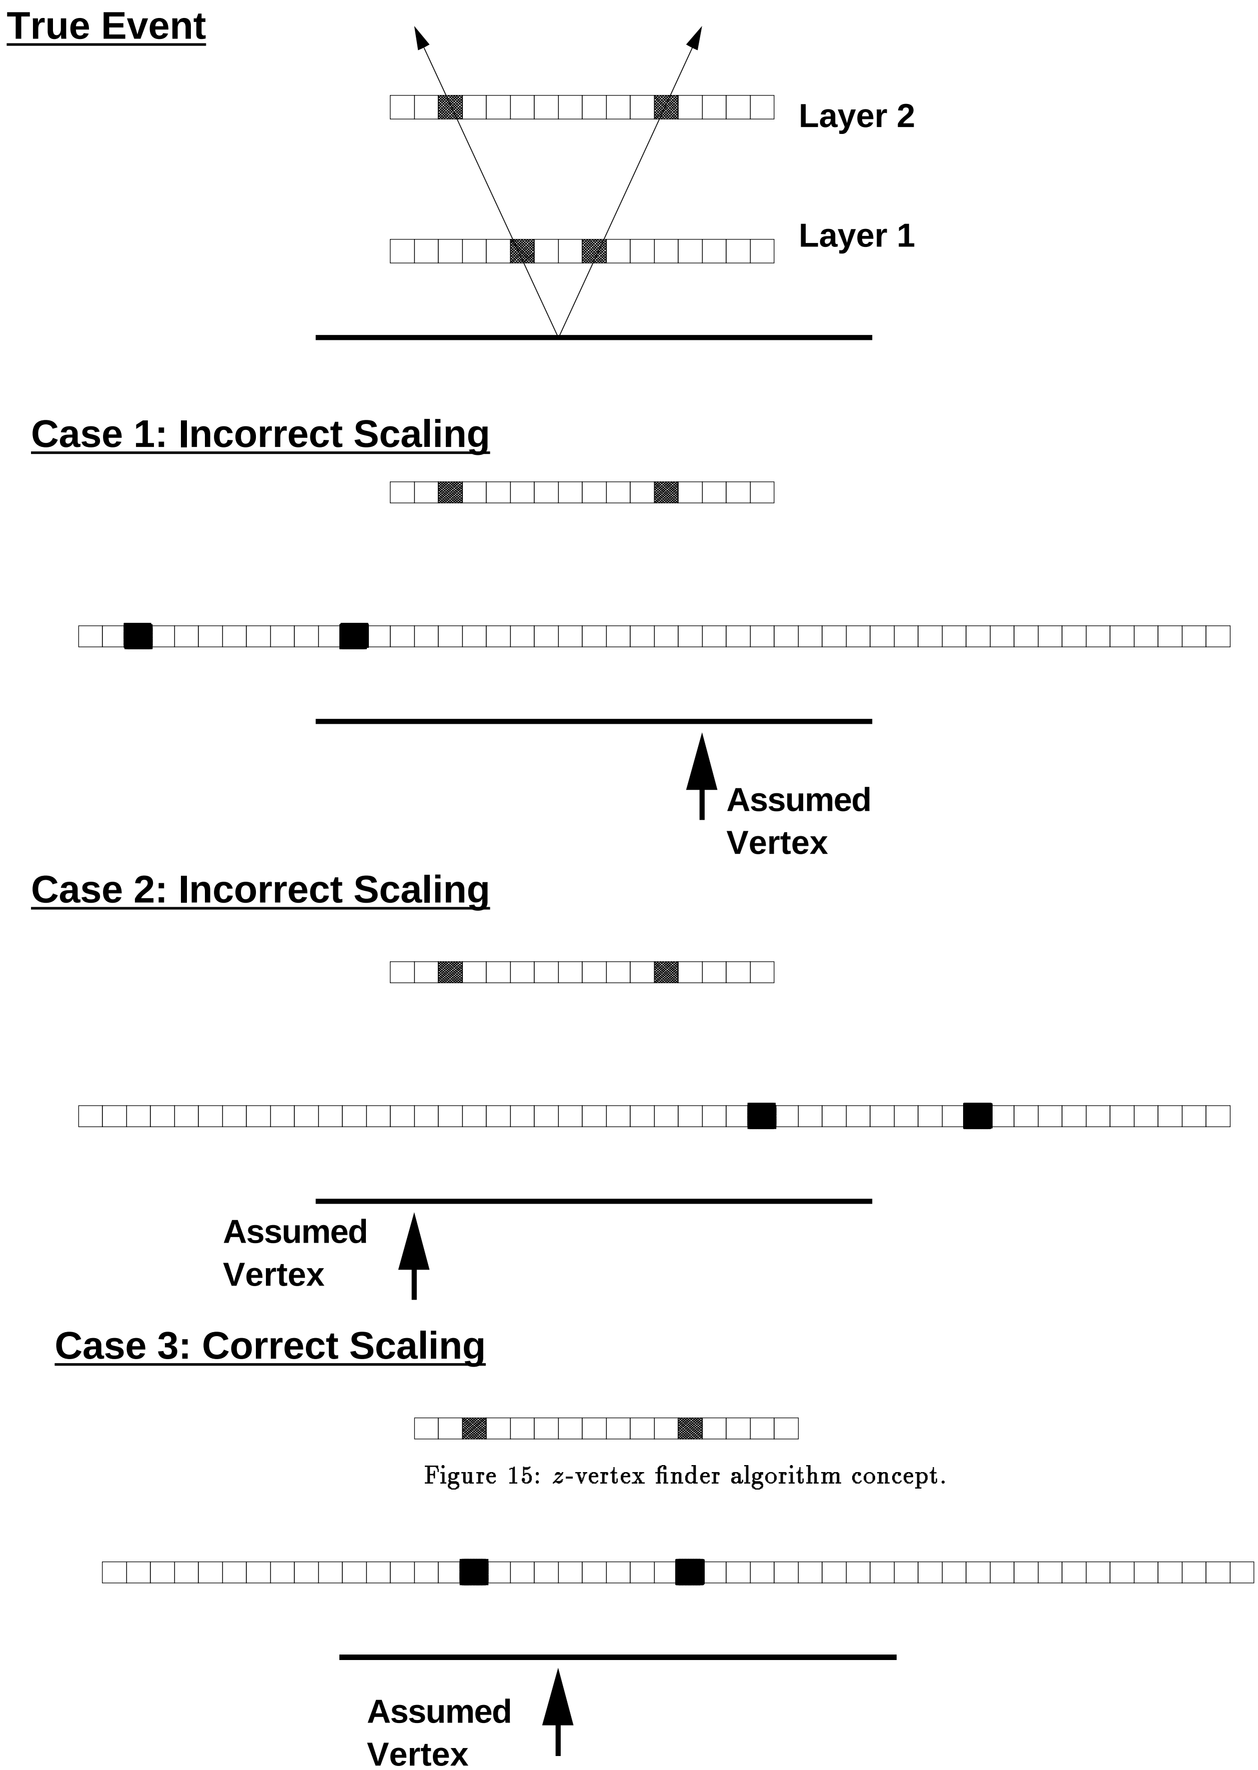

Figure 15: z-vertex finder algorithm concept.

The region of interest along the beamline is divided into virtual bins, with appropriate bin sizes. Again, a value of 1 mm has been used in this exercise. Now, the hits in the inner layer are scaled as described in equation 1. It can be seen that for the correct value of z, an equality $z'_{1,i} = z_{3,i}$ is formed. By counting such coincidences, the vertex can be located above the background from random coincidences. This transformation consists of a simple scaling of all hits (first term) followed by consecutive shifts by a constant number of bins (for z changing in integral amounts).

The SMT geometry naturally divides itself into 12 overlapping sectors. The z vertex finder takes full advantage of this and performs its computations on 12 cards in parallel. One such sector is shown in figure 16.

The first task is to form centroids using raw data. This function is carried out on the trigger cards. The output of the cluster finder is buffered internally in a DSP, which also scales the hits for layer 1 by the factor appropriate for this sector. The scaling factor will be determined from off-line alignment programs after detector turn-on. The DSP should be able to keep up with the output of the cluster finder. The output of the DSP is a 10-bit number for layer 1 but only 8-bits for layer 2. These hits are decoded and the corresponding bins are latched in large arrays as shown.

So far, all of the functions are carried out in a pipeline and hence the estimated elapsed time is about 10 μs after a Level-1 accept. Now, the hits from the two layers are presented to a parallel coincidence and adder circuit. We estimate that the task requires four FPGAs. The total input pin count is 1024 (256×2 for layer 3 and 256×2 for layer 1). Each FPGA will use 256 pins for the input and 6–8 pins for the output.

The gate count for the coincidence (including neighbor logic) is 2048 and about 4000 for the parallel adder. This will not only fit quite easily into various commercial parts but given the small gate count, it will be able to perform this function in one clock cycle. The outputs of all four FPGAs is added by a simple adder and loaded into a histogram memory. These last two blocks can also be implemented in another FPGA.

Next, the hits in layer 1 are shifted by 3 bins (corresponding to the scale factor) and again fed into the coincidence FPGAs. For a total of 512 z bins, this stage requires 512 clock ticks to fully load the histogram memory. Assuming a 53 MHz clock, this takes about 10 μs. Hence, the total time spent since the Level-1 accept is now 20 μs.

The histograms from all 12 sectors have to be now added. A simple check of bins above some threshold count can be used to define vertices. However, with the remaining approximately 30 μs of the assigned Level-2 preprocessor time, a more sophisticated peak finding algorithm can also be implemented.

2.6 Communication with Global Level-2 Processor

2.6.1 Track Preprocessor Crate Overview

The results of the track fits performed on the processor card will be sent to MagicBus Transciever cards (MBT) [14] in the CFTpp crate using the Cypress Hot Link serial output bus. One hotlink is envisaged per processor card. With one processor card per crate, the output of the entire STTpp should fit into one MBT in the CFTpp crate.

The exact layout of the Level-2 CFTpp crate is still being decided. However, it is forseen to look similar to the standard DØ Level 2 preprocessor crate. Its main working elements will be

CFT MBTs Receive CFT Level-1 track lists from CFT concentrator crate. Send this information out onto MagicBus.

STT MBT Receive STTpp data from processor cards. Send this information out onto MagicBus.

Alpha Processors Process CFT and SMT information. Transmit results over MagicBus to Level-2 MBTs.

Level-2 MBTs Send processed Level-2 track trigger information to the global Level-2 processor via Cypress Hot Link.

2.6.2 Input from Processor Cards

Data from the processor cards will contain the information shown in table 3. The exact size of some of the data quantities will depend on the algorithm used for the track fit, and is thus not yet defined. However, each should fit in a 16-bit word. Monte Carlo studies [12] indicate that the average number of tracks with

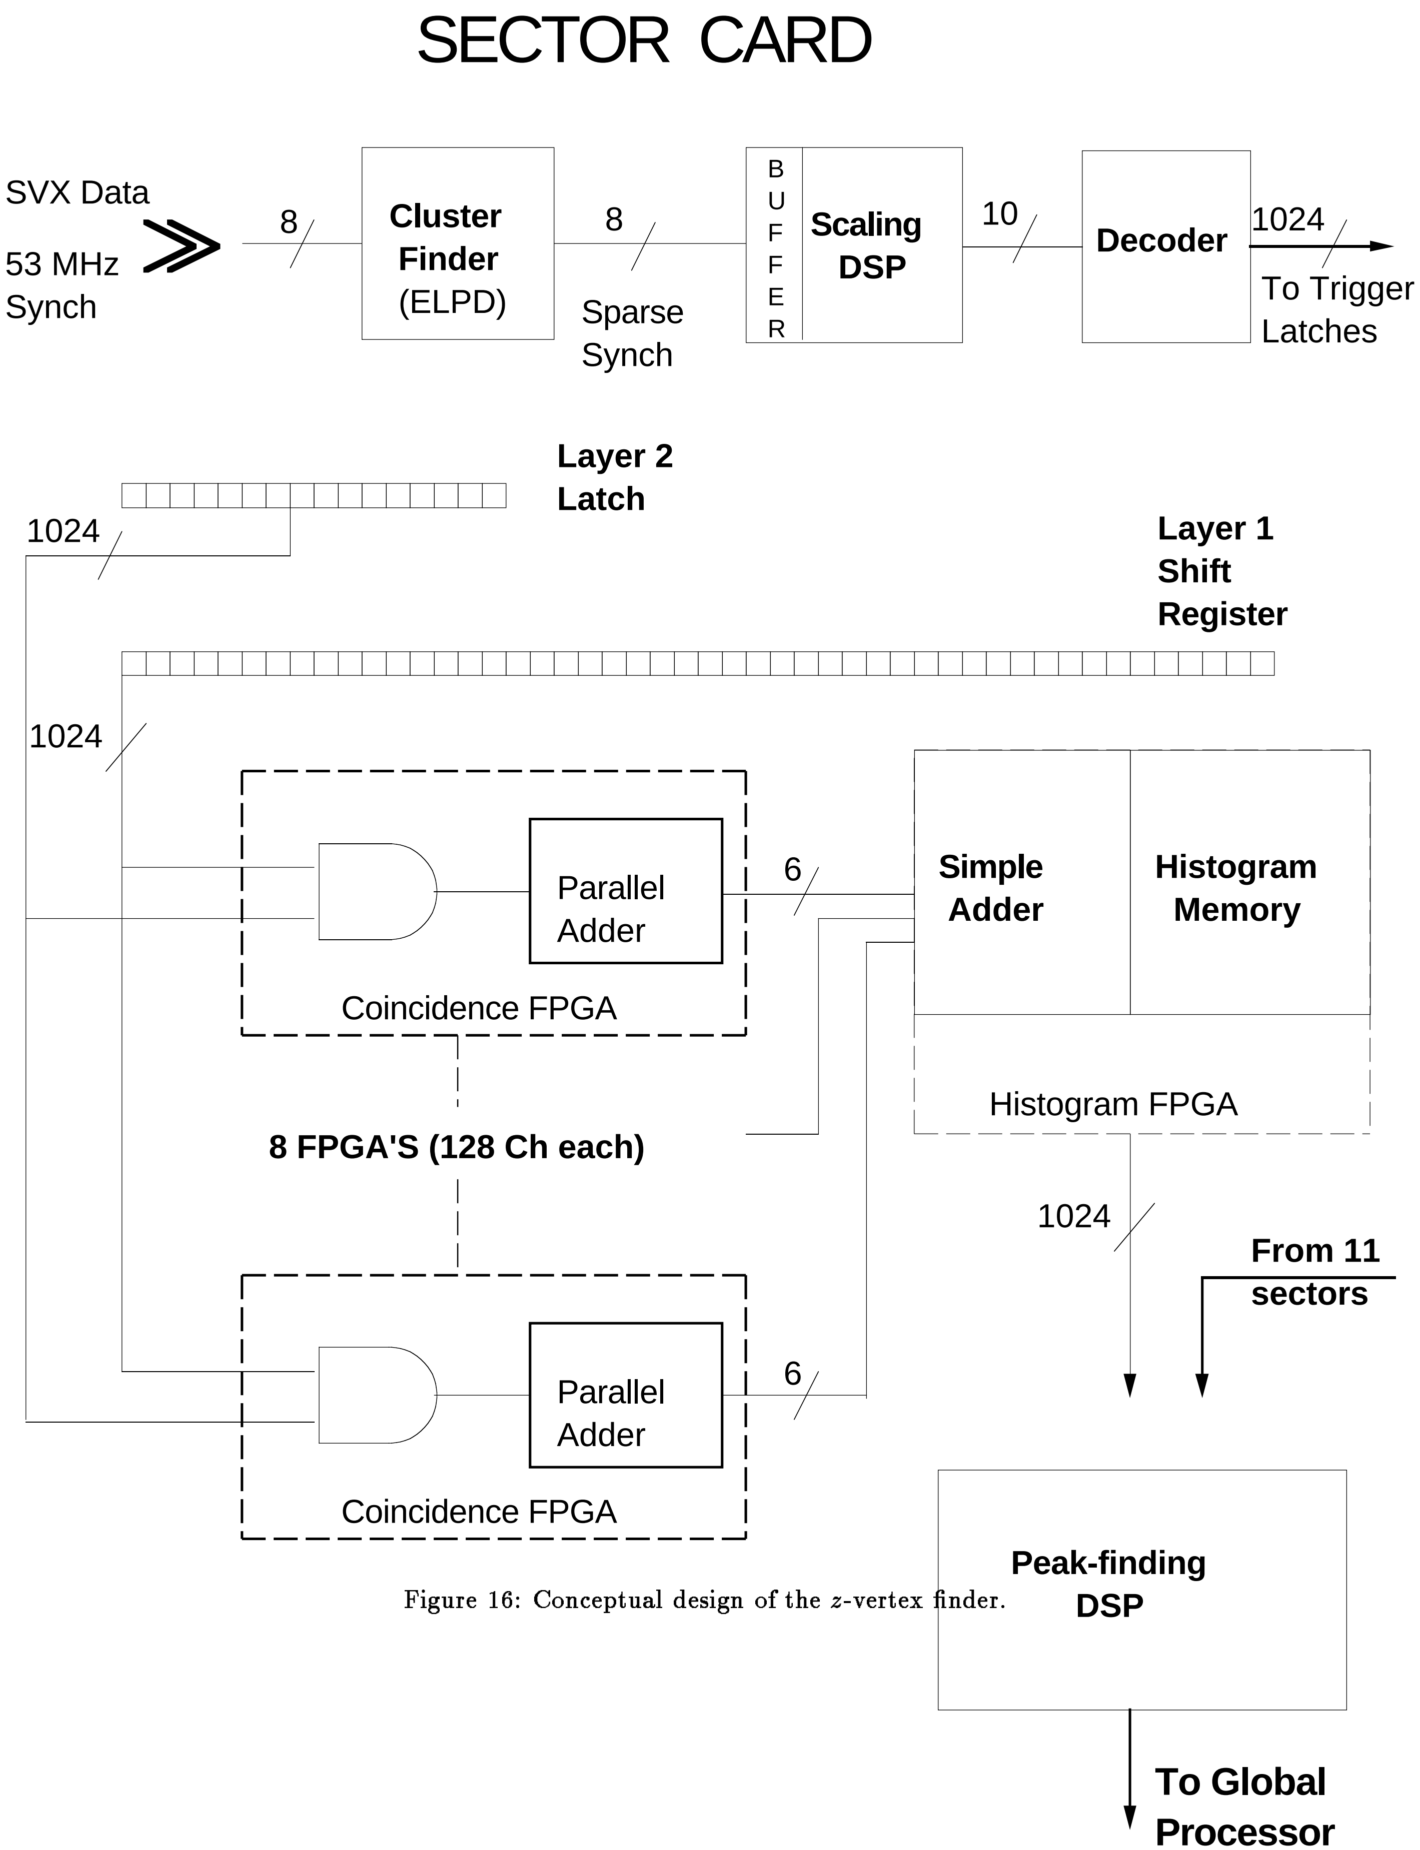

Figure 16: Conceptual design of the z-vertex finder.

STT information should be around 1 per sector while the maximum possible would be 32. Transfer speed to the MBT will be 160 Mb/s meaning that transfer times for the track parameters would average 0.8 μs and could reach as high as 26 μs.

In addition to the track parameters it is also desirable to have cluster information available at the CFTpp. The implications of this option are discussed in more detail in the next section.

Table 3: Data transmitted to the CFTpp crate per track. Also given is the anticipated size of each quantity per track in bits and where the quantity is originally defined.

Quantity	Bits	Source
CFT track info and number	24	fiber road card
b (signed)	16	processor card
p_T	16	processor card
ϕ_0	16	processor card
χ^2	16	processor card
x position of beam	16	processor card
y position of beam	16	processor card
layers used bit mask	4	processor card
padding (?)	4	
total	**128**	

2.6.3 Data Processing in the CFTpp and the global Level-2 Processor

Output of the results of the CFTpp processing to the global Level-2 processor will be done over Cypress Hot Links from MBT(s) in the CFTpp crate to an MBT in the global Level-2 processor. Algorithms for combining the STT information with the track trigger are currently under discussion, as are the details of exactly what level of processing will be done with the data in the CFTpp and global Level-2 processor.

2.7 Communication with Level-3

2.7.1 Readout Requirements

According to DØ rules all data created by a processor must be read out by Level-3 on every event. The STTpp creates two kinds of data: hit clusters and tracks. The raw SMT data are read out through the VRB crates and the CFT Level-1 track candidates are read out from the CFTpp. The following sections describe three alternative schemes for readout to Level 3.

2.7.2 Readout from Trigger Cards and CFTpp

To save unpacking time in Level-3 it may be desirable to read all hit clusters out. These data are only available on the trigger cards, so that in this case the trigger cards have to be read out. Assuming that there are as many clusters on the stereo side of a double sided detector as there are on the axial side, we expect a total of 260 clusters on average and a maximum of 1000 clusters in each STTpp crate. Every cluster position requires 18 bits (see figure 14) and all clusters correspond to 0.6–2.3 kB of data that have to be read out from the trigger cards.

For each road we also need to save the results of the track fit, which can be characterized for example by

- CFT Level-1 information or track number,
- transverse momentum or radius of curvature,
- azimuthal direction of the track,
- signed impact parameter,

- x position of beam,
- y position of beam,
- χ^2 from fit,
- bit mask that indicates which hits were used in fit.

Each of these should fit into a 16-bit word. For 2–32 tracks that means 32 bytes to 500 bytes of data.

The complete track information is transferred to the CFTpp and read out to Level-3 from there. The serial link transfers data at a rate of 160 Mbits/s=20 MB/s so that this takes 1.6–25 μs. In this scheme the information which roads a cluster lies in would not be saved.

2.7.3 Readout from Trigger Cards only

For diagnostic purposes it is necessary to read out cluster positions along with their associations to roads, but it is probably sufficient to read out the clusters that are in a road. This only requires reading out the processor cards. In this scheme we need to read out 42–500 clusters and 2–32 CFT tracks. This corresponds to $95 + 34 \approx 130$ bytes to $1.1 + 0.5 \approx 1.6$ kB of data

We may also read out the stereo clusters from the trigger cards which amounts to another 0.3–1.1 kB.

If all information is read out from the STTpp crate to Level-3, only "good" tracks have to be transferred from the processor cards to the CFTpp. We define "good" tracks as having hits in at least three layers of the SMT. This reduces the number of tracks to half the number of roads. We then have to transfer on average one track per STTpp crate or 32 bytes, which takes 0.8 μs. For the maximum of 32 roads it takes 26 μs.

2.7.4 Readout from CFTpp only

Alternatively, we can transfer the clusters and tracks in all roads to the CFTpp and have Level-3 read it out from there. This would remove the need for any Level-3 buffer controls and VBDs in the STTpp crate. Transferring 0.13–1.6 kB of data takes 6.5–80 μs.

3 Performance

3.1 Queueing Simulations

3.2 Track Reconstruction Algorithms

The track reconstruction algorithm will be implemented in the processors on the track fit cards. It has to determine the track parametrization that best fits the list of hits in a given road. This task consists of two parts: (1) selecting the hits that belong to the track and (2) fitting the parameters. As the queuing studies described above show, the time budget for this task is about 50 μs. In this section we discuss the algorithm options and the processors required to implement them within this time budget.

Once the set of hits associate with a given CFT road has been found, the next step in the processing is determining the track parameters in the $r\phi$ plane: the impact parameter b, the track direction ϕ_0 at the point of closest approach and the track curvature κ. There are two parts to this task: (1) determining which of the hits belong on the track and (2) determining the track parameters. In practice, the order and separability of these two steps depends critically on the algorithm being used. This section contains a description of the track definition, a description of algorithms being tested, some comparisons between the algorithms and a brief summary of initial timing studies.

3.2.1 Track Parameter Derivation

The track parameters are the impact parameter b, the track direction ϕ_0 at the point closest to the beam, and the track curvature κ. We assume

- no energy loss or multiple scattering from interactions in the beam pipe or the tracking detectors, so that the trajectory is a circle in the transverse plane;
- $b\kappa \ll 1$ and $\phi - \phi_0 \ll 1$ for a silicon hit at the position (r, ϕ), so that the trajectory can be approximated by a linear function of the parameters.

These requirements are valid for all tracks with $b \leq 2$ mm and $p_T < 1.5$ GeV[5]. The expression for the track trajectory is:

$$\phi(r) = b/r + \kappa r + \phi_0. \tag{2}$$

Here r is the radial position of a point on the track and ϕ is its azimuth. The best values for b, κ and ϕ_0 are determined by minimizing the χ^2 given by

$$\chi^2 = \sum_{i=hits} \left[\frac{\phi_i - (b/r_i + \kappa r_i + \phi_0)}{\sigma_i} \right]^2. \tag{3}$$

The set of hits consists of one hit inside the road per SMT layer and the the two CFT hits that define the road. The parameters that minimize χ^2 can be determined analytically and therefore fast, because the fitting function is linear in the parameters.

3.2.2 Hit Filtering Algorithms

In order to obtain the best possible performance, it is important to use only the silicon hits that belong to the charged particle track. If there is more than one hit per layer in the road we have to determine the "correct" hit. The number of hits per layer per road depends on the event topology, the luminosity and the noise in the detector. Table 4 shows the average number of hits/layer/road and the fraction of layers having two or more hits for different physical processes and luminosity conditions. On average there are ???? interactions at ??? pb^{-1}. Given the large number of layers with multiple hits, attention must be paid to choosing hits. We are considering a number of different algorithms. These are

1. **Static Road-Center**: The hits closest to the center of the road defined by the CFT hits and the beam position are used in each layer.

[5] The condition $b\kappa \ll 1$ is valid for a much larger range. However, the condition $\phi - \phi_0 \ll 1$ breaks down as impact parameters approach 2 mm.

Table 4: Hit multiplicity per road per layer (N) and fraction of tracks with two or more hits per road per layer (F_2) for different number of interactions/beam crossing.

Process	# of Interactions	N	F_2
$Z \to b\bar{b}$	1		
	3		
	7		
$WHqb$	1		
	3		

2. **Dynamic Road-Center**: The hits closest to the center of the road defined by the CFT hits and the hits in SMT layer 4 are used in each layer. This algorithm requires looping over all hits in SMT layer 4.
3. **All Combinations**: All combinations of hits are considered and the one with the best fit is chosen.
4. **Best Combination at Layer**: Moving from the outer SMT layer in at every layer the hit is chosen that gives the best fit.
5. **Reduced All Combinations**: Ignoring the highest multiplicity layer, all combinations of hits are considered and the one with the best fit is chosen.

The last two algorithms are based on the *all combinations* algorithm, but truncate the number of trial combinations in a controlled manner.

We judge the performance of each algorithm using two figures of merit determined from Monte Carlo event samples: (1) the number of Monte Carlo tracks contributing hits to a reconstructed track (N_{MC}) and (2) the χ^2 of the fitted parameters with respect to the Monte Carlo track that contributes the largest number of hits. Table 5 presents these comparisons for the samples listed in table 4.

Table 5: Figures of merit for algorithms 1–5.

Process	# of Interactions	Algorithm 1		Algorithm 2		Algorithm 3		Algorithm 4		Algorithm 5	
		N_{MC}	χ^2	N_{MC}	χ^2	N_{MC}	χ^2	N_{MC}	χ^2	N_{MC}	χ^2
$Z \to b\bar{b}$	1										
	3										
	7										
$WHqb$	1										
	3										

The simulations described below use only algorithm (1). The comparison in table 5 indicate that algorithm (3) or its variants perform somewhat better than algorithm (1) at high luminosity. To make the the final processor choice it is important to understand the gains these algorithms offer, because they require significantly more processing power. Studies of these algorithms are in progress.

3.2.3 Timing Studies

We are considering a number of different processor options. The leading choices are (1) a Digital AXP-based multi-processor board, (2) TI????? floating point DSPs, or (3) arrays of programmable logic devices.

Timing results based on a 500 MHz AXP processor show that a single track is processed in ???? μs using algorithm (1). For algorithm (3), the average time is ????? and 95sample and for a $WH \to q\overline{q}'b\bar{b}$ sample the corresponding times are ???? and ????. Thus, for a maximum of 32 tracks per section, the AXP solution works provided we have ???? processors per sector. The time budget cannot be met for algorithm (3).

Studies for a TI?????-based DSP solution indicate that a fixed-point DSP cannot provide the necessary performance because floating point arithmetic must be emulated in a fixed point processor. Initial estimates indicate that algorithm (1) could be performed in roughly 5 μs on the corresponding floating point DSP. Further studies are being conducted for DSP- and FPGA-based solutions. The FPGA solution is attractive because a high degree of parallel processing can be achieved. This could allow algorithm (3) or its variants to be performed in nearly constant time. The key issue is the speed of floating point multipliers and adders implemented in programmable logic. The preliminary design and cost estimate presented in this document are based on DSPs with one processor/track road.

3.3 Expected Resolutions

We use Monte Carlo simulations of single electrons and muons with $1.5 \leq p_T \leq 15$ GeV in the pseudorapidity range $|\eta| \leq 1.0$ to measure the resolution achieved by the track reconstruction in the STTpp [35]. We simulated the response of the detector to these particles using the GEANT program [28]. The track reconstruction algorithm for the STTpp at L2 used in these simulations [42] corresponds to algorithm (1) in section 3.2 with the following exception: if reduced $\chi^2 > 4$, the SMT hit with the largest contribution to χ^2 is dropped and the remaining hits were fit again. All tracks with hits in at least three SMT layers and reduced $\chi^2 \leq 4$ were considered "good" tracks.

We measure p_T in three different ways:

1. "CFT L1": using all eight CFT hits (this corresponds to the p_T measurement that the CFT Level-1 trigger performs internally;
2. "CFT L2": using the CFT hits in the inner and outer layers only (this is the information reported to the Level-2 trigger and thus corresponds to the measurement of p_T that is available in the Level-2 trigger without the STTpp);
3. "STTpp": using the CFT hits in the inner and outer layers and the SMT hits (this corresponds to the measurement performed by the STTpp).

Figure 17 shows that the STTpp improves the p_T resolution in the Level-2 trigger by a factor 2–3, depending on p_T.

The improved p_T resolution allows the use of an E_T/p_T cut for electron identification at Level-2. Fig. 18 shows a comparison of E_T/p_T with p_T calculated as in CFT-L2 and as in STTpp for electron tracks with $p_T = 5$ and 20 GeV. The better resolution results in an improvement in the background rejection for the E_T/p_T cut at Level-2 by a factor of 2 for $p_T = 20$ GeV and by a factor 4 for $p_T = 5$ GeV, relative to using CFT information only to measure p_T.

Figure 19 shows the resolution of the impact parameter measurement from the STTpp using single muons. The dependence of the resolution on p_T can be parametrized by

$$\sigma_b(p_T)^2 = \sigma_0^2 + \left(\frac{a}{p_T}\right)^2 \tag{4}$$

with $\sigma_0 = 12.6$ μm, and $a = 49.7$ GeV$\cdot\mu$m. This parameterization also reproduces the impact parameter resolution for tracks from the primary vertex in $b\bar{b}$ and $t\bar{t}$ events.

We define the impact parameter significance

$$S_b = \frac{b}{\sigma_b(p_T)} \tag{5}$$

using $\sigma_b(p_T)$ from equation 4. As shown in Fig. 20, the distribution of impact parameter significance for tracks coming from the primary vertex is satisfactorily described by a Gaussian for both $b\bar{b}$ and $t\bar{t}$ events, although there are non-Gaussian tails. The width is very close to 1 for $b\bar{b}$ events, with a slight degradation and more pronounced tails in $t\bar{t}$ events, due to the higher track multiplicity. Furthermore, in both $b\bar{b}$ and $t\bar{t}$ events, the impact parameter significance distribution for tracks from the primary vertex is significantly narrower than that for all tracks (see Fig.21).

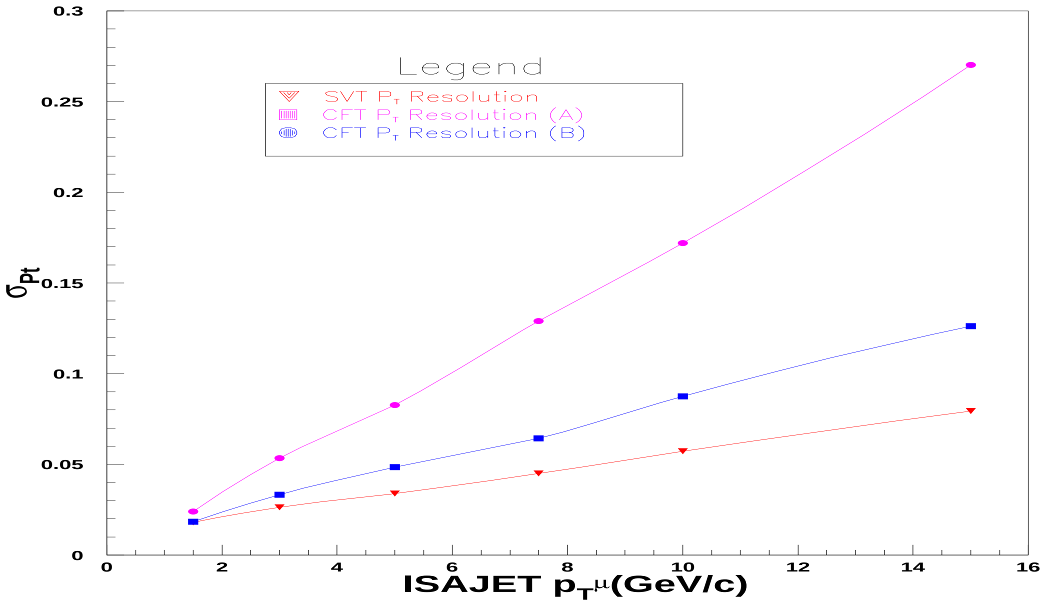

Figure 17: p_T resolution for single muons in the Level-1 and Level-2 trigger.

3.4 Simulation of Physical Processes

3.4.1 Trigger Simulation

To simulate the performance of the STTpp, we generate events from processes of interest using the event generators ISAJET[25], PYTHIA[26], or COMPHEP[44]) and simulate the detector response using the GEANT program. For all simulations we assume that the interaction vertex position along the beam direction (z_v) was normally distributed with mean(z_v)=0 and rms(z_v)=25 cm, as expected for Run II. We find that the expected transverse size of the beam of 30–40 μm has only a small effect on the results[30] and therefore assume infinitely narrow beams.

We estimate signal efficiencies and trigger rates with and without the STTpp using Monte Carlo simulations. Without STTpp, triggers at Level-1 and Level-2 are based on information from the calorimeter, muon system, and CFT, only. The additional requirements that can be imposed by the STTpp are:

- $\geq N$ tracks with $p_T > T$;
- $\geq N$ tracks with $b > T$;
- $\geq N$ tracks with $S_b > T$,

where N and T are programmable parameters. The track reconstruction algorithm for the STTpp at L2 used in these simulations [42] corresponds to algorithm (1) in section 3.2 with the following exception: if reduced $\chi^2 > 4$, the SMT hit with the largest contribution to χ^2 was dropped and the remaining hits were fit again. All tracks with hits in at least three SMT layers and reduced $\chi^2 \leq 4$ were considered "good" tracks.

The background event sample used to determine the trigger rates consists of dijet events, generated in six p_T bins. For each p_T bin, samples were created with 1, 3, 5, or 7 additional interactions. The additional interactions were simulated by low p_T dijet events. These 24 samples are combined weighted by their respective cross section times the probability of seeing the respective number of interactions for ??? pb^{-1}. The results for +0, +2, +4, +6 and +8 interactions, for which no events were generated, were found by interpolating the results from the generated files.

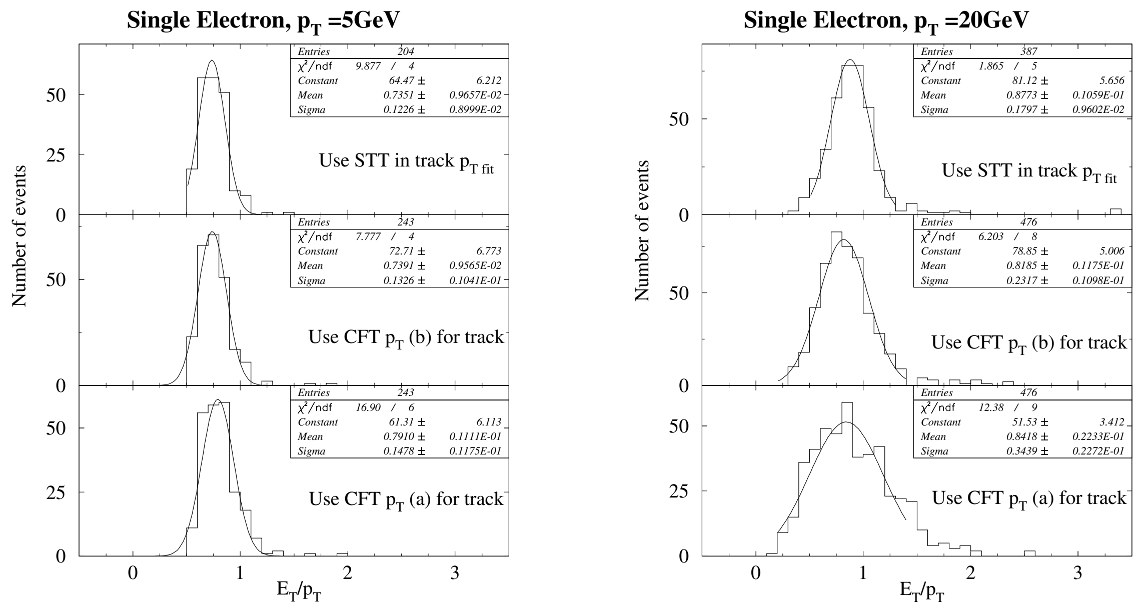

Figure 18: E_T/p_T for electrons with p_T = 5 GeV (left) and p_T = 20 GeV (right), assuming a calorimeter energy resolution of $13\%/E_T^{1/2}$.

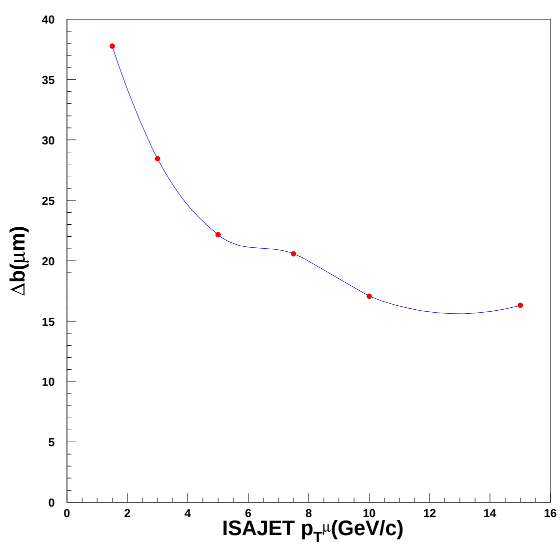

Figure 19: Impact parameter resolution vs p_T for single muons.

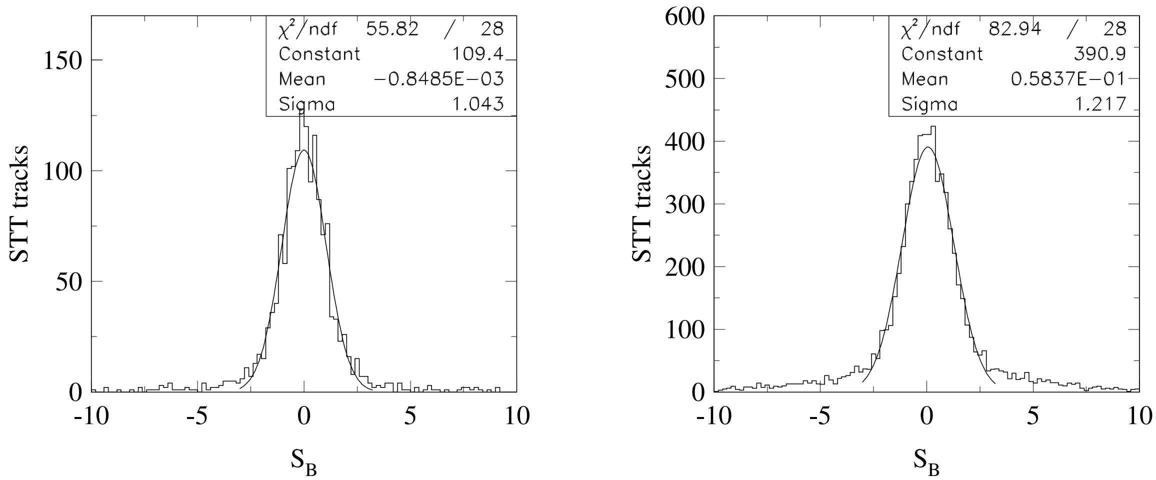

Figure 20: Impact parameter significance distribution for tracks matched to MC tracks coming from primary vertex, for $b\bar{b}$ events (left), and for $t\bar{t}$ events (right)

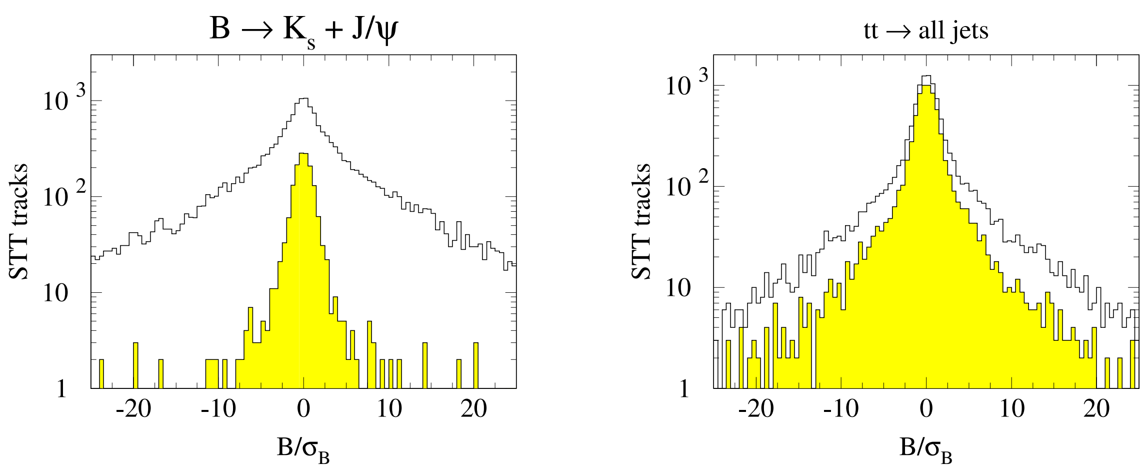

Figure 21: Impact parameter significance distribution for matched tracks, for $B \to K_s + J/\psi$ events (left) and $t\bar{t} \to$ *all jets* events (right). The shaded histograms correspond to tracks from the primary vertex.

3.4.2 Top Quark Pair Production

To study the benefits of the STTpp for the selection of $t\bar{t}$ events [30], two samples of events were generated using ISAJET:

- a sample of 2126 events in which both W bosons decay into hadrons, denoted as $t\bar{t} \to$ *all jets*;
- a sample of 1175 events in which one of the W bosons is forced to decay into $\mu\nu$, denoted as $t\bar{t} \to \mu + jets$.

To trigger efficiently on $t\bar{t} \to$ *all jets* and $t\bar{t} \to \mu + jets$ decays we have to require the logical or ($\cup$) of a series of trigger conditions:

1: ≥ 4 jets with $p_T > 15$ GeV;

2: ≥ 3 jets with $p_T > 10$ GeV and $\not{p}_T > 20$ GeV;

3: ≥ 1 jet with $p_T > 45$ GeV;

4: ≥ 2 jets with $p_T > 20$ GeV;

5: ≥ 1 muon with $p_T > 6$ GeV and ≥ 2 jets with $p_T > 10$ GeV and ≥ 1 jet with $p_T > 20$ GeV.

The change in signal efficiency and background rate for additional requirements based on the STTpp is demonstrated in Fig. 22 for trigger condition 5 and for the $t\bar{t} \to$ *all jets* sample. The results are similar for the $t\bar{t} \to \mu + jets$ sample. About 8% of the events are lost when STTpp-based requirements are imposed due to the acceptance of the SMT along the z axis. An efficiency well above 80% is obtained when requiring only one track with a large impact parameter significance. The trigger rate can be reduced by a factor of ≈ 2 when requiring ≥ 1 track with $S_b > 3$, and a reduction factor of 4 is reached when ≥ 2 tracks with $S_b > 3$ are requested. The $t\bar{t} \to$ *all jets* trigger efficiencies corresponding to these two criteria are 86.9% and 77.7%, respectively.

We varied the transverse beam position to study the impact of finite beam size and found that a 30 μm uncertainty in the transverse vertex position leads to a 5% reduction in the signal efficiency for a fixed background rate. A systematic shift of the transverse vertex position in the of ≤ 100 μm does not introduce a signal efficiency loss as long as the STTpp criterion is adapted to maintain the background rate at a constant level.

Trigger conditions 4 and 5 account for most of the background rate but they are required to achieve high signal efficiency. If we require a logical and ($\cap$) of these two trigger conditions with the additional condition

6: ≥ 2 tracks in STTpp with $S_b > 2.5$

we can reduce the background rate substantially with minimal loss in signal efficiency. The rates and efficiencies for the two cases are listed in Table 6.

Table 6: Background rates and $t\bar{t}$ trigger efficiencies.

Trigger Condition	Background Rate		Efficiency	
	Level-1	Level-2	$t\bar{t} \to \mu + jets$	$t\bar{t} \to$ *all jets*
$1 \cup 2 \cup 3 \cup 4 \cup 5$	240 Hz	110 Hz	96%	100%
$1 \cup 2 \cup 3 \cup (4 \cap 6) \cup (5 \cap 6)$	240 Hz	40 Hz	93%	99%

3.4.3 Single Top Production

The Monte Carlo event samples of for this process were generated with the COMPHEP generator. We select a combination of electron, muon and jet triggers with high efficiency for single top events and investigated the effect of additional STTpp conditions on signal efficiency and background rejection [36]. The results for ≥ 1

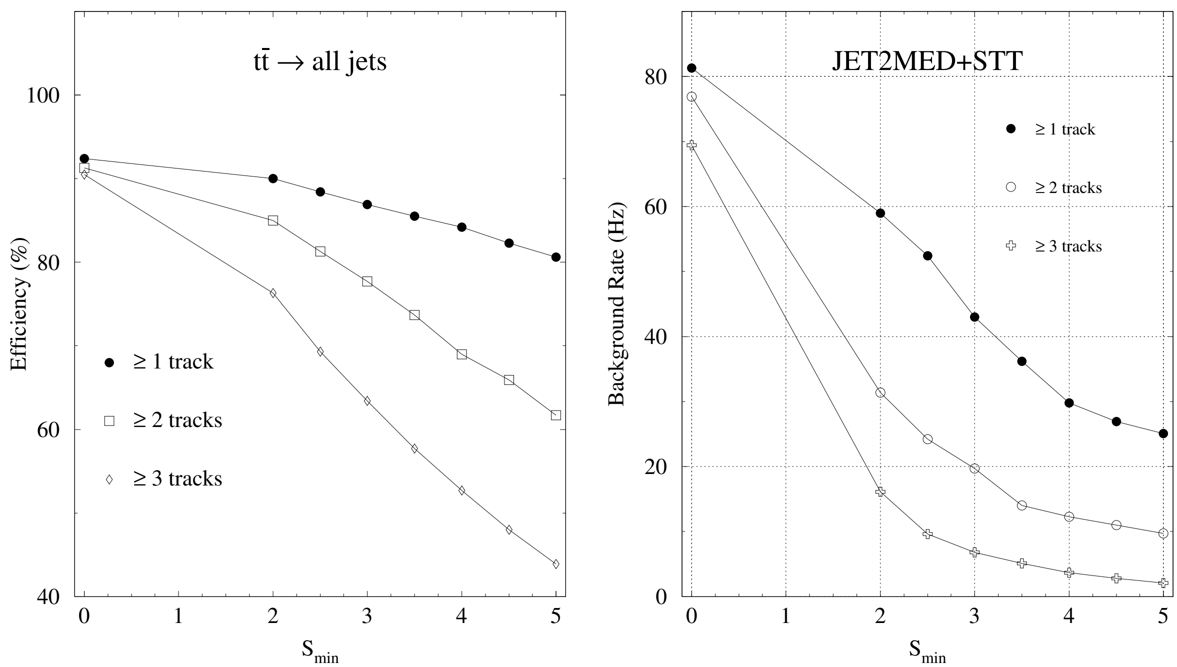

Figure 22: Signal efficiency for $t\bar{t} \rightarrow$ *all jets* events (left) and background rate (right) for the events satisfying trigger condition 5 together with the STTpp trigger condition $\geq N$ tracks with $S_b > S_{min}$.

track with $|S_b > S_{min}$ as the STTpp condition are summarized in Table 7. For both leptonic and hadronic decay modes of the W, the STTpp allows a reduction of the trigger rate with little loss in signal efficiency. Using the STTpp may become indispensible if it turns out that the background rates are underestimated by the simulation.

Table 7: Background rejection and signal efficiency for single top production.

Trigger Condition	Background Rate	Additional Rejection		Additional Efficiency	
		$S_{min} = 2$	$S_{min} = 4$	$S_{min} = 2$	$S_{min} = 4$
electron triggers	40 Hz	1.6	3.0	87%	74%
muon triggers	7 Hz	1.4	7.0	90%	78%
jet triggers	200 Hz	1.4	3.0	94%	80%

3.4.4 Associated Higgs Boson Production

We can easily trigger on associated Higgs production if the vector boson decays into electrons or muons. If the vector boson decays into jets, however, triggering becomes much more diffcult. To study this process, we generated a sample of 1700 $p\overline{p} \rightarrow WH^0 \rightarrow q\overline{q}b\overline{b}$ events using PYTHIA with $M_H = 90$ GeV [32]. We studied a number of calorimeter-based trigger conditions requiring between 3 and 5 jets with p_T thresholds between 5 and 15 GeV. In order to maintain reasonable signal efficiency we need to keep the required jet multiplicity and the p_T threshold low. A trigger requiring ≥ 3 jets with $p_T > 10$ GeV at Level-2 achieves a signal efficiency of 74% at a background rate of 63 Hz.

Additional trigger conditions using track information offer the possibility to reduce the trigger rate while maintaining acceptable signal efficiency (see Table 8). For relative trigger efficiencies of 90%, the calorimeter trigger rate can be reduced by a factor of two to three. The use of STTpp information increases the rejection by 25% over the best trigger without STTpp information. As the trigger efficiency drops, the effect of STTpp information becomes much more significant. At a relative efficiency of 80%, the STTpp triggers have rejection which is a factor of 1.5 better than the CFT-based triggers. For a relative efficiency of 50%, the STTpp triggers have rejection factors of up to 40, while the CFT trigger rejection is only a factor of 10. The best algorithm is that which requires $\geq N$ STTpp tracks with $p_T > 3$ GeV. From these results, it is clear that the STTpp can significantly help in the $W/Z + H$ search.

Table 8: Background rejections for four (arbitrarily chosen) signal efficiencies ϵ, relative to a trigger using only calorimeter information.

Trigger Condition	$\epsilon \approx 90\%$	$\epsilon \approx 80\%$	$\epsilon \approx 50\%$	$\epsilon \approx 10\%$
$\geq N$ CFT tracks with $p_T > 3$ GeV	2.5	4.7	10	35
$\geq N$ STTpp tracks with $p_T > 3$ GeV	3.1	7.5	40	1000
$\geq N$ CFT tracks with $p_T > 3$ GeV and $\overline{b_{25}} > X$	3.0	6.4	26	500
$\geq N$ CFT tracks with $p_T > 3$ GeV and $\overline{S_{25}} > X$	2.8	7.1	33	650

$\overline{b_{25}}$ - average impact parameter for tracks 2–5
$\overline{S_{25}}$ - average impact parameter significance for tracks 2–5
(tracks ordered in decreasing p_T)

3.4.5 Z Boson Decay to $b\overline{b}$

For this study [34], we use 10000 $Z \rightarrow b\overline{b}$ events generated with ISAJET. Most $Z \rightarrow b\overline{b}$ events have at least two jets with $p_T > 20$ GeV and we therefore define a set of calorimeter-based triggers which require the presence of at least two jets above a certain p_T threshold. For 40% efficiency the rate is about 25% of the available bandwidth, which is unacceptably high, and even for 20% efficiency the L2 rate corresponds to almost 10% of the bandwidth.

To control the rate, we use CFT information at trigger level-1. For an acceptable efficiency of 35% out of Level 2, the background rate still remains at about 20% of the bandwidth. Without additional rejection power this trigger can only be operated with a prescale and will not be able to exploit the full integrated luminosity.

Using the information provided by the STTpp provides the required additional background rejection. With the STTpp, we gain a factor of 2 in background rejection compared to CFT-based triggers for the same overall efficiency of 35%. For a background rate of as low as 20 Hz, we still have a signal efficiency of around 20%, five times higher than the calorimeter-based trigger with the same background rate. Thus, the STTpp gives us the capability of triggering on $Z \rightarrow b\bar{b}$ decays at a low rate while maintaining an acceptable efficiency.

With a trigger efficiency of 20% and an estimated efficiency for tagging both b jets during the offline analysis of 25% we expect about 50000 signal events per fb^{-1}. Assuming the double b-tag reduces the gluon-jet and light-quark background by an order of magnitude we expect 10^7 background events per fb^{-1}. For an integrated luminosity of 1 fb^{-1} the expected signal significance is thus about 17.

Table ?? summarizes these background rates and trigger efficiencies and Fig. 23 shows a plot of background rate versus efficiency for various trigger conditions relative to requiring ≥ 2 jets with $p_T > 15$ GeV.

Table 9: Background rates and signal efficiencies for $Z \rightarrow b\bar{b}$.

Trigger Condition	Background Rate		Efficiency
	Level-1	Level-2	
≥ 2 jets with $p_T > 10$ GeV	8000 Hz	1400 Hz	67%
≥ 2 jets with $p_T > 15$ GeV	620 Hz	260 Hz	41%
≥ 2 jets with $p_T > 20$ GeV	190 Hz	85 Hz	20%
≥ 2 jets with $p_T > 15$ GeV and ≥ 2 CFT tracks with $p_T > 3$ GeV	450 Hz	200 Hz	35%
≥ 2 jets with $p_T > 15$ GeV and ≥ 1 STTpp track with $S_b > S_{min}$	620 Hz	110 Hz	35%
≥ 2 jets with $p_T > 15$ GeV and ≥ 1 STTpp track with $S_b > S_{min}$	620 Hz	20 Hz	20%

3.5 Primary Vertex Determination

The Level-3 system for Run II has 25 ms to reconstruct an event. This may not be long enough to unpack the hits from the SMT and use them to find the number and z-positions (z_v) of the primary interaction vertices. We therefore propose to find the primary interaction vertices at Level-2 in a dedicated processor and transmit them to Level-3 for use in computations of derived quantities, like the p_T of jets and electrons. Knowledge of the z-position of the primary vertex at Level-3 is extremely beneficial. During Run I, the Level-2 electron algorithm used the z-position of the primary vertex from Level-0 in order to get a better estimate of the electron p_T.

The hits in the 90° strips of the silicon ladders in layers 1 and 3 of the inner four barrel segments of the SMT offer the possibility to determine the z-position of the primary event vertex Level-2 [33]. The resolution of the measurement is expected to be 550 μm for events with one high-p_T interaction. For events with two additional interactions, the vertex resolution is 610 μm. For a sample of $Z \rightarrow b\bar{b}$ events without additional interactions and with the event vertex within $-25 < z_v < 25$ cm, the efficiency for reconstructing a vertex is 98%. If we superpose two additional interactions on these $Z \rightarrow b\bar{b}$ events, the efficiency of reconstructing at least one vertex is 99%. The efficiency for identifying the correct vertex as the high-p_T vertex is 78%. In these events there maybe one or more vertices within the acceptance region and we find that the efficiency to reconstruct a vertex in this environment is about 85%.

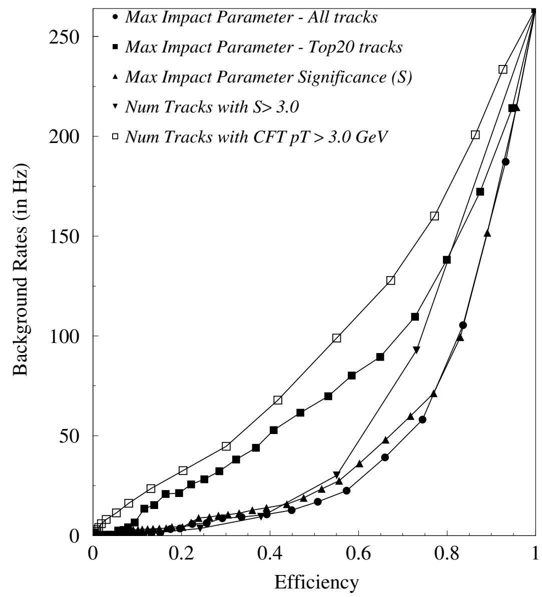

Figure 23: Background rate versus efficiency for $Z \to b\bar{b}$ decays relative to the trigger condition $\geq$ 2 jets with $p_T > 15$ GeV.

Due to the limited acceptance of the SMT (only the four inner barrel sections are equipped with 90° strips), the reconstruction of primary vertices with this method is limited to ± 30 cm around the center of the detector. The rms width of the luminous region during Run II is expected to be 25 cm. About 30% of the events have the high-p_T vertex outside this region. For these either no vertex is found or the vertex from a soft $p\bar{p}$ interaction inside this acceptance region is misidentified as the high-p_T vertex.

The influence of the knowledge of the vertex position on the p_T trigger turn-on is shown in Fig. 24. Even if the correct vertex position is only found for $|z_v| < 30$cm (dash-dotted curve in the figure), the p_T threshold still sharpens significantly. This improvement implies that we can trigger more efficiently and with lower rates on single electrons at Level 3. The electron energy resolution is assumed to be $13\%/\sqrt{E}$.

3.6 Beam Alignment and Stability

3.6.1 Initial Alignment of SMT with Beam

In the STTpp, the impact parameter is determined with respect to a nominal beam position. A misalignment of the detector and beam axes will cause a fake impact parameter, leading to increased background trigger rates and changes in the trigger efficiency.

In the vicinity of the interaction point the beam has a Gaussian width σ which varies with the distance from the nominal crossing point ($z = 0$):

$$\sigma(z) = \sigma(0) \cdot \sqrt{1 + \left(\frac{z}{\beta^*}\right)^2}, \tag{6}$$

where $\sigma(0)$ is the beam size at the nominal interaction point, and z is the distance from the nominal interaction point. During Run II $\beta^* = 35$ cm and $\sigma(0) = 40.5$ μm are expected. Figures 25–?? show the effect of beam tilt on the impact parameter distribution, assuming that the z distribution of the primary vertices is a Gaussian with a width of 25 cm. Table 10 summarizes the effect of beam tilt. For beam tilt angles larger than 200 μrad, the fraction of tracks with large fake impact parameters becomes too large, and gives rise to unwanted triggers.

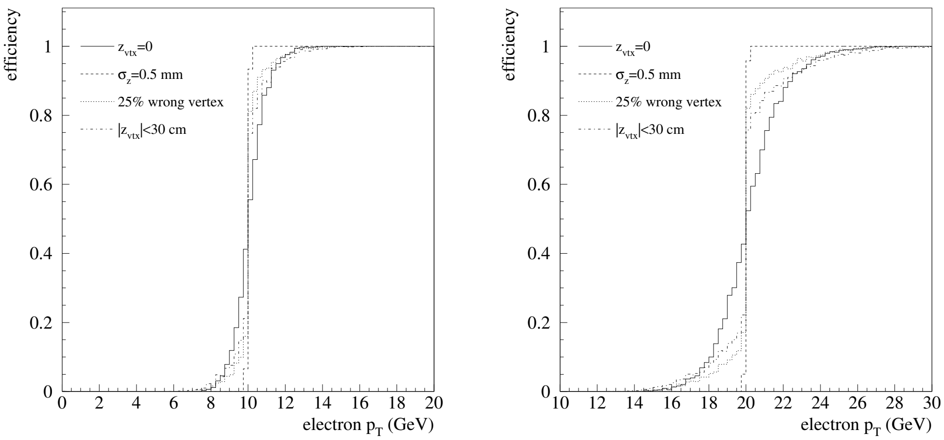

Figure 24: Turn-on curves for electrons with $p_T = 10$ GeV (left plot) and 20 GeV (right plot). For the solid curve $z_v = 0$ is assumed, for the dashed curve the vertex is reconstructed with 100% efficiency, for the dotted curve the vertex reconstruction efficiency (78%) is taken into account, and for the dash-dotted curve the limited detector acceptance is also taken into account.

tilt	mean(b)	fraction with		
		$b > 100\ \mu$m	$b > 200\ \mu$m	$b > 400\ \mu$m
0 μrad	32 μm	1%	0%	0%
100 μrad	35 μm	2%	0%	0%
200 μrad	43 μm	5%	0%	0%
400 μrad	66 μm	20%	2%	0%
600 μrad	92 μm	38%	8%	0%
800 μrad	118 μm	51%	17%	0%
1000 μrad	146 μm	58%	27%	3%
2000 μrad	218 μm	77%	57%	26%

Table 10: Fake impact parameter due to misalignments.

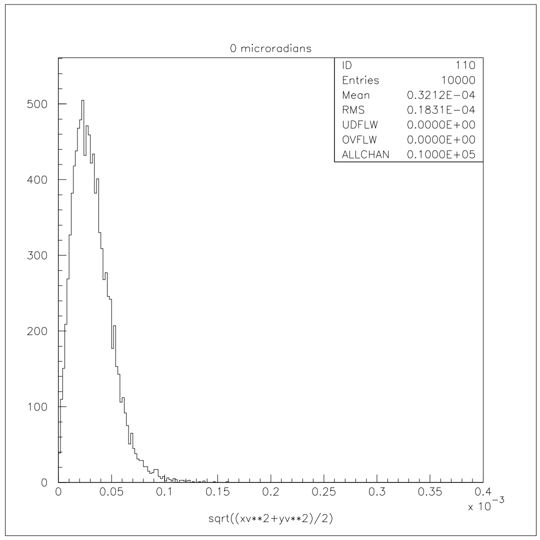

Figure 25: Fake impact parameter distribution with respect to (0,0) for tracks coming from primary vertex, for zero tilt angle.

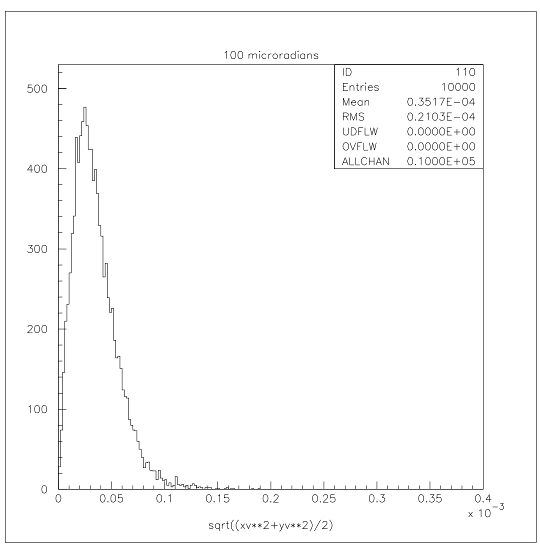

Figure 26: Fake impact parameter distribution with respect to (0,0) for tracks coming from primary vertex, for tilt angle = 100 μrad.

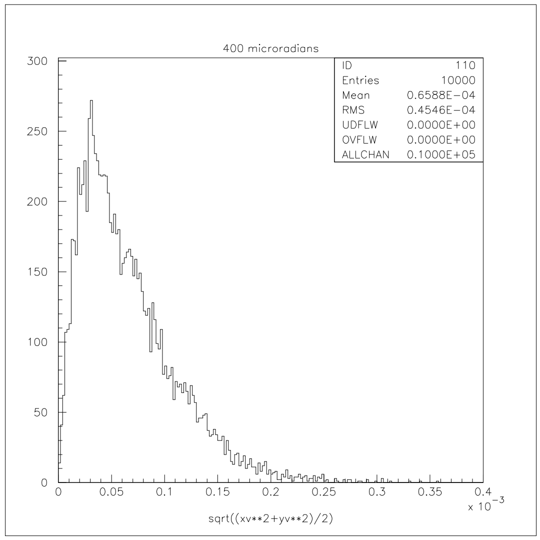

Figure 27: Fake impact parameter distribution with respect to (0,0) for tracks coming from primary vertex, for tilt angle = 400 μrad.

For optimal operation of the trigger, any residual offset between the detector and beam axes should be less than 1 mm, and the angle between the detector axis and the beam direction should be smaller than 200 μradians. A larger offset would cause inefficiencies from tracks that cross SMT sectors and a larger angle would cause too large a fake impact parameter. To achieve this precision, an iterative procedure is planned.

Before the detector is rolled into the collision hall, a nominal beam position is determined based on the position of the low-beta quadrupole magnets. This actual beam position is expected to be within about 1 mm of this nominal value. The detector is centered on the nominal beam position to a precision of about 0.5 mm. After the detector is in place the SMT can be used to obtain a precise measurement of the beam trajectory relative to the detector. Limited corrections (± 1 mm and ± 250 μrad) can be made to the beam trajectory through orbit tuning. If the misalignment exceeds these margins either the detector or the low-beta quadrupole magnets have to be repositioned.

3.6.2 Beam Stability

Any variations in the beam position must be kept below 30 μm (approximately the width of the beam) else they cause an increase in the trigger rate due to the apparent impact parameter. Variations in the relative angle of beam and detector axes must be kept below 200 μrad.

During Run I, CDF monitored the beam position using their silicon vertex detector. The average offset of the beam axis from the detector axis was about 0.5 mm, and the average angle between beam and detector axis was about 500 μrad. There was a long-term drift in both position and angle, with discrete steps before and after shutdown periods. Typical store-to-store variations in beam centroid position were about 100 μm horizontally and about 150 μm vertically. Angle variations from store to store were typically less than 100 μrad, and of the order of 200 μrad over the course of the run. Variations during a store were much less important, typically 40 μm horizontally and vertically. Angle variations during a store were too small to be measurable [**?**].

In order to compensate for variations in beam position, CDF in collaboration with the accelerator division, has implemented a feedback system, using the measured beam position to set dipole corrector magnets near B0 to stabilize beam position and angle [18]. The maximum possible range of variation is determined by the range of the magnet power supply ($\pm$50A), corresponding to a maximum angle correction of ± 100 μrad. Tests performed by CDF and the accelerator division in June 1995 showed that, with 5 minute updates, the beam position could be kept stable to about 5 μm, and the beam angle to about 50 μrad. In a further test in February 1996 the beam was moved 50 μm vertically and 50 μm horizontally in the B0 interaction region. Closed orbit measurements verified that the bumps were local to the B0 interaction region without change elsewhere around ring, an indication that the beam can be controlled independently at D0 and B0.

Similar correction magnets exist at the D0 interaction region and DØ will implement a similar feedback system to maintain stability of the beam position to the required precision.

4 Cost Estimate and Schedule

Table 11 shows the cost estimate for development and construction of the full system. The project requires 5 person-years (py) of engineering support and $ 740k in equipment funds.

Table 11: Cost estimate for STTpp project.

Item	engineering	unit cost	number	total cost
trigger card	1.5 py	$ 5200	60	$ 312k
FPGAs for trigger card	0.5 py	$ 273	500	$ 137k
processor card	1.5 py	$ 10000	12	$ 120k
receiver/controller card	0.5 py	$ 5200	6	$ 31k
transceiver card[1]	—	$ 5000	8	$ 40k
alpha processor[1]	—	$ 15000	2	$ 30k
VME crates/power supplies	—	$ 10000	7	$ 70k
system integration	1.0 py	—	—	—
total project cost	5.0 py			$ 740k

[1] under development for the DØ trigger system

Table 12 shows a schedule for the design, testing, implementation and commissioning of the STTpp. The device will not be available for the initial commissioning phase of Run II, which starts in April 2000. At the beginning of the run the luminosity is expected to be below design and the rejection afforded by the STTpp is not required. The full impact of the STTpp will come to bear later in the run, when the accelerator reaches its design luminosity.

Table 12: Schedule for development of STTpp.

date		milestone
April	1998	collaboration review of STTpp design
September	1998	award of grant by NSF
June	1999	prototype subprocessor crate
October	1999	finalize design, begin production
March	2000	system test
April	2000	begin Run II
August	2000	installation at DØ
October	2000	STTpp operational

5 Conclusion

We have presented an initial design of a Silicon Track Trigger preprocessor for the DØ Level-2 trigger system in Run II. This processor will associate hits in the SMT with CFT tracks, found at Level-1 and fit them jointly for precise impact parameter and curvature determinations. Knowledge of the impact parameter at Level-2 will allow more efficient data collection for physical processes with b quarks in their final state. Such processes are of fundamental importance to many of the questions that we want to address in Run II. The proposed processor also offers the possibility to compute the position along the beam trajectory of the primary $p\overline{p}$ interactions in an event. This information is important to improve momentum resolution for electrons and jets in Level-3 and can serve as seed information for a fast tracking algorithm in Level-3. The design presented shows that implementation of this functionality is possible with commercially available components. Its details will have to be worked out in the coming months.

References

[1] D. Amidei and R. Brock, eds., FERMILAB-PUB-96/082.

[2] S.L. Glashow, Nucl. Phys. **22**, 579 (1961); S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967); A. Salam, *Proceedings of the 8^{th} Nobel Symposium*, ed. N. Svartholm (Almqvist and Wiksells, Stockholm 1968), p. 367.

[3] "A Measurement of the W Boson Mass", B. Abbott *et al.*, FERMILAB PUB-97/422-E, submitted to Phys. Rev. Letters and FERMILAB PUB-97/423-E, submitted to Phys. Rev. D.

[4] D. Ward, EPS, Jerusalem, 1997.

[5] S. Abachi *et al.*, Phys. Rev. Lett. **79**, 1197 (1997);
S. Abachi *et al.*, Phys. Rev. Lett. **80**, 2063 (1998).

[6] "Measurement of the Top Quark Mass", F. Abe *et al.*, FERMILAB-Pub-97/284-E, submitted to Phys. Rev. Lett..

[7] E. Eichten, K. Lane, J. Womersley, Phys. Lett. **B405**, 305 (1997).

[8] M. Johnson *et al.*, DØ note 3168; "Preliminary Hardware Design for a Silicon Vertex Trigger", M. Johnson, February 1998.

[9] "Road Definition for Silicon Track Trigger", DØ note 3363, U. Heintz, 1998.

[10] reference for VRB.

[11] reference for VTM receiver module.

[12] "Dataflow Requirements for the Silicon Track Trigger", DØ note ????, M. Narain, 1998.
The expected averages are based on an average of 4 interactions/crossing, corresponding to $\mathcal{L} = 2 \times 10^{32}$ cm^{-2}s^{-1} and a crossing interval of 396 ns. The expected maximum is based on the endpoint of the spectra for 8 interactions.

[13] `http://d0sgi0.fnal.gov/˜foco/vertex/tbtalk.ps`

[14] `http://www.pa.msu.edu/hep/d0/l2/mbt.htm`

[15] DØ note 3338, M. Narain, S. Mani, 1998.

[16] John Marriner, private communication.

[17] "Interaction region measurements during run II startup", Paul Derwent, Brenna Flauger, CDF note 4059, March 1997.

[18] "Beam Motion and Feedback", Paul Derwent, CDF note 3859, August 1997.

[19] "Fitting the beam position with the SVX", Hans Wenzel, CDF note 1924, 1993.

[20] "The primary interaction vertex", Hans Wenzel, CDF note 4066, 1997.

[21] "Effect of primary vertex uncertainty on SVTpp efficiency", Frédéric Stichelbaut, December 1998, (`\http://ww-d0.fnal.gov/\˜wahl/minutes/`).

[22] "A preliminary consideration of hit finders for the DØ impact parameter trigger", Henryk Piekarz, DØ note 3209, March 1997.

[23] "The run II handbook", Fermilab Beam Division, `http://www-bd.fnal.gov/lug/runII_index.html`.

[24] "Benefits of a Silicon Vertex Trigger for ttbar Events", F. Stichelbaut, DØ note 3312, September 1997;
"Studies of SVT Trigger using SUGRA Four b Quark Signature", D. Norman, DØ note 3318, September 1997;
"Towards a Run II Trigger for $p\overline{p}$ to W Higgs", J. Hobbs, DØ note 3321, October 1997;
"Triggering on $Z \rightarrow b\overline{b}$ Events in Run II", M. Narain, DØ note 3339, October 1997.

[25] ISAJET ref

[26] PYTHIA ref

[27] COMPHEP ref

[28] GEANT ref

[29] "Benefits of a Silicon Vertex Trigger for $t\overline{t}$ Events", Frédéric Stichelbaut, DØ Note 3312.

[30] Update of the silicon track trigger studies for `events`, Frédéric Stichelbaut, DØ note 3444.

[31] "Studies of SVT trigger using SUGRA four b quark signature", Doug Norman, DØ note 3318.

[32] " Towards a Run II Trigger for $p\overline{p} \rightarrow WH^0 \rightarrow q\overline{q}b\overline{b}$", John Hobbs, DØ Note 3321

[33] "Primary vertex finding al Level 2 using the silicon vertex detector", Meenakshi Narain and Sudhindra Mani, DØ note 3338.

[34] "Triggering on $Z \rightarrow b\overline{b}$ events in run II", Meenakshi Narain, DØ note 3339.

[35] "Triggers for B physics in run II", Frédéric Stichelbaut, Andrzej Zieminski, Meenakshi Narain, DØ note 3354.

[36] "Single top and the silicon tracker trigger", Terry Heuring, DØ note 3400.

[37] "Silicon track trigger performance studies", Terry Heuring, Dan Karmgard, DØ note 3415.

[38] Private communication, Jerry Blazey and Cecilia Gerber. See instructions at `http://niuhep.physics.niu.edu/~blazey/ntuple_AUG96.html`

[39] "L1 axial CFT trigger hardware and firmware design for the baseline trigger algorithm, Fred Borcherding", DØ note 3082.

[40] "The impact of a displaced vertex trigger on the CFT L1 axial trigger, Fred Borcherding", DØ note

[41] "Central fiber tracker trigger simulation", Kathy Streets, Kin Yip, DØ note 3301.

[42] "An algorithm for a silicon track trigger", P. Gutierrez, E. Smith, H. Muramatsu, DØ note XXXX

[43] Meena Narain, algorithm and code presented during STT working group meetings.

[44] E. Boos, M. Dubinin, V. Ilyin, A. Pukhov, and V. Savrin, SNUT-94-116, hep-ph/9503280;
A.P. Heinson, A.S. Belyaev, E.E. Boos, *Phys. Rev.* **D56** 3114.

[45] S. Abachi *et al.*, Phys. Rev. Lett. **74**, 2632 (1995).

[46] F. Abe *et al.*, Phys. Rev. Lett. **74**, 2626 (1995).

[47] See, for example, E. Laenen, J. Smith, and W.L. van Neerven, Phys. Lett. **B321**, 254 (1994).

[48] "Status of single top analysis for $e\ +\ jets/\mu$" A. Belyaev *et al..*, DØ note 3441;
Jeff McDonald, private communication.

[49] See, for example, S. Dawson, *Nucl. Phys.* **B249**, 42 (1985); C.-P. Yuan, *Phys. Rev.***D 41**, 42 (1990); R.K. Ellis and S. Parke, *Phys. Rev.***D 46**, 3785 (1992); G. Bordes and B. van Eijk, *Nucl. Phys.* **B435**, 23 (1995).

[50] I. Bigi *et al.*, *Phys Lett.* **B181**, 157 (1986).

[51] For example, see R.S. Chivukula, S.B. Selipsky, E.H. Simmons, *Phys. Rev. Lett.***69**, 575 (1992); D. Carlson, E. Malkawi, C.-P. Yuan, *Phys. Lett.* **B337**, 145 (1994).

[52] Particle Data Group, O. Caso *et al.*., "Review of Particel Physics", *European Physics Journal* **C 3** 1 - 794 (1998).

[53] "Searches for New Particles at Present Colliders", P. Janot, in Proceedings from International Europhysics Conference on High Energy Physics, August, 1997.

[54] "Standard Model Higgs at LEP", Peter McNamara, invited talk at ICHEP98.

[55] "Searching for new particles at existing accelerators", Daniel Treille, invited talk at ICHEP98.

[56] "The TeV 2000 Report", D. Amidei and R. Brock, eds., FERMILAB-PUB-96/082.

[57] "The TEV33 Committee Report", Amidei, *et al.*, Feb. 1996.

[58] E. Eichten, K. Lane, J. Womersley, *Phys. Lett.* **B405**, 305 (1997).