

The D0 Silicon Track Trigger

Kenneth Herner*

State University of New York, Stony Brook

on behalf of the DØ Collaboration

E-mail: kherner@fnal.gov

The Silicon Track Trigger is a dedicated preprocessor within the Level 2 trigger system of the D0 detector. It is designed to provide fast pattern recognition using inputs from the central tracking detectors to search for high-impact parameter tracks. Triggering on high-impact parameter tracks enhances the heavy-quark content of the data, benefiting Higgs searches, Top physics, and b -physics.

17th International Workshop on Vertex detectors

July 28 - 1 August 2008

Utö Island, Sweden

*Speaker.

1. Introduction

The D0 Silicon Track Trigger (STT or L2STT) is a dedicated preprocessor in the Level 2 trigger system of the DØ detector. DØ is one of two general purpose detectors at the Fermilab Tevatron; it is described in detail in Refs. [1] and [2]. The STT uses inputs from two of the detector subsystems, the Central Fiber Tracker (CFT) and Silicon Microstrip Tracker (SMT). The CFT information comes to the STT via the Level 1 Central Track Trigger (L1CTT) [3], part of the Level 1 trigger system at DØ. Figure 1 shows an overview of the DØ trigger system, while Figure 2 shows a schematic of the Level 1 and Level 2 trigger systems, showing how the STT fits within the entire system.

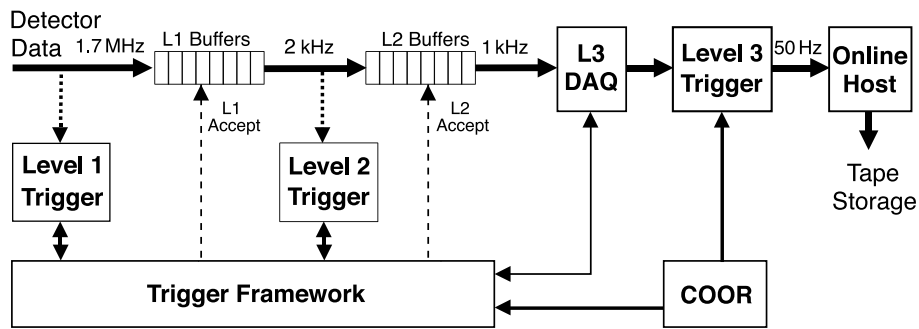


Figure 1: Overview of the DØ trigger system [2].

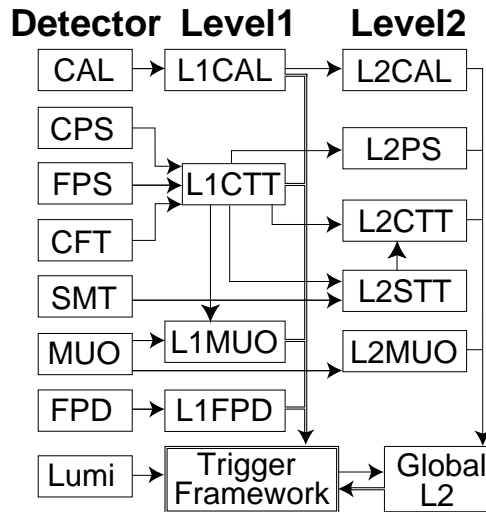


Figure 2: Diagram of the Level 1 and Level 2 trigger system [2].

The SMT is arranged in a barrel and disk geometry as shown in Figure 3. The barrels are cylinders consisting of five layers of silicon sensors concentric with the beam pipe. Each barrel is approximately 12 cm long in z . In the innermost layer of silicon there are in fact eight barrels, with the inner four barrels being approximately 6 cm long each in z such that the outer edges of the inner

POS (VERTTEX 2008) 025

four barrels roughly align with the outer edges of the innermost two barrels in the other layers. The STT does not use information from the disk detectors.

The overall concept of the STT is to match hits from the SMT barrel detectors to L1CTT tracks to provide fast identification of high-impact parameter tracks. Such tracks often come from displaced vertices, a trademark of b decays. By triggering on the presence of such tracks one can enhance the heavy flavor content of a data sample, important for Higgs searches, top physics, and b physics.




Figure 3: SMT layout [2]. The two outer H-disks were removed during the 2006 detector upgrade.

Figure 4 shows the STT pattern recognition concept. Tracks from L1CTT are used as seeds, or “roads”. SMT hits within $\pm 2\text{mm}$ of the road are considered for use in a fit. A “hit” is defined as a cluster of individual SMT channels that together pass a certain threshold (sec. 2.2.1). There must be suitable hits in at least four of the five SMT layers for the fit to proceed. The (r, ϕ_0) hit positions go into a χ^2 minimization function with the outputs being the track impact parameter b , radius of curvature κ , and ϕ_0 . We define ϕ_0 as the angle between the track’s tangent line and the x -axis at the point of closest approach to the origin, with a correction for the measured beam position from the previous data-taking run (the beam spot generally changes only slightly between runs) in order to get a more accurate impact parameter measurement.




Figure 4: STT pattern recognition. Clusters from SMT hits are associated to tracks from L1CTT; the STT then fits the SMT hits to a track by minimizing a χ^2 function in b , κ , and ϕ . [4].

The STT impact parameter resolution is a function of three main effects: the detector resolution, the Tevatron beam cross section size, and multiple scattering in the beam pipe and SMT. The detector resolution contributes about $19 \mu\text{m}$ to the total resolution, the beam size about $35 \mu\text{m}$, and the multiple scattering term is track p_T -dependent. The final total impact parameter resolution is

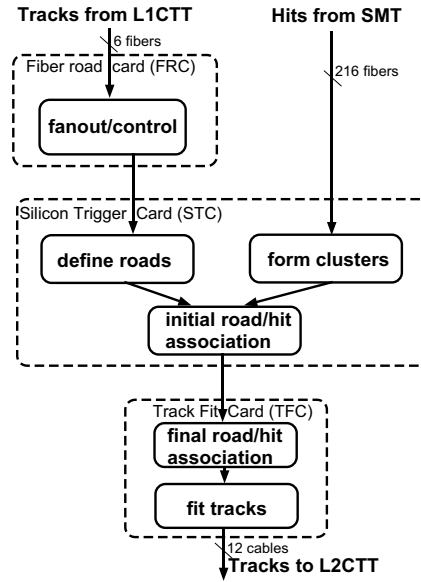


Figure 5: Dataflow through the STT. The Level 3 data paths have been omitted [4]

then

$$\sigma_b = \sqrt{35^2 + 19^2 + \left(\frac{54\text{GeV}}{p_T}\right)^2} \mu\text{m} \quad (1.1)$$

where p_T is the track p_T in GeV. There are other small effects such as the z position of the $p\bar{p}$ interaction and beam tilt (which affects the $x-y$ position of the beam, but these effects together are less than $5 \mu\text{m}$).

2. Technical Design

There are six STT readout crates at D0, each covering approximately 1/6th of the total ϕ ¹ space. A crate consists of one Fiber Road Card (FRC), ten Silicon Trigger Cards (STCs), and two Track Fit Cards (TFCs). If the trigger framework accepts the event at Level 1 (with an accept rate of 500-1500 Hz depending on instantaneous luminosity), the L1 CTT tracks come in through the FRC, where they pass to the TFCs and STCs. Raw SMT hits go directly to the STCs, which form the SMT clusters from the raw hits and associate the clusters to the CTT roads. The output passes to the TFCs where the final hit/road association is done and the MT clusters are fit to tracks. Finally the TFC output passes to the L2CTT preprocessor (shown in Figure 2) to take part in the trigger decision. On a L2 accept, the TFC output is also sent to Level 3 for bookkeeping and/or debugging purposes. Figure 5 shows the flow of data through the STT. The typical data volume from the CTT per readout crate is between 15 and 100 bytes per event, depending on the number of tracks. Typical latency for readout of SMT channels is $15 \mu\text{s}$, with an additional $5 \mu\text{s}$ for transfer. The system is designed to allow for $100 \mu\text{s}$ total latency but in practice the total latency is well under the limit.

¹ ϕ is the azimuthal angle in the plane transverse to the beam direction.

2.1 Fiber Road Card

The Fiber Road Card (FRC) functions as the link between the crate and the rest of the trigger system, executing the tasks sent by the trigger framework such as resets, initializations, and sending a copy of the STT output to Level 3 on L2 accepts. It also receives the L1CTT tracks in events where there is an L1 trigger accept, and sends them to the appropriate STCs and TFC(s) based on their ϕ position. It also serves as a conflict manager between the single board computer (SBC) in the crate and the in-crate CPU used for initialization.

Figure 6 is a block diagram of the FRC. Communication with the trigger framework is handled through the Serial Command Link Formatter (SCLF). The L1CTT tracks come in on a VME Transition Module (VTM) where the Trigger/Road Data Formatter (TRDF) receives the information and combines it with data from the SCLF. After a check for a bunch crossing (BX) number match between the two data sources the TRDF packages the CTT and needed SCLF data into a block consisting of between six and fifty-four 32-bit words, depending on the number of CTT tracks. Only the first 46 tracks (sorted in decreasing p_T) from the CTT are sent on in each of the six crates. The data block then goes to the first two of three PCI buses where Link Transmitter Boards (LTBs) send the block to the STC and TFC cards. The Buffer Manager (BM) in each FRC is responsible for generating error signals, passing initialization commands and monitoring data collection requests to the in-crate CPU, buffering the L3 data for transmission, and scheduling VME accesses to accomplish these tasks.

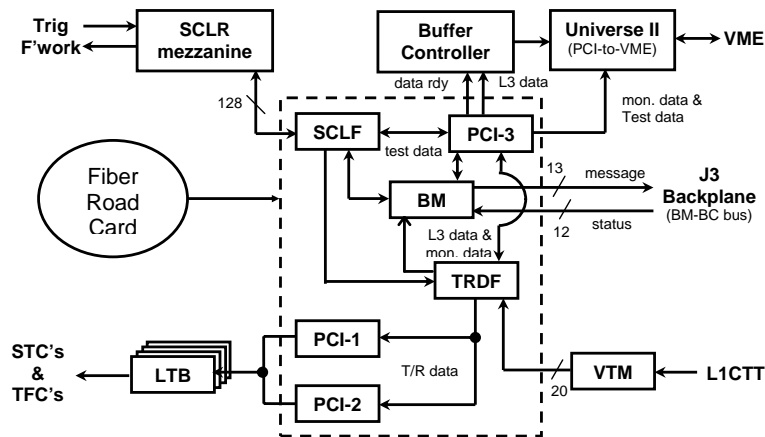


Figure 6: Block diagram of the Fiber Road Card (FRC). The dashed rectangle denotes the on-board FRC components [4].

2.2 Silicon Trigger Card

The STC receives SMT data over the HP G-link protocol into VTMs via optical fibers, and the L1CTT data from the FRC. Its tasks include performing pedestal subtraction on the SMT channels, removing dead or noisy channels, clustering the channels together, and finally associating those clusters to the roads defined by the L1CTT tracks. Figure 7 shows a block diagram of the STC.

Each of the ten STCs within a crate receives SMT input data from a unique VTM. Each VTM has four channels connected via optical fibers to four SMT input FIFOs, or SMT-IFs, as shown.

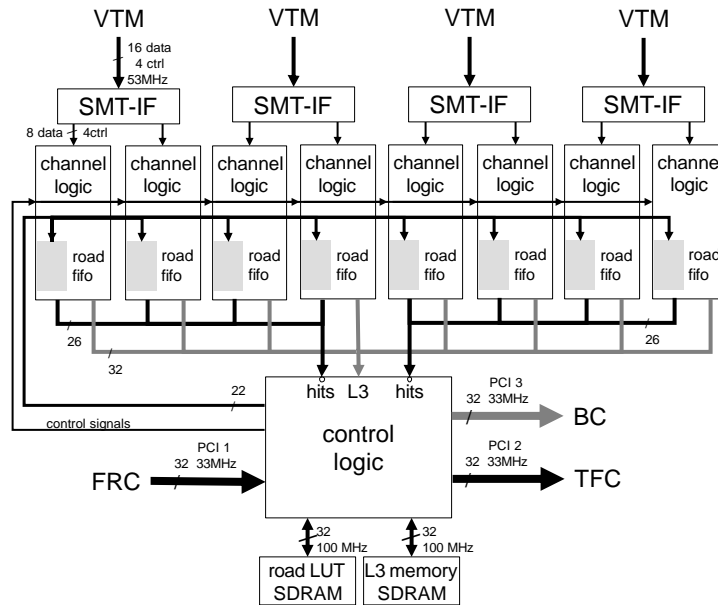


Figure 7: Block diagram of the Silicon Trigger Card (STC)[4].

Each fiber carries data for two detector elements. There are eight bits for each detector element, along with four control bits; thus there are 20 bits per fiber. The data from each SMT-IF are split into two streams of eight SMT plus four control bits, and go into channel logic blocks. The channel logic blocks perform the pedestal subtraction, clustering, and hit-road association. The clusters associated with roads are thereafter called “hits”. The control logic block is the heart of the STC; it manages the overall event processing, issuing control signals to the channel logic and the PCI interfaces. The input and output data pass through several PCI interfaces; the first PCI interface receives the L1CTT tracks from the FRC; the second takes the hits from the channel logic and transfers them to the TFC, while the third receives download-able parameters (such as SMT channel pedestals) and transfers the L3 data (clustering information for one SMT detector per STC for monitoring, or other detailed diagnostics if selected) to the Buffer Controller for output to L3.

2.2.1 Clustering and Hit-road Association

The core functions of the STC, the clustering and hit-road association, occur in the channel logic. The first thing that happens to the SMT data received by the SMT-IFs is pedestal subtraction. For each SMT channel the STC subtracts a certain number of ADC counts, known as the pedestal, to correct for noise effects. The pedestal for each channel is a download-able parameter that is loaded at initialization time. The pedestal-subtracted channels are also checked against a list of known bad channels (also downloaded at initialization); channels on this list have their ADC counts set to 0, effectively eliminating them. The channels then pass through a centroid-finding algorithm that compares neighboring strips and finds a centroid. The centroid position is then used as the cluster position. The total pulse height under the centroid is a measure of the energy deposited in the SMT and also goes to the TFC, but the dynamic range is scaled to fit into three bits.

All clusters’ centroid positions are compared to all roads to associate clusters to roads if the

cluster is within 2mm of a road. The comparison works by looking at two 11-bit addresses that correspond to the chip and strip numbers of the road edges as defined by that particular CTT track. If the centroid position, encoded in a 13-bit address, has a leading 11 bits that lie within the two addresses for the road, then that cluster gets associated to the road and is then called a “hit.” It is possible for a centroid to be associated with more than one road, and for roads to have multiple hits within each SMT layer. The associated hits then pass to the Track Fit Card (TFC) for final hit-road association and track fitting, where only one hit per layer is allowed.

2.3 Track Fit Card

Each STT crate has two Track Fit Cards. Each TFC receives input from the FRC and the STCs. Due to the detector geometry, there are three STCs per crate that send data to only the first TFC, three that send only to the second, and four that send to both. The TFC performs filtering on the initial hit/road association and chooses the final associations, fits the clusters, and outputs the result to L2CTT for use in the trigger decision.

Figure 8 shows the data flow through the various parts of the TFC. The input data from the STCs and the FRC comes in on the PCI-1 bus. The data from the FRC are an exact copy of the L1CTT input tracks, with header and trailer data added. The STC data include the centroid for each hit, the 3-bit pulse height, and an index referencing the CTT track to which the hit has been associated. The index has the same order as the L1CTT input tracks. The STC and FRC are initially held in Input Dual-Port Memory (IDPM), with each road and its associated clusters occupying consecutive memory locations. As shown in the figure the IDPM has two independent banks, one holding odd-numbered roads and the other holding even-numbered roads. The entire event is held in the IDPM until all STC clusters have been read in. During the read-in, the SMT cluster centroid information is converted into an effective r, ϕ position for use in the fitting by means of a coordinate conversion lookup table (CCLUT in figure 8.) This position then gets sent for fitting. The fits take place inside Digital Signal Processors (DSPs), with four DSPs assigned to each IDPM bank as shown in figure 9. We describe the fitting procedure in detail below. As fits are completed they go to Output Dual-Port Memory (ODPM) until all fits in an events are done. Then the output fits pass to the PCI-2 bus where they are formatted with D0 standard headers and trailers attached. The bus also synchronizes the event information, preserving the event time-ordering, and then the fits pass out over the Hotlink Transmitter, going to L2CTT.

2.3.1 Track Fitting

The first step in the track fitting is called final filtering, which finalizes the cluster/road association; only one cluster per SMT layer is chosen. In the first part of the final filtering the road width narrows from $\pm 2\text{mm}$ to $\pm 1\text{mm}$, eliminating clusters outside of 1mm. Then, starting from the outermost possible layer, an algorithm searches for combination of hits consistent with a straight line in the $r - z$ plane. Effectively, this means that the hits selected in each layer must all be within the same SMT barrel, or if there is a change, there can be only one change over all layers, and this change must not be more than one barrel. Clusters surviving these constraints pass to the second stage of the final filtering, which measures the $r_i \delta \phi$ distance between each cluster and the CTT

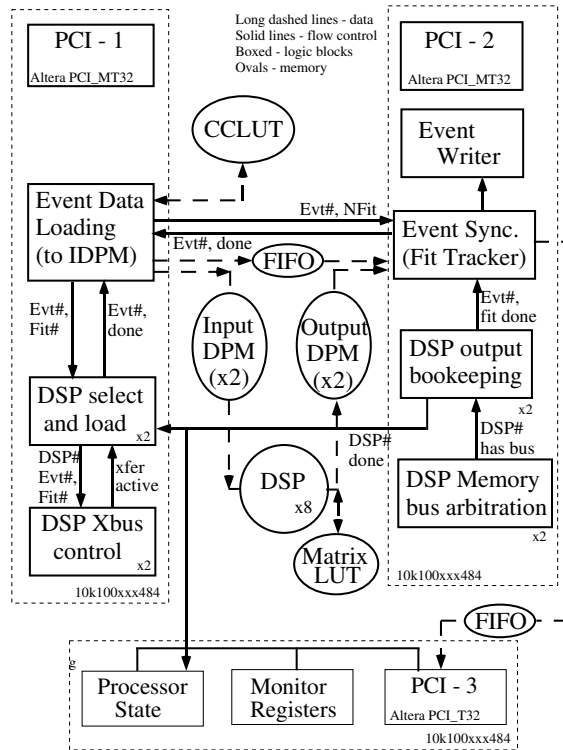


Figure 8: Dataflow through the Track Fit Card (TFC)[4].

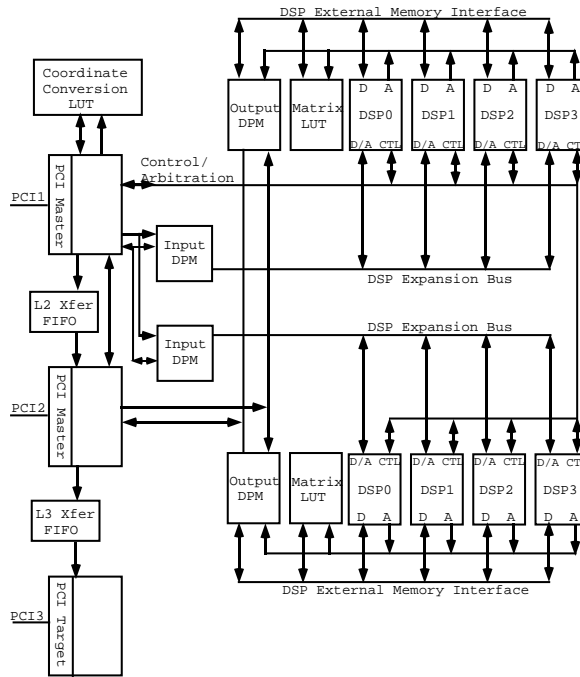


Figure 9: Block diagram of the Track Fit Card (TFC)[4].

POS (VERTTEX 2008) 025

track at radius r_i . $r_i \delta\phi$ is defined as

$$r_i \delta\phi = |r_i(\phi_i - \phi_{0CTT}) - \kappa_{CTT} r_i^2 - b_{CTT}| \quad (2.1)$$

where r_i and ϕ_i are the r and ϕ positions of the i -th cluster, and ϕ_{0CTT} and κ_{CTT} are the track azimuthal angle and curvature relative to the measured beam position. b_{CTT} is the CTT track impact parameter measured with respect to the detector origin. In each SMT layer, the chosen cluster is the one that has the smallest $r_i \delta\phi$ value.

After the final filtering there are SMT hits from either four or five layers. We use these, plus the CTT track (r, ϕ_0) position where it crosses the CFT A layer and H layer in the final track fitting. We assume the track has a circular trajectory in the r, ϕ plane near the origin. One needs three parameters to define such a circular trajectory; they are the signed impact parameter b , the track ϕ_0 , and curvature κ . We approximate a circle close to the origin by the expression

$$\phi(r) = \frac{b}{r} + \kappa r + \phi_0 \quad (2.2)$$

We base the impact parameter sign on whether the origin is inside or outside this circle. We say the product $b\kappa$ is positive when the origin is inside the circle, and negative when it is outside the circle.

We can use equation 2.2 to describe the fitted track trajectory. To obtain the track parameters we minimize the χ^2 function

$$\chi^2 = \sum_{clusters} \left[\frac{r_i \phi_i - r_i \phi(r_i)}{\sigma_i^2} \right]^2 \quad (2.3)$$

A linear least-squares minimization yields the equation for the three parameters of interest:

$$b = \sum_j M_{1j} \Phi_j \quad (2.4)$$

$$\phi_0 = \sum_j M_{2j} \Phi_j \quad (2.5)$$

$$\kappa = \sum_j M_{3j} \Phi_j \quad (2.6)$$

with $j=1,2,3$, $\Phi_n = \sum_k r_k^n \phi_k / \sigma_k^2$ ($k = 1, \dots, N_{points}$), and σ_k is the uncertainty on the (r, ϕ) position of the k -th hit (typically $\approx 15 \mu\text{m}$). All of the fitting steps take place in the DSPs. However, the only have support for integer adds and multiplies; they cannot do floating-point calculations. In addition, the DSP time budget is $100 \mu\text{s}$ per fit; it is not feasible to analytically solve these equations within the time budget as it would require calculating the M_{ij} matrices. If, however, we “re-center” the coordinate system such that the ϕ angles are expressed as differences from the ϕ value of one of the points (ϕ_1) we can rewrite Eqs. 2.4 through 2.6 as

$$b = \sum_k M'_{1k} \delta\phi_k \quad (2.7)$$

$$\phi_0 = \sum_k M'_{2k} \delta\phi_k \quad (2.8)$$

$$\kappa = \sum_k M'_{3k} \delta\phi_k \quad (2.9)$$

with $\delta\phi_k \equiv \phi_k - \phi_1$. Now the matrix elements and $\delta\phi_k$ values can be rescaled into 16-bit signed integers with ease. Additionally, we divide the ϕ space into 1,440 slices, and we use the same matrix values for each slice. This gives acceptable precision. Further, there are actually several thousand separate matrices for each ϕ slice to take into account the different possibilities of barrel combinations and whether or not four or five layers go into the fit. All of the matrices are computed offline and stored in a Matrix Lookup Table (Matrix LUT in figures 8 and 9.) Finally the track fitting works as follows: for each (r, ϕ) hit position, after ϕ conversion, the appropriate integer multiplications take place within the DSP and the Matrix LUT is accessed to obtain the appropriate elements. Through the use of an Inverse Matrix LUT, one gets back the fitted track b , κ , and ϕ_0 .

The final χ^2 value of the fit is a measure of quality. Since the approximations made in the fit neglect multiple scattering, we scale the fit χ^2 by the factor $1/\sqrt{4+(8/p_T)^2}$ to account for χ^2 degradation from multiple scattering. The constants were initially determine using Monte Carlo but were later tuned on data to provide 95% χ^2 acceptance at all momenta. The resulting χ^2 is then scaled to an integer ranging from 0 to 31. If a fit contains hits from all five layers but the χ^2 is 31, i.e. a bad fit, we drop the SMT hit with the highest individual contribution to the χ^2 and re-fit with the remaining four SMT hits. This is called “two-pass fitting.” Five-layer fits with $\chi^2 < 31$ are always preferred over four layer fits, however, even if a four-layer fit might have a lower χ^2 .

3. Performance

To measure the STT’s physics performance we match tracks from the STT L3 output to reconstructed offline tracks (RECO tracks), using the reconstructed tracks as the standard. If the STT is performing well then there should be an STT track overlapping the reconstructed track when the reconstructed track is within the STT fiducial acceptance. Since RECO tracks are not required to have SMT hits or lie in a certain η range ², we place the following cuts on RECO tracks before considering them as “good”, to match the STT design constraints:

- $p_T > 1.5$ GeV
- $|\eta| < 1.6$, i.e. within the CTT acceptance
- track $\chi^2 < 4$
- hits in ≥ 7 CFT axial layers
- hits in ≥ 4 SMT layers
- overlap with L1CTT track

These cuts ensure that the RECO track is within the STT’s acceptance and that conditions exist such that an STT track can be found, namely, a CTT track and hits in at least four SMT layers. They also protect against fake RECO tracks that would artificially lower the measured efficiency.

² $\eta \equiv -\ln \tan(\frac{\theta}{2})$, where θ is the polar angle with respect to the proton beam axis.

Good STT tracks are those STT tracks output to L3 that have $p_T > 1.5$ GeV and a fit χ^2 passing a chosen cut. We can then define the efficiency ε as

$$\varepsilon = \frac{N_{STT-RECO\ match}}{N_{good\ RECO}} \quad (3.1)$$

and conversely the purity P as

$$P = \frac{N_{STT-RECO\ match}}{N_{good\ STT}} \quad (3.2)$$

the match criterion is a χ^2 consisting of the track's impact parameter (measured with respect to the beam position), κ , and ϕ values. A low χ^2 implies that the tracks are similar. We measure the efficiency and purity in two different samples: in $Z \rightarrow \mu\mu$ events where we have clean, high p_T tracks, and in a sample dominated by multijet events, consisting of mostly low- p_T tracks.

3.1 Performance in $Z \rightarrow \mu\mu$ events

$Z \rightarrow \mu\mu$ events provide a good sample in which to measure the STT efficiency on high- p_T tracks. These are very often clean tracks with little else to degrade the resolution, so we should see a high efficiency here. To measure the efficiency in this sample we only consider the two RECO tracks matched to the two muons identified as the Z daughters. To get only these two tracks we impose additional requirements on the RECO tracks on top of those in section 3. We raise the p_T threshold to 15 GeV, and require that the invariant mass of the two tracks is consistent with the Z mass, namely $76 < M_{\mu\mu} < 106$ GeV. We also modify the good STT track definition by raising the p_T threshold to 15 GeV as well. The left-hand plot in Figure 10 shows the STT efficiency as a function of the STT track fit χ^2 in $Z \rightarrow \mu\mu$ events. Here the track fit χ^2 is the integer fit χ^2 value divided by 2. We see that for a χ^2 cut of 5, a typical value used in a trigger, the STT is 85% efficient, and 90% if one does not cut on the fit χ^2 . The efficiency does not reach 100% due in part to the inefficiencies of the SMT itself. Some SMT channels, while considered for offline analysis, are not considered in the STT online triggering due to unacceptably high readout time. These disabled channels represent less than 5% of the total. There is also a 2% loss due to tracks that cut across the boundary between the two TFCs in each readout crate, as the TFCs do not work together when fitting tracks.

3.2 Performance in multijet events

It is also instructive to measure the efficiency in multijet events, dominated by low p_T tracks. In this case we define good STT and RECO tracks according to the standard definitions given above. The right-hand plot in Figure 10 shows the efficiency vs. purity for several different cut values of the STT track χ^2 . The data sample is a typical run with no detector problem and a low to intermediate instantaneous luminosity ($\approx 6 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$.) We see that for a scaled χ^2 cut value of 5, the STT is 64% efficient with 87% purity. The efficiency increases to 70% with 84% purity for no χ^2 cut.

3.3 Use in trigger design

The STT has several possible uses in the trigger. In some triggers one simply requires an STT track without regard to fit quality or impact parameter. In this case the idea is to use the silicon

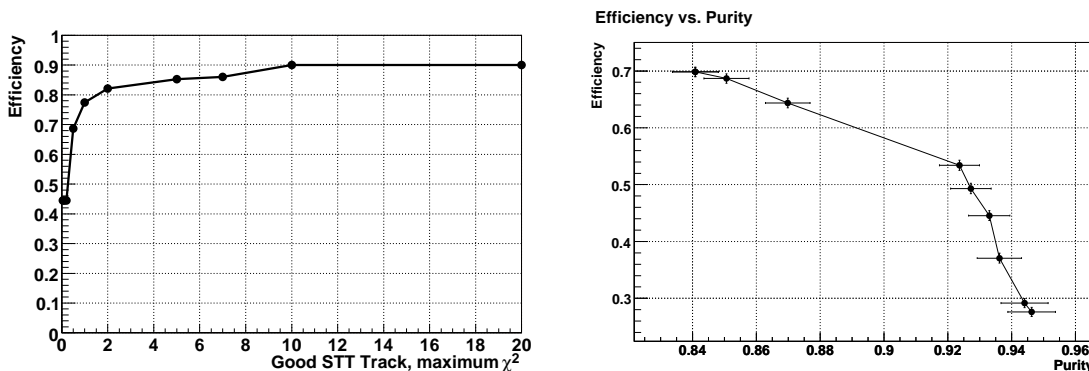


Figure 10: Left: Efficiency for a good STT track to be matched to a good RECO track as a function of the STT track fit χ^2 [4] in $Z \rightarrow \mu\mu$ events. Right: Efficiency for a good STT track to be matched to a good RECO track as a function of the STT track fit in multijet events. The different point represent different values of the STT track χ^2 . Looser cuts have low purity and high efficiency, while tighter cuts have lower efficiency but higher purity.

information to improve the momentum resolution on the track. In other triggers, there is a cut on the STT track's impact parameter significance, defined as b/σ_b . This is an excellent way to increase the heavy flavor content of the triggered sample. One such example in the D0 trigger list is a dedicated multijet trigger. At Level 2, the trigger requirement is an OR of four separate requirements. Two of the four terms require an STT track with impact parameter significance ≥ 3 , $\chi^2 < 5.5$, and $p_T > 5$ GeV. As we saw in equation 1.1, the resolution σ_b is of the order $40 \mu\text{m}$, meaning a significance cut of 3 would pick out tracks with an impact parameter of at least $120 \mu\text{m}$ at that resolution. For comparison, simulated tracks from b -quarks in $ZH \rightarrow \nu b \bar{b}$ events showed an average impact parameter of $80 \mu\text{m}$ [4]. By requiring an STT track in the trigger, one can enhance the b -jet content of the sample while avoiding the use of other topological variables involving the jets that lower the trigger efficiency.

4. Summary

The D0 Silicon Track Trigger is a dedicated preprocessor in the D0 Level 2 trigger system. It is designed to take inputs from the central tracking detectors and identify high-impact parameter tracks with a correction for the beam position. It has high efficiency and is a useful tool to trigger on event containing heavy quarks, a benefit to several areas of the D0 physics program.

References

- [1] D0 Collaboration, S. Abachi et al., *The D0 Detector*, Nucl. Instrum. Methods Phys. Res. **A 338**, 185 (1994).
- [2] D0 Collaboration, V.M. Abazov et al., *The Upgraded D0 Detector*, Nucl. Instrum. Methods Phys. Res. **A 565**, 463 (2006).
- [3] J. Anderson et al., *The D0 Central Track Trigger*, IEEE Trans. on Nucl. Sci. **51**, 345 (2004).
- [4] T. Adams et al., *The D0 Run II Impact Parameter Trigger*, physics/0701195.