Real Time Secondary Vertexing at CDF

J. Adelman [a], A. Annovi [b], M. Aoki [c], A. Bardi [b], J. Bellinger [d], M. Bitossi [b], M. Bogdan [a], R. Carosi [b,1], P. Catastini [b], A. Cerri [e], S. Chappa [f], M. Dell'Orso [b], B. Di Ruzza [h], I. Furic [a], P. Giannetti [b], P. Giovacchini [b], T. Liu [f], T. Maruyama [c], I. Pedron [g], M. Piendibene [b], M. Pitkanen [f], B. Reisert [f], M. Rescigno [h], L. Ristori [b], H. Sanders [a], L. Sartori [g], M. Shochet [a], B. Simoni [b], F. Spinella [b], S. Torre [b], R. Tripiccione [g], F. Tang [a], U. K. Yang [a], A. M. Zanetti [i]

[a] *Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA*
[b] *INFN, Sezione di Pisa, I-56127 Pisa, Italy*
[c] *University of Tsukuba, Tsukuba, Japan*
[d] *University of Wisconsin, Madison, WI 53706, USA*
[e] *LBNL, Berkeley, CA 94720, USA*
[f] *Fermilab, Batavia, IL 60510, USA*
[g] *INFN, Sezione di Ferrara, I-44100 Ferrara, Italy*
[h] *INFN, Sezione di Roma, I-00173 Roma, Italy*
[i] *INFN, Sezione di Trieste, I-34012 Trieste, Italy*

Abstract

The Online Silicon Vertex Tracker (SVT) is the trigger processor dedicated to the 2-D reconstruction of charged particle trajectories at the Level 2 of the CDF trigger. As the Tevatron luminosity rises, multiple interactions increase the complexity of events and thus the SVT processing time, reducing the amount of data CDF can record. The SVT upgrade aims to increase the SVT processing power to restore at high luminosity the original CDF Data Acquisition capability. In this paper we review the tracking algorithms implemented in the SVT and we report on the first step in the SVT upgrade.

Key words: Vertexing; Trigger; Online Tracking; Data Acquisition: Real time pattern recognition; Position-sensitive detectors
PACS: 07.05.Hd, 29.40.Gx

One challenge in the Collider Detector at Fermilab (CDF) experiment [1] is to extract signals of interest efficiently from much larger backgrounds. The total inelastic $p\bar{p}$ cross-section at the Tevatron is about 50 mb, while the b-quark cross-section within CDF's acceptance (transverse momentum $p_T > 6$ GeV/c, rapidity $|y| < 1$) is about 10 μb, and the t-quark cross-section is about 5 pb. At luminosities above 0.35×10^{32} cm^{-2}s^{-1}, the mean number of interactions per beam crossing exceeds

[1] Corresponding author. E-mail address: roberto.carosi@pi.infn.it (R. Carosi).

1. Reducing the 1.7 MHz beam-crossing rate to the 70 Hz DAQ output rate implies a trigger rejection of 25000.

Good background rejection in the trigger requires the fast identification of distinctive signatures. In the CDF trigger, many important signatures exploit fast charged-particle track reconstruction in the bending plane of the spectrometer, transverse to the beam axis.

CDF uses a three-level trigger. On each beam-crossing (396 ns), the entire front-end digitizes. A 5.5 μs pipeline of programmable logic forms axial drift chamber tracks and can match these with calorimeter and muon-chamber data. On Level 1 accept, front-end boards store the event to one of four buffers. Level 2 processing, with about 30-40 μs latency, adds fast offline quality silicon tracking and calorimeter clustering. The final Level 2 decision is made in software on a single-board computer to provide full flexibility in the final selection. At Level 3, a farm of 250 commodity PCs runs full event reconstruction. This is the first stage at which three-dimensional tracks (e.g. for invariant mass calculation) are available. Events passing Level 3 are written to disk.

Output rates at L1/L2/L3 are approximately 25000/350/70 Hz. Silicon tracking at L2 allows CDF to collect large samples of fully hadronic bottom and charm decays, by requiring two drift chamber tracks at L1, requiring each track to have a significant (at least 100 μm) transverse impact parameter at L2, and performing full software tracking at L3 to confirm the hardware tracking.

CDF's Level 1 drift chamber hardware track processor, XFT [2], with a latency of 1.9 μs, finds tracks of $p_T > 1.5$ GeV/c with 96% efficiency. XFT's resolutions, $\sigma(p_T) = 0.017 \cdot p_T^2$ GeV/c and $\sigma(\phi_0) = 5$ mrad, are only a factor 10 coarser than those of the offline reconstruction.

For each event passing Level 1, the Silicon Vertex Trigger (SVT) [3] associates each XFT track with silicon hit data from four detector planes, and produces a transverse impact parameter measurement of 35 μm resolution (50 μm when convoluted with the beam spot) with a mean latency of 25-30 μs, 9 μs of which spent waiting for the first silicon data. SVT's impact parameter resolution for $p_T \approx 2$ GeV/c is comparable to that of offline tracks that do not use Layer 00 (mounted on the beam pipe), which is not available in SVT. The resolutions on the other transverse parameters are $\sigma(p_T) = 0.003 \cdot p_T^2$ GeV/c and $\sigma(\phi_0) = 1$ mrad.

The SVT is a system of 150 VME boards containing FPGAs, RAMs, FIFOs, and one ASIC design. CPUs are used only for initialization and monitoring.

Three key features allow SVT to carry out in 15 μs a silicon track reconstruction that typically requires about 0.1 s in software: a highly parallel/pipelined architecture, custom VLSI pattern recognition, and a linear track fit in fast FPGAs.

The silicon detector's modular, symmetric geometry tends itself to parallel processing. SVT's first stage, converting a sparsified list of channel numbers and pulse heights into charge-weighted hit centroids, processes 12×6×5 (azimuthal × longitudinal × radial) silicon planes in 360 identical FPGAs. The overall structure of SVT reflects the detector's 12-fold azimuthal symmetry. Each 30° azimuthal slice is processed in its own asynchronous, data-driven pipeline that first computes hit centroids, then finds coincidences to form track candidates, then fits the silicon hits and drift chamber track for each candidate to extract circle parameters and a goodness of fit.

In SVT's usual configurations, a track candidate requires a coincidence of an XFT track and hits in four (out of five available) silicon layers. To define a coincidence, each detector plane is divided into bins ("superstrips") of programmable width, about 500 μm. For each 30° slice, the set of 32000 most probable coincidences ("patterns") is computed offline in a Monte Carlo program and loaded into 256 custom VLSI associative memory (AM) chips. For every event, each binned hit is presented in parallel to the 256 AM chips, and the hit mask for each of the 128 patterns per chip is accumulated in parallel. When the last hit has been read, a priority encoder enumerates the patterns for which all layers have a matching hit. The processing time is thus linear in the total number of hits in each slice and in the number of matched patterns. Actually, the last matched pattern is read out only a few clock cycles after the hit readout is completed.

There is no exact linear relationship between the transverse parameters (curvature c, azimuthal an-

gle ϕ, transverse impact parameter d) of a track in a solenoidal field and the coordinates at which the track meets a set of detector planes. However for $p_T > 2$ GeV/c, $|d| < 1$ mm, $|\phi| < 15°$, a linear approximation is valid within a few percent. By linear regression to Monte Carlo data, we derive the coefficients and the intercepts relating the transverse parameters to the four silicon hits and the two parameters (curvature and azimuthal angle) measured by XFT, c_{XFT} and ϕ_{XFT}. The same regression produces the coefficients and intercepts corresponding to the fit's 3 degrees of freedom, with which we calculate the three constraints and the χ^2.

In the start-of-run download, all the constants are precomputed at the edge of each pattern and stored. For each pattern, the fitter board computes corrections to the parameters and constraints with respect to the pattern edge, using 8-bit multiplication in 6 parallel FPGAs, in 300 ns per fitted track. Tracks passing a goodness-of-fit cut $\chi^2 < 15$ propagate downstream.

The most critical parameter provided by the SVT is the transverse impact parameter with respect to the beam axis. The SVT is supposed to work with the beam in its nominal position, i.e., parallel to the z axis of the CDF detector and at $x = 0$ and $y = 0$. In practice, some misalignments and time variations of the beam position are possible, thus, corrections are needed. The beam position in the transverse plane can be calculated using the correlation between d and ϕ. If the beam-spot position in the transverse plane is (x_0, y_0), the relationship between d and ϕ for primary tracks is $d = -x_0 \sin(\phi) + y_0 \cos(\phi)$. The impact parameter with respect to the position of the beam is $d' = d + x_0 \sin(\phi) - y_0 \cos(\phi)$. The average beam position and the impact parameter correction are performed online by fitting data. Figure 1 shows the $d - \phi$ correlation for online track candidates which satisfy a cut $\chi^2 < 15$.

The SVT processing time increases as the instantaneous luminosity of the Tevatron increases. Multiple interactions in the same accelerator bunch crossing increase the number of silicon hits, and thus increase the time to process all of them. A more serious problem is the increase in the number of track candidates to be fitted. In each pattern,

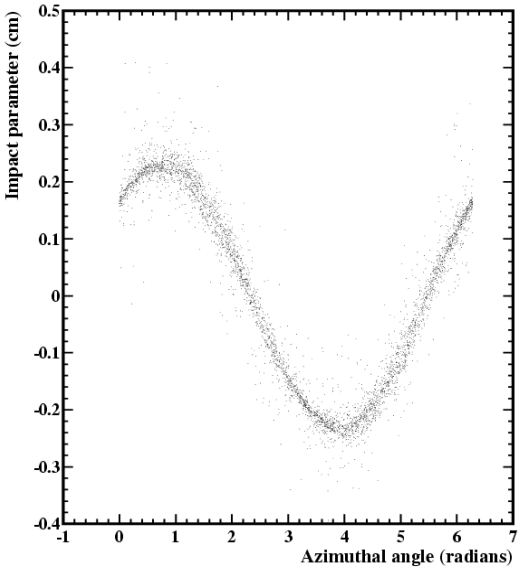

Fig. 1. Impact parameter versus azimuthal angle for candidate tracks with $\chi^2 < 15$.

the number of required fits (combinations) is the product of the number of hits in each silicon layer. As the hit density in the silicon increases, the number of fits can become quite large.

The SVT upgrade [4] aims to both reduce the number of fits and perform each fit more quickly. The latter is achieved by increasing the speed of the track fitting boards. To achieve the former, the maximum number of patterns per azimuthal slice is increased from 32000 to 128000 in a first step, and to 512000 in a second step. This allows a factor two smaller width for the superstrips without reducing the pattern coverage and the SVT efficiency. Smaller superstrips contain a smaller number of silicon hits and therefore a smaller number of track candidates processed by the track fitter.

The increase in the number of patterns is achieved with a new AM system (AM++) [5], based on a redesigned and more powerful AM board. The new track fitter boards (TF++), can handle the larger number of patterns and are implemented on "Pulsar" hardware [6], mostly porting the FPGA algorithm of the old board. The Pulsar operates at a 70 MHz clock speed

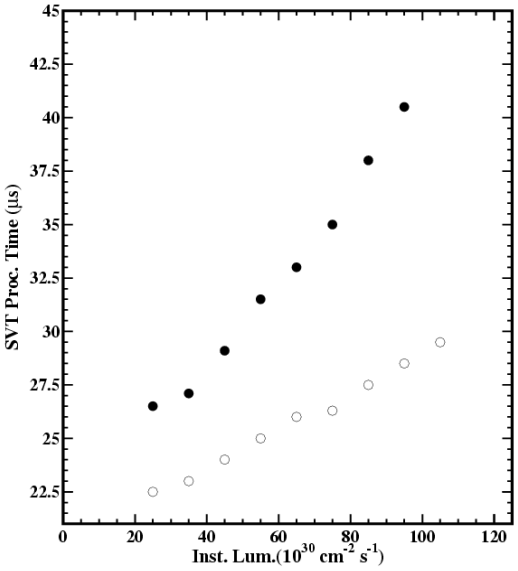

Fig. 2. Average SVT processing time versus instantaneous luminosity, before (full dots) and after (open circles) the upgrade

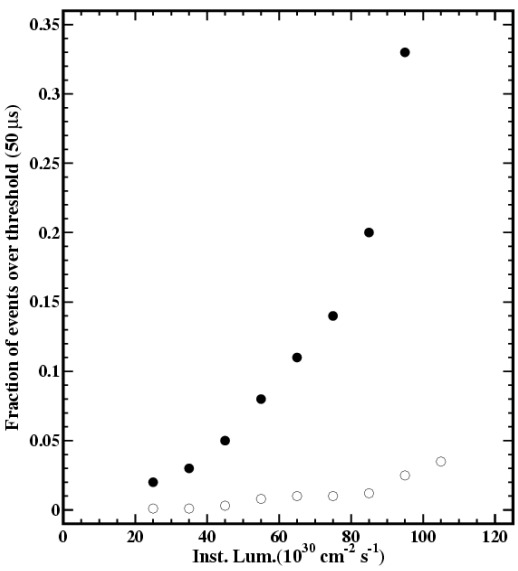

Fig. 3. Fraction of events with SVT processing time > 50 μs versus instantaneous luminosity, before (full dots) and after (open circles) the upgrade

(vs. 30 MHz for the old board), thus significantly reducing the overall processing time.

The first step in the SVT upgrade uses a partial increase in the number of patterns, to 128000 per azimuthal slice. This provides a smaller improvement than the full increase at 512000 patterns, but still significant. Figure 2 shows the average SVT processing time at different Tevatron instantaneous luminosities, before the upgrade (full dots) and after the first step (open circles).

The gain in the SVT processing time is about 12 μs at about 10^{32} cm^{-2} s^{-1} instantaneous luminosity. This gain is due to the reduced number of fits per event (7 μs) and the faster electronics (5 μs). Figure 3 shows the fraction of events requiring more than 50 μs for the SVT processing. Because of the limited event buffering in the data acquisition at Level 2, fluctuations in the processing time are important, and are the main cause of Level 2 dead time, as shown by simulation programs. We observe a fractional tail of 3% at 10^{32} cm^{-2} s^{-1} instantaneous luminosity. For comparison, the result from a detailed simulation of the full upgraded SVT (with 512000 patterns) at $3{\cdot}10^{32}$ cm^{-2} s^{-1} is 6.6%.

The reduction of both the average and the fractional tails of the SVT processing time has allowed an increase of the Level 1 output rate from 20 to 25 kHz without any increase of the dead time (5%).

References

[1] R. Blair, et al., CDF Collaboration, FERMILAB-PUB-96-390-E.

[2] E. J. Thompson, et al., *IEEE Trans. Nucl. Sci.* **49** (2002) 1063.

[3] W. Ashmanskas, et al., *Nucl. Instr. and Meth. A* **518** (2004) 532; W. Ashmanskas, et al., *IEEE Trans. Nucl. Sci.* **49** (2002) 1177; R. Amendolia, et al., *IEEE Trans. Nucl. Sci.* **39** (1992) 795.

[4] J. Adelman, et al., 2005 IEEE Nuclear Science Symposium, 26-29 October 2005, Puerto Rico.

[5] A. Bardi, et al., 2005 IEEE Nuclear Science Symposium, 26-29 October 2005, Puerto Rico.

[6] K. Anikeev, et al., *IEEE Trans. Nucl. Sci.* TNS-00650-2004.R2, accepted for publication on Sept. 2005.