Initial experience with the CDF SVT trigger

B. Ashmanskas[a], A. Barchiesi[b], A. Bardi[c], M. Bari[d], M. Baumgart[a],
Stefano Belforte[d,*], J. Berryhill[a], M. Bogdan[a], R. Carosi[c], A. Cerri[e],
G. Chlachidze[f], R. Culberston[a], M. Dell'Orso[c], S. Donati[c], I. Fiori[g], H. Frisch[a],
S. Galeotti[c], P. Giannetti[c], V. Glagolev[f], A. Leger[h], Y. Liu[h], E. Meschi[c],
L. Moneta[h], F. Morsani[c], T. Nakaya[a], G. Punzi[c], M. Rescigno[b], L. Ristori[c],
H. Sanders[a], S. Sarkar[b], A. Semenov[f], M. Shochet[a], T. Speer[h], F. Spinella[c],
H. Vataga[c], X. Wu[h], U. Yang[a], L. Zanello[b], A.M. Zanetti[d]

[a] *Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA*
[b] *INFN, Sezione di Roma I and University La Sapienza, I-00173 Roma, Italy*
[c] *INFN, University and Scuola Normale Superiore of Pisa, I-56100 Pisa, Italy*
[d] *INFN, Sezione di Trieste, I-34012 Trieste, Italy*
[e] *Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*
[f] *Joint Institute for Nuclear Research, Dubna, Russia*
[g] *University of Padova and INFN, Sezione di Padova, I-35031 Padova, Italy*
[h] *University of Geneva, CH-122 Geneva 4, Switzerland*

Abstract

The Collider Detector at Fermilab (CDF) Silicon Vertex Tracker (SVT) is a device that works inside the CDF Level 2 trigger to find and fit tracks in real time using the central silicon vertex detector information. SVT starts from tracks found by the Level 1 central chamber fast trigger and adds the silicon information to compute transverse track parameters with offline quality in about 15 μs. The CDF SVT is fully installed and functional and has been exercised with real data during the spring and summer 2001. It is a complex digital device of more than 100 VME boards that performs a dramatic data reduction (only about one event in a thousand is accepted by the trigger). Diagnosing rare failures poses a special challenge and SVT internal data flow is monitored by dedicated hardware and software. This paper briefly covers the SVT architecture and design and reports on the SVT building/commissioning experience (hardware and software) and on the first results from the initial running.
© 2002 Elsevier Science B.V. All rights reserved.

PACS: 07.05.Hd; 29.40.Gx

Keywords: Trigger; Data acquisition; Real time pattern recognition; Tracking; position sensitive detectors

1. The CDF upgrade and the Silicon Vertex Tracker

The Collider Detector at Fermilab (CDF) is a general purpose detector for the study of high

energy proton–antiproton interactions produced in the Fermilab Tevatron Collider particle accelerator. CDF completed its first period of data taking (Run1) in 1994. The Tevatron has now been upgraded (Run2) to higher energy (from 800 to 1000 GeV per beam) and luminosity (from 2×10^{31} to 1×10^{32}). CDF has been upgraded as well to cope with the new accelerator conditions (higher instantaneous luminosity and reduced bunch spacing from 3.5 μs to 132 ns) and extend its physics reach.

Four upgrades are relevant to this paper: (1) A 3 level digital trigger: Level 1 is completely pipelined and deadtime-less, completes in 4 μs and has an accept frequency $\simeq 50$ kHz, the Level 2 operates after detector readout and completes in ~ 20 μs with accept rate 200–300 Hz, Level 3 will work after event building to bring the accepted rate down to ~ 50 Hz for data storage and offline analysis. (2) A new silicon vertex detector (SVXII [1]) with 5 layers of micro-strip sensors arranged in a 12-fold azimuthal geometry (ϕ) and segmented in 6 longitudinal barrels (3 mechanical units, each read out at both ends) along the z-axis. (3) A new fast trigger processor (XFT [2]) for the central drift chamber that finds tracks in the transverse (r–ϕ) projection in time to be used in the L1 trigger decision and with a measured resolution, of $\delta\phi = 4$ mrad, $\delta P_t/P_t = 1.6\%$. (4) The SVT (Silicon Vertex Tracker) device [3], used in the L2 trigger to combine XFT tracks with SVXII data to fit tracks with offline-like quality. The design goals for the SVT, from simulation of the algorithm on Run1 data, are $\delta\phi \simeq 1$ mrad, $\delta P_t/P_t \simeq 0.3\%$ and $\delta d \simeq 35$ μm, where d is the track *impact parameter* i.e. the minimum distance between the particle trajectory and the primary vertex in the transverse plane.

By providing precision information on the impact parameter, the SVT allows CDF to select events containing long lived particles (*B*-baryons in particular) already in the trigger process. This novel feature, available for the first time in a hadron collider experiment, greatly enriches the CDF physics program, in particular uniquely provides access to hadronic B decays like $B^0 \to \pi\pi$ or $B_s \to D_s\pi \to \geqslant 2$ hadrons. A typical trigger path will require two tracks with $P_t \geqslant 2$ GeV at L1 from XFT with a ~ 30 kHz expected accept rate that will be reduced to ~ 30 Hz at L2 after including impact parameter cuts, typically: 1 mm $> d > 100$ μm.

The SVT is a parallel pipelined data driven device. Tracks are found independently in the twelve $30°$ sectors of the SVXII (wedges). Raw SVX data flow from the front-end to the Hit Finder boards that find clusters of strips with a significant energy deposit and compute the coordinates of the centroids (Hits). The P_t and ϕ of tracks from XFT are received, fanned out to the 12 SVT wedges and fed both to the Associative Memory and to the Hit Buffer boards together with all SVX Hits. The Associative Memory finds track candidates (Roads) using a coarse spatial resolution in the SVX (250 μm) and sends them to the Hit Buffer. Each Hit Buffer stores all hits and tracks in a wedge for each event in an internal memory, then for each Road output by the Associative Memory, retrieves the hits and the XFT track belonging to that road and sends them to the Track Fitter boards, that perform quality cuts and estimate track parameters using the full available spatial resolution in a linearised fit. A set of Merger boards, each providing 4-way fan-in and 2-way fan-out, ties together all boards in a chain that starts from 288 optical links from SVX and ends in a single data path to the CDF Level 2 processor.

Data flow through the SVT pipeline on uniform point-to-point connections implemented via 23-bit LVDS data channels running at about 700 Mbit/s each. The SVT is made of over one hundred $9U \times 400$ mm VME boards housed in 8 crates, connected by about 50 of those data paths and by 13 short busses on custom backplanes. Up to 4 events may be present at the same time in the SVT pipeline, with data being stored on on-board FIFOs. The output list of high precision tracks is sent to the CDF L2 processors for the final decision on the average 15 μs after the L1 decision, depending on the event size.

SVT construction has been completed and all hardware is installed and fully operational since the beginning of 2001.

2. Operational challenges

The SVT design has been fully simulated before construction, using both the CAD digital simulation of all boards, and the simulation of the tracking algorithm using the CDF data from Run1. Nevertheless there are issues that could not be studied with simulation and can only be addressed when operating the system with beam data, e.g. a full system test and the z-alignment.

During full system test we focused on spotting possible hardware, firmware or programming problems that may cause the SVT to malfunction in rare situations. While all boards had been extensively tested before installation, those bench tests were limited to either very low rate (limited by how fast a computer can feed data to the SVT) or very small data variety (using the SVT boards to send a preloaded set of data in a loop at full speed). Moreover one has to make sure of correctness and proper timing of input data from SVXII and XFT. Since only about $\frac{1}{1000}$ of the events processed by SVT are available on tape, we cannot rely on offline analysis only. Most SVT diagnostics is carried out in real time by dedicated hardware that continuously monitors the data integrity. Special care has been taken to allow quick reduction of system problems to single boards or links: each board verifies that all input data come from the same event and checks data parity on the input cables, all internal FIFOs are monitored for overflows (data loss) and data flowing on each link are continuously copied to large circular memories (Spy Buffers) providing a snapshot of all data processed in the SVT in the last few hundred events. These data are then compared with the SVT simulation by processes running on the VME crate controllers. Operation of these memories is coordinated via one dedicated Spy Control board in each crate. Whenever an erroneous or illegal condition is detected, error flags are asserted locally in the boards and inside the data stream.

The z-alignment issue is related to the fact that the SVT only fits tracks in the transverse plane (although separately in each of the six z-barrels). Therefore the z-misalignment of the detector, either internal or relative to the beam line, spreads the beam profile, causing a worsening of the resolution, a higher trigger rate for a given impact parameter cut, and a higher background contamination. To make that spread small compared to the natural beam width, each SVXII strip and the beam line must be all parallel within 100 μrad. Assembly of the SVXII barrels met (even exceeded) this specification, more challenging is the SVX-beam relative alignment: the detector must be properly positioned in the collision hall, and the beam orbit must be steered in each store to be put on the same axis, and kept there. During Run1 the beam angle at the collision point showed variations beyond the required limits, the plan for Run2 is to use SVT feedback to continuously adjust the beam position. For this SVT must provide accurate beam position information in the x–y plane in each of the 6 barrels with a time constant of about 1 s.

3. SVT performance with data

The SVT has been operated since one year, in the October 2000 CDF commissioning run (prototype SVX detector, no XFT tracks to SVT) and in the 2001 physics run period that started in April. Initial focus has been on the system integration: debugging the SVX and the XFT communication to the SVT, developing control and monitor software infrastructure, interfacing to DAQ start-up and error handling, debugging the L2 processors. This activities are not completed yet, and while SVT data checking and diagnostics features are proving to be very useful in commissioning the surrounding systems, it has not been possible yet to use the SVT to provide trigger decisions and measure trigger rates. The SVT has processed data routinely, with its output track list logged in the event data, so the tracking performance of the SVT could be measured. Good tracking resolution is the first and most important requirement for SVT operation, therefore it was the object of most initial effort. Indeed, if the detector-algorithm combination does not provide the needed precision, there is no point in looking at efficiency, stability, rates, background, and so on. Especially with regards to the tracking resolution, the SVT

device performed remarkably well even without any parameter tuning and in spite of having operated with wide roads, nominal detector geometry, large cut on track χ^2, no correction for z-tilt between SVXII and the beam, no correction for SVX dead/hot channels, pedestals, gains, no tuning of the clustering algorithm.

Fig. 1 shows the d–ϕ plot from data taken on August 2001. Even when SVT was used in an unforeseen situation, without XFT tracks, acting as standalone tracker in the innermost 4 SVXII layers, it clearly separated tracks coming from the beam (the sinusoidal wave reflects the d–ϕ relationship when the beam is offset in the x–y plane with respect to the origin of the coordinate system). Vertical stripes in the left plot of Fig. 1 indicate silicon hot channels. The tracks are of course much cleaner when XFT information could be added (right plots). From distributions like the one in Fig. 1 it is possible to extract the impact parameter resolution. A typical d distribution is shown in Fig. 2 on data collected in one run in July 2001. The width of the distribution ($\sim 50\ \mu m$) is the convolution of the SVT resolution with the intrinsic beam spread. This distribution is not corrected for difference between ideal and real

detector geometry and for the beam z-tilt. The beam tilt is measured in real time by a task running on the SVT VME crate controller, that reads the


Fig. 2. Impact parameter distribution (in cm) using XFT and SVX data.

Fig. 1. Impact parameter (vertical, cm) vs. polar angle (horizontal, rad). Left: using only SVX data. Right: using both SVX and XFT: uncorrected (top), after correcting for beam offset fit on top plot (bottom). The region without data points around $\phi = 4$ rad on the right plots is due to HV problems in a sector of the drift chamber.

track list from the SVT Spy Buffers and fits the d–ϕ distribution to obtain the $\{x_0, y_0\}$ beam position in each of the six barrels, a linear fit to this 6 pairs returns then the beam line direction. This proved to work smoothly and reliably providing the information needed for the online beam orbit steering. In this way the z-beam slope was found to be $\delta x/z \simeq 600$ μrad, $\delta y/z \simeq 150$ μrad, in agreement with offline studies using the central drift chamber and well beyond the SVT specification. The deviation is also larger then the beam orbit adjustment range and thus moving the SVXII detector is required. Using the offline reconstruction to compute the z of each hit it is possible to correct for the beam tilt. This study allowed to single out the contribution to the impact parameter resolution from all sources and showed that present data, once correct for z-tilt and detector alignment, demonstrate a 45 μm impact parameter resolution (convolution of SVT resolution and beam width), in agreement with expectations from the Run1 data.

4. Learning from the SVT experience

During the SVT construction, installation and commissioning, we did not encounter completely unexpected problems or difficulties, but, as it happens usually, some items remained under-manned, first priority being to have the hardware working and installed in time. Now we can identify among the large list of things that we would have liked to do, but did not, the ones that would have been most useful.

For what concerns the hardware, it is very important never to stop the design process until the very last moment. We benefitted a lot from early prototyping of the SVT boards,[1] but also learnt that it is very useful to delay freezing the project as much as possible in order to use state of the art technology to add flexibility and simplify the system by using less, more uniform and powerful components. Major drawbacks in our

case have been: a large number of obsolete/overcrowded/non-reprogrammable FPGAs/CPLDs; not all internal pipeline stages accessible for diagnostics; difficult/impossible to add useful new features; not enough common hardware solutions and common behaviour among different boards, hence the need for many hardware experts and some operational difficulty. We tried to uniform board designs, but should have pushed this more. Nevertheless the hardware construction has been largely a success, failure rate has been very small, and moving from bench test to real system operation exposed only a few minor problems which have been fixed with small firmware changes.

Looking at the SVT architecture and overall planning, the single most important lesson is: plan from the beginning for debugging and commissioning, to make this part fast, effective and easy. A key point is to allow one to use the system in a completely realistic way even without the rest of the detector/DAQ, using the same control and monitor software used during data taking. We made it possible to inject data into the SVT, process them in the boards, and read the results in a standalone mode. However, the data flow in this mode could not be synchronised with the DAQ/Trigger timing. During commissioning we realised this shortcoming and had to improvise a way to synchronise the flow of test data with the CDF DAQ/Trigger without using SVX and/or XFT as data sources. Only a fraction of SVT hardware can be exercised in this mode at present, but we are working to extend such a capability to the full system. We have also learned that it would have been very useful to complement SVTs asynchronous SpyBuffer monitoring mechanism with the option to read intermediate data into the event record, both to use uniform offline tools to access internal and external data and to tie the internal data to the specific events on tape.

On the good side there are several features built into the SVT which proved to be very useful. The common data communication protocol and capability to inject/read data from any SVT board allows to easily recombine SVT boards in various configuration (often simply via software) to accommodate many different tests both of the

---

[1] Many glitches fixed on SVT prototypes could not be found by CAD simulation and would have been disastrous had they only been discovered at installation time.

SVT and neighbour systems. The data integrity and synchronisation controls in SVT helped to find and fix low rate problems even in other systems. The capability to operate SVT boards independently of the DAQ/trigger framework makes it possible to develop, test and repair boards away from the CDF control room with no (or minimal) additional hardware. Finally, although not explicitly planned for, it turned out to be very easy, and immensely useful, to operate SVT when a subset of the detector was missing, allowing the detector and the trigger commissioning to proceed in parallel.

In conclusion, the Silicon Vertex Tracker has been successfully built, installed and operated in CDF and it performs as expected. Tracks are found and fit with the required precision and the beam position feedback to the accelerator in real time is possible. While the trigger rates still have to be measured and the device parameters optimised, the impact parameter resolution is already as good as expected and is sufficient for a successful operation of the trigger.

The commissioning and the debugging of such a complex device as the SVT has shown the importance of flexible re-programmable hardware, uniform simple components and data integrity control. In addition it has become clear that the design must include from the early stages the functionalities needed to debug, test and commission of the hardware, at the same level as what is needed for the steady operation of the system.

References

[1] J.A. Valls, Nucl. Phys. B 78 (1999) 311.
[2] C. Ciobanu, et al., IEEE Trans. Nucl. Sci. 46 (1999) 933.
[3] S. Belforte, et al., IEEE Trans. Nucl. Sci. 42 (1995) 860.