

The Silicon Vertex Trigger upgrade at CDF[☆]

J. Adelman^a, A. Annovi^b, M. Aoki^c, A. Bardi^b, M. Bari^b, J. Bellinger^d, M. Bitossi^b,
 M. Bogdan^a, R. Carosi^b, P. Catastini^b, A. Cerri^e, S. Chappa^f, M. Dell'Orso^b, B. Di Ruzza^g,
 Ivan K. Furić^{a,*}, P. Gianetti^b, P. Giovacchini^b, T.H. Liu^f, T. Maruyama^c, I. Pedron^h,
 M. Piendibene^b, M. Pitkanen^f, B. Riesert^f, M. Rescigno^g, L. Ristori^b, H. Sanders^a, L. Sartori^b,
 M. Shochet^a, B. Simoni^b, F. Spinella^b, S. Torre^b, R. Tripiccone^h, F. Tang^a,
 U.K. Yang^a, A.M. Zanettiⁱ

^aEnrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA

^bINFN, University and Scuola Normale Superiore of Pisa, I-56100 Pisa, Italy

^cUniversity of Tsukuba, Tsukuba, Japan

^dUniversity of Wisconsin, Madison, WI 53706, USA

^eLawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

^fFermi National Accelerator Laboratory, Batavia, IL 60510, USA

^gINFN, Sezione di Roma, I-00173 Roma, Italy

^hINFN, Sezione di Ferrara, I-44100 Ferrara, Italy

ⁱINFN, Sezione di Trieste, I-34012 Trieste, Italy

for the CDF Collaboration

Available online 29 November 2006

Abstract

The Silicon Vertex Trigger (SVT) in the CDF experiment at Fermilab performs fast and precise track finding and fitting at the second trigger level and has been a crucial element in data acquisition for Run II physics. However, as 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 the original CDF DAQ capability at high luminosity. We describe the SVT upgrade, consisting of a new Associative Memory 16 times larger than the existing one, and new faster Track Fitter and Hit Buffer boards to take advantage of these patterns. We describe the existing system, the upgrade, tests and performance.

© 2006 Published by Elsevier B.V.

PACS: 07.05.Hd; 29.40.Gx

Keywords: Vertexing; Trigger; Online tracking; Data acquisition; Real-time pattern recognition; Position-sensitive detectors

1. Introduction

Precision real-time tracking in CDF [1] is motivated by selecting events with b quarks in the trigger. Historically, events containing b quark decays were selected by identifying the lepton from the b decay. Utilizing precision real-time tracking and vertexing allows for the recognition

of track and vertex displacement. This increases the efficiency of selecting events containing b quark decays, and also allows sampling hadronic decays of b quarks, which were inaccessible when triggering on leptonic decays.

Real-time pattern recognition and precision tracking is already implemented within the CDF trigger system, and is known as the Silicon Vertex Trigger (SVT) [2]. The precision of the SVT track fitting system already approaches that of offline tracking routines, while the processing time is at the level of tens of microseconds. Three key features allow the SVT to perform precision tracking so quickly: a highly parallel and pipelined

[☆]Talk presented at the 10th Pisa Meeting on Advanced Detectors, La Biodola, Isola d'Elba, May 21–27, 2006.

*Corresponding author.

E-mail address: ikfuric@hep.uchicago.edu (I.K. Furić).

architecture, custom Very Large-Scale Integration (VLSI) chip-based pattern recognition and a Track Fitter implemented with Field-Programmable Gate Array (FPGA) technology.

Upgrades to the CDF trigger are necessary as the Tevatron luminosity increases. The increasing luminosity spawns more complex events and increases trigger processing time, in particular in the track trigger. This reduces the amount of data that CDF can record. A large upgrade program is underway to upgrade the CDF trigger in order to maintain its performance at instantaneous luminosities of $3 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. In this program, the SVT upgrade plays a central role.

2. CDF trigger overview

CDF uses a three-level trigger. On each beam crossing (132 ns), the entire front end digitizes (the silicon detector samples and holds). A $5.5 \mu\text{s}$ 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 (the silicon detector digitizes and transmits to the silicon trigger and event builder). Level 2 processing, with about $30 \mu\text{s}$ latency, adds fast silicon tracking, calorimeter clustering and EM calorimeter shower-max data. The final Level 2 decision is made in software on a commodity PC CPU, so a wider range of thresholds and derived quantities is possible (e.g. transverse mass of muon track pairs), even for information that is in principle available at Level 1. Upon Level 2 accept, front-end VME crates transmit to the event builder. 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. The maximum output of the upgraded system at L1/L2/L3 are, respectively, 25,000/1000/100 Hz. Drift chamber tracking is performed at Level 1, and silicon tracking is performed at Level 2.

CDF collects large samples of fully hadronic bottom and charm decays by requiring two drift chamber tracks at Level 1, requiring each track to have a significant (at least $120 \mu\text{m}$) impact parameter at Level 2, and performing full software tracking at Level 3 to confirm the hardware tracking. CDF Level 1 drift chamber hardware track processor, XFT [3], is a cornerstone of the CDF trigger, and is undergoing its own upgrade project [4]. For every bunch crossing, with $1.9 \mu\text{s}$ latency, it finds tracks of $p_T > 1.5 \text{ GeV}/c$ with 96% efficiency. The XFT obtains coarse hit data (two time bins) from each axial drift chamber wire, finds line segments in the 12 measurement layers of each axial superlayer, then links segments from these four superlayers to form track candidates. The XFT resolutions, $\sigma(1/p_T) = 1.7\%/\text{GeV}/c$ and $\sigma(\phi) = 5 \text{ mrad}$, are only about a factor of 10 coarser than those of the offline reconstruction.

3. SVT track processing

For each event passing Level 1, the SVT swims each XFT track into the silicon detector, associates silicon hit data from four detector planes and produces a transverse impact parameter measurement of $35 \mu\text{m}$ resolution ($50 \mu\text{m}$ when convoluted with the beam spot) with a mean latency of $24 \mu\text{s}$, $9 \mu\text{s}$ of which is spent waiting for the first silicon data. The SVT impact parameter resolution for $p_T = 2 \text{ GeV}/c$ tracks is comparable to that of offline tracks that do not use information from the Layer00 [5] system (layer of silicon mounted on the beam pipe), which is unfortunately not available for SVT processing. For fiducial online muon tracks from J/ψ decays, having $p_T > 2.0 \text{ GeV}/c$ and hits in the four silicon planes used by the SVT, measured SVT efficiency is 90%. The most suitable definition of efficiency in a given context depends on what one aims to optimize: relaxing the requirements on which layers contain offline silicon hits reduces the efficiency to 70%.

The SVT is a system of 150 custom 9U VME boards containing FPGAs, RAMs, FIFOs and one ASIC design. CPUs are used only for initialization and monitoring. SVT's input comprises 144 optical fibers, 1 Gbit/s each, and one 0.2 Mbit/s LVDS cable; its output is one 0.7 Mbit/s LVDS cable. The silicon detector's modular, symmetric geometry lends itself to parallel processing. The SVT's first stage converts a sparsified list of channel numbers and pulse heights into charge-weighted hit centroids, and processes $12 \times 6 \times 5$ (azimuthal \times longitudinal \times 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 and finally fits the silicon hits and drift chamber track for each candidate to extract circle parameters and a goodness of fit. In SVT's preupgrade configuration, a track candidate requires a coincidence of an XFT track and hits in a specified four (out of five available) silicon layers. To define a coincidence, each detector plane is divided into bins of programmable width (superstrips), typically $250 - 700 \mu\text{m}$ wide, and XFT tracks are swum to the outer radius of the silicon detector and binned with 3 mm typical width. For each 30° slice, the set of 32K 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 five layers have a matching hit. The processing time is thus linear in the total number of hits in each slice and linear in the number of matched patterns.

4. High-luminosity challenge

With Tevatron improvements, the luminosity of collisions increases. In the last year of running, a significant

fraction of the time, the detector was routinely gathering data at luminosities above $1.5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. The mean number of $p\bar{p}$ collisions per beam crossing increases with luminosity. Multiple collisions tend to mimic signal trigger signatures. The cross-section for multiple collisions faking trigger signatures increases with the mean number of collisions per beam crossing. Consequently, trigger rates increase in a rapid non-linear fashion with increasing luminosity. At the same time, the accepted events become more complicated, leading to increasing SVT track processing times.

For sufficiently complex events, the trigger system can still be occupied processing previous events in a full pipeline when another event gets accepted, leading to data loss. This condition is known as trigger dead time. Using a parametrized simulation of the increase of event processing time with luminosity, we have estimated the trigger dead time at collision luminosity of $3 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. By design, the acceptable trigger dead time due to Level 2 processing is 5%. The trigger system would incur the maximum acceptable trigger dead time at $3 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ and an input rate of only 13 kHz, effectively halving the rate of useful physics events from SVT-based triggers. The challenge of the SVT upgrade is to decrease the track processing time while maintaining the same physics performance (track finding efficiency, displacement resolution). Increasing the number of patterns from 32 to 512 K allows using patterns with smaller superstrip width. Narrower patterns contain fewer silicon hits and therefore reduce the number of hit combinations that are found and need to be fitted. Upgrading track processing boards to increase the clock frequency makes individual track fits faster. The clock frequency of the upgraded boards is 70 MHz, while the original boards operated at roughly 30 MHz. The upgraded system has to be able to process events at collision luminosity of $3 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ and Level 2 input rate of 23 MHz, while inducing a dead time of 5% or less.

The SVT upgrade is designed to require a minimum of new hardware. The increase in the number of roads is provided by the upgraded AM chip (AM++) [6]. The AM chip is capable of storing 128 patterns per chip; 4096 AM chips would be necessary to store information about 512 K patterns, which was not realistic. Instead, the AM++ chip is developed to store 5120 patterns on the same area (roughly 1 cm^2). The AM board was also upgraded to interface with the new AM++ chips. The upgraded AM boards (AMS/RW) have enough space to double the number of mounted chips. The AMS/RW boards also remove redundant pattern coincidences which occur when patterns are generated allowing that a hit is not present in one of the layers of the silicon detector. Two upgraded AM boards, each containing 64 chips, are used to store the 512 K pattern bank for a single wedge of the SVT.

The central board used in the rest of the SVT upgrade is the Pulsar board [7]. The Pulsar board is a general purpose 9U VME interface board. The board is equipped with

virtually all CDF signal connectors, and three Altera APEX 20K400BC-652-1XV [8] FPGAs. Each FPGA is coupled to a $128 \text{ K} \times 36$ synchronous-pipelined Burst SRAM equipped with No Bus Latency™ logic. Two of the FPGAs provide interfaces to two mezzanine cards each. The capacity of the on-board memories was insufficient to store 512 K patterns, so two mezzanine cards were developed with memories capable of storing $4 \text{ M} \times 48$ [9] and $512 \text{ K} \times 24$ [10] bits of data, respectively. The presence of SVT data interfaces on the Pulsar board drastically speeds up the upgrade process, since development is reduced to FPGA programming.

5. SVT upgrade results

The modular design of the SVT allowed for a staged upgrade. In the first stage, the AM system and the Track Fitter were upgraded. The pattern bank was expanded to 128 K patterns. In the second stage, the pattern bank was expanded to its full capacity of 512 K patterns. The Hit Buffer board was upgraded at the same time. All stages of the upgrade were tested in data taking with the CDF detector. During commissioning, all installed boards were required to bitwise match offline simulation outputs at the level of one million processed events.

Measuring the performance of the entire system in data taking is somewhat non-trivial. The instantaneous luminosity is continuously dropping during collisions. CDF's trigger system takes advantage of this by dynamically allocating more bandwidth to time-intensive trigger paths, such as those involving the SVT. Consequently, the effect of an upgrade can reliably be estimated by comparing SVT processing time before and after the upgrade at the same luminosity, with the same trigger path mixture. Fig. 1 shows the improvement on the fraction of events with processing time above $50 \mu\text{s}$ as a function of luminosity, for different stages of the SVT upgrade. This fraction of events

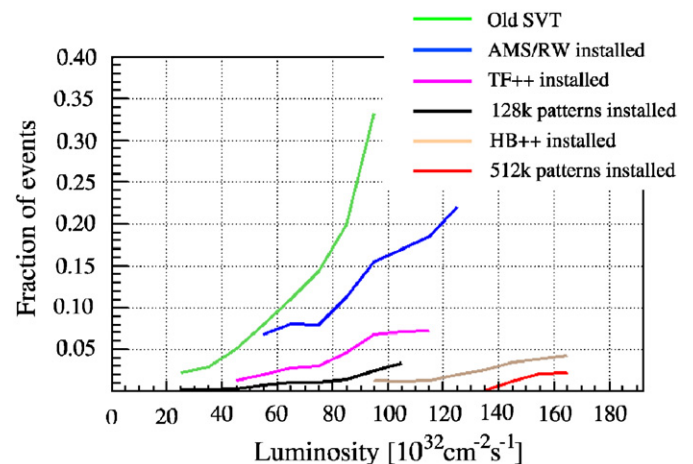


Fig. 1. Contribution of various incremental upgrades to the reduction of the fraction of events with processing time over the $50 \mu\text{s}$ threshold.

is interesting because over-threshold events directly contribute to trigger dead time.

The first improvement, installing the AMS/RW boards, reduced processing time by reducing the number of track fit candidates and reducing the pattern recognition time. The second upgrade, installing the upgraded Track Fitter (TF++) board, significantly reduces the fraction of over-threshold events by speeding up the track fitting process with faster clocks and a six-fold increase in the number of fitting engines on the new board. Next, the use of 128 K patterns reduces the number of fit combinations per recognized pattern. The upgraded Hit Buffer (HB++) further increased the processing speed by virtue of the faster clock speed on the upgraded board. Finally, the full power of the upgrade is visible after enabling all 512 K patterns. The fraction of events over threshold is well below 5% at the highest luminosities available for these tests. Data taking without an upgraded SVT system at these luminosities would clearly suffer huge rate penalties, as the corresponding fraction of events over threshold is roughly 25% at half the maximum tested luminosity, with a steeply rising tendency. The SVT upgrade has been a clear success in the tested range of luminosities; we are looking

forward to evaluating the performance of the system at $3 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

References

- [1] R. Blair, et al., CDF Collaboration, FERMILAB-PUB-96-390-E, 2006.
- [2] W.J. Ashmanskas, CDF Collaboration, et al., Nucl. Instr. and Meth. A 518 (2004) 532;
W.J. Ashmanskas, CDF Collaboration, et al., IEEE Trans. Nucl. Sci. NS-49 (2002) 1177;
R. Amendolia, CDF Collaboration, et al., IEEE Trans. Nucl. Sci. NS-39 (1992) 795.
- [3] E.J. Thomson, CDF Collaboration, et al., IEEE Trans. Nucl. Sci. NS-49 (2002) 1063.
- [4] See I. Fedorko, et al., CDF Collaboration, these proceedings.
- [5] C. Hill, CDF Collaboration, et al., Nucl. Instr. and Meth. A 530 (2004) 1.
- [6] A. Annovi, et al., CDF Public Note 7339, 2006.
- [7] K. Anikeev, CDF Collaboration, et al., IEEE Trans. Nucl. Sci. NS-53 (2006) 653 (<http://hep.uchicago.edu/~thliu/projects/Pulsar>).
- [8] ALTERA Pub., APEX 20K programmable logic device family, Data Sheet v5.1, 2004.
- [9] (http://edg.uchicago.edu/~tang/Memory/sram_M4M.html).
- [10] (http://edg.uchicago.edu/~tang/Memory/sram_M512K.html).