

First Steps in the Silicon Vertex Trigger upgrade at CDF

J. Adelman¹, A. Annovi², M. Aoki³, A. Bardi⁴, F. Bedeschi⁴, S. Belforte⁵, J. Bellinger⁶, E. Berry¹, M. Bitossi², M. Bogdan¹, M. Carlsmith⁶, R. Carosi⁴, P. Catastini⁹, A. Cerri⁸, S. Chappa⁷, W. Chung⁶, M. A. Ciocci⁹, F. Crescioli², M. Dell'Orso², B. Di Ruzza¹¹, S. Donati², I. Furic¹, S. Galeotti⁴, P. Giannetti⁴, C. M. Ginsburg⁶, P. Giovacchini⁴, R. Handler⁶, Y. K. Kim¹, J. D. Lewis⁷, T. Liu⁷, R. Mahlum⁷, T. Maruyama³, F. Morsani⁴, G. Ott⁶, I. Pedron¹⁰, M. Piendibene⁴, M. Pitkanen⁷, L. G. Pondrom⁶, G. Punzi², B. Reiser⁷, M. Rescigno¹¹, L. Ristori⁴, H. Sanders¹, L. Sartori¹⁰, F. Schifano¹⁰, F. Sforza⁹, M. Shochet¹, B. Simoni², F. Spinella⁴, P. Squillacioti⁹, F. Tang¹, S. Torre⁹, R. Tripiccion¹⁰, G. Volpi⁹, U. K. Yang¹, L. Zanello¹¹, A. M. Zanetti⁵

¹University of Chicago, Illinois, USA, ²University of Pisa, Italy, ³University of Tsukuba, Japan, ⁴INFN Sezione di Pisa, Italy, ⁵INFN Sezione di Trieste, Italy, ⁶University of Wisconsin, USA, ⁷Fermilab, Batavia, Illinois, USA, ⁸LBL, California, USA, ⁹University of Siena, Italy, ¹⁰University of Ferrara and INFN, Italy, ¹¹University of Rome and INFN, Italy

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 at high luminosity the original CDF DAQ capability. We describe the first steps in the SVT upgrade, consisting of a new Associative Memory with 4 times the number of patterns, and a new Track Fitter to take advantage of these patterns. We describe the system, its tests and its performance.

I. INTRODUCTION

CDF has a powerful multi-level trigger [1]. At high instantaneous luminosity, the Level-2 Trigger processing time can limit the Level-1 accept rate, which cannot be increased without increasing the dead time above 5%, the maximum allowed in CDF data taking.

The Level-2 trigger is organized as a 3-stage pipeline consisting of only 4 buffers. Large fluctuations in the Level-2 (L2) processing time can fill all four buffers and cause deadtime. The need to keep the deadtime low imposes a severe limit on the Level 2 input rate.

The task of the Silicon Vertex Trigger (SVT) [2] is very complex, and its processing time is a significant fraction of the total Level-2 processing time. Triggers requiring the SVT, mostly dedicated to studying b-physics (B-triggers), were the major source of deadtime before this upgrade.

The SVT links hits from five layers of the Silicon Vertex Detector (SVX) to tracks found in the central drift chamber (COT) to reconstruct complete tracks in real time precisely enough to measure b-quark decay secondary vertices. The COT tracks are reconstructed by XFT, the CDF Level-1 track processor [3]. The SVT task proceeds through steps. First, the position of detected charge centroids (hits) is determined from the coordinate of all fired channels in each detector layer. Pattern recognition consists of associating *hits* to track candidates at coarse resolution (*roads*), after which fits of all the possible combinations of hits from different layers within a road are performed, and the track parameters precisely determined.

SVT processing time increases as the instantaneous luminosity rises. Multiple interactions in the same accelerator bunch crossing increase the number of hits in the SVX, and thus increase the time to process all of the silicon hits (time/hit = ~35 ns). An additional, more serious problem is an increase in the number of track candidates to be fit (time/fit = ~300 ns). In each road, 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 aims to both reduce the number of fits and perform each fit more quickly. To achieve the former, the width of roads is reduced by increasing the total number of roads per azimuthal wedge from 32×10^3 (32k) to 512×10^3 (512k). The latter is achieved by increasing the speed of the track fitting hardware (Track Fitter). Table 1 shows the expected performance improvement with these upgrades. The increase in the amount of data CDF can record is quite significant. These results were obtained with a simulation model trained to predict the SVT performance at high luminosity. Data were taken at three luminosities between $1 \times 10^{31} \text{ cm}^{-2} \text{ sec}^{-1}$ and $5 \times 10^{31} \text{ cm}^{-2} \text{ sec}^{-1}$ with a two-track B-meson trigger. The hit and combination distributions were extrapolated to $3 \times 10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$ and used as input to a queuing program already used in the past to understand the rate limitations of the multi-level CDF trigger system. The program accounts for the timing of each stage of the trigger and includes effects due to the limited number of L2 buffers and the effects of the L2 upgrade [4]. Table 1 shows the maximum Level-1 trigger rate corresponding to a 5% deadtime. The validity of these results was tested at current luminosities. The rates in the table have uncertainties of approximately $\pm 30\%$ due to the extrapolation of a factor of 6 in luminosity.

Current SVT	Upgrade SVT
13 kHz	23 kHz

Table 1: The maximum Level-1 trigger rate at $3 \times 10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$ instantaneous luminosity so that the Level-2 deadtime does not exceed 5%.

The first step in the SVT upgrade uses a partial increase in the number of roads, to 128×10^3 per wedge (128k). This provides a smaller improvement than the full increase to 512k, but is still significant. Table 2 shows that using 128k roads reduces the average number of executed fits by a factor of 2.

# of stored roads	32k	128k
Typical road width (μm)	~ 500	~ 300
Average # of fits	37.1	15.4

Table 2: The typical road width at each SVX layer and the average number of fits at $1 \times 10^{32} \text{ cm}^{-2} \text{ sec}^{-1}$ instantaneous luminosity with 32k (old SVT) and 128k (1^{st} stage SVT upgrade) roads.

II. HARDWARE DESCRIPTION

Figure 1 shows the upgraded SVT. Boards with highlighted titles were replaced during the first step of the upgrade. In SVT, raw silicon hits are transmitted on optical fibers to Hit Finder boards that find hit clusters in each silicon layer and calculate their centroids. Merger boards (MRG) combine the silicon data from 3 Hit Finders (HF) with tracks in the drift chamber found by the XFT processor [3] working at Level-1. The combined information is sent to both the Associative Memory system (AM) [5], which does the pattern recognition, and a Hit Buffer (HB) [6] that stores hits until a pattern is found with sufficient hits to be deemed a track candidate. There are three pattern recognition boards per wedge: one Sequencer (AMS) that controls the operation and two AM boards (AMB) that find track candidates. The Sequencer converts silicon hit coordinates into coarser resolution superstrips, typically $500 \mu\text{m}$ wide, a resolution appropriate for pattern recognition. The AM boards contain AM chips, which are content addressable memories. The AM finds roads containing a drift chamber

track and the required number of silicon hits and passes them to the HB board.

As each road enters the HB, the silicon hits are retrieved at full resolution and passed to the Track Fitter (TF) along with the road number. The Track Fitter performs a linear fit to track candidates, and outputs tracks that pass goodness-of-fit cuts. The recently added Road Warrior (RW) [7] removes, prior to fitting, duplicate track candidates that arise because the road list include roads with one missing silicon hit. The RW has been implemented on the CDF Pulsar board [8], which was designed for and is now being used in the CDF Level-2 upgrade. The Pulsar design is very flexible, with room for mezzanine cards, various input and output protocols and three powerful FPGAs for data manipulation.

The SVT upgrade is designed to require a minimum of new hardware. The increase in the number of roads is achieved with a new AM system - the AM++ [9] - based on a redesigned AMS and one more powerful AMB. The new Sequencer uses a Pulsar board. Firmware for the FPGAs and simple mezzanine cards containing memory chips are the major needs for the AMS [10], which also carries out the RW function. By moving the RW into the AMS, time is saved in the HB on roads that would be later removed by the RW. Finally, the first step of the upgrade includes a new Track Fitter board (TF++), which can handle the larger number of patterns. The TF++ uses a Pulsar board with the same mezzanine cards as the AMS. The TF++ boards are implemented on Pulsar hardware, mostly porting the FPGA algorithm of the old TF [11]. The Pulsar operates at a 70 MHz clock speed (vs. 30 MHz for the old TF), thus significantly reducing the overall processing time.

An important feature of electronics used in high-energy physics is the ability to debug and trace errors that appear in data. For this purpose every input and output of each SVT board has circular memories maintaining a parallel image of the channel's data flow. If an SVT board finds an error, all

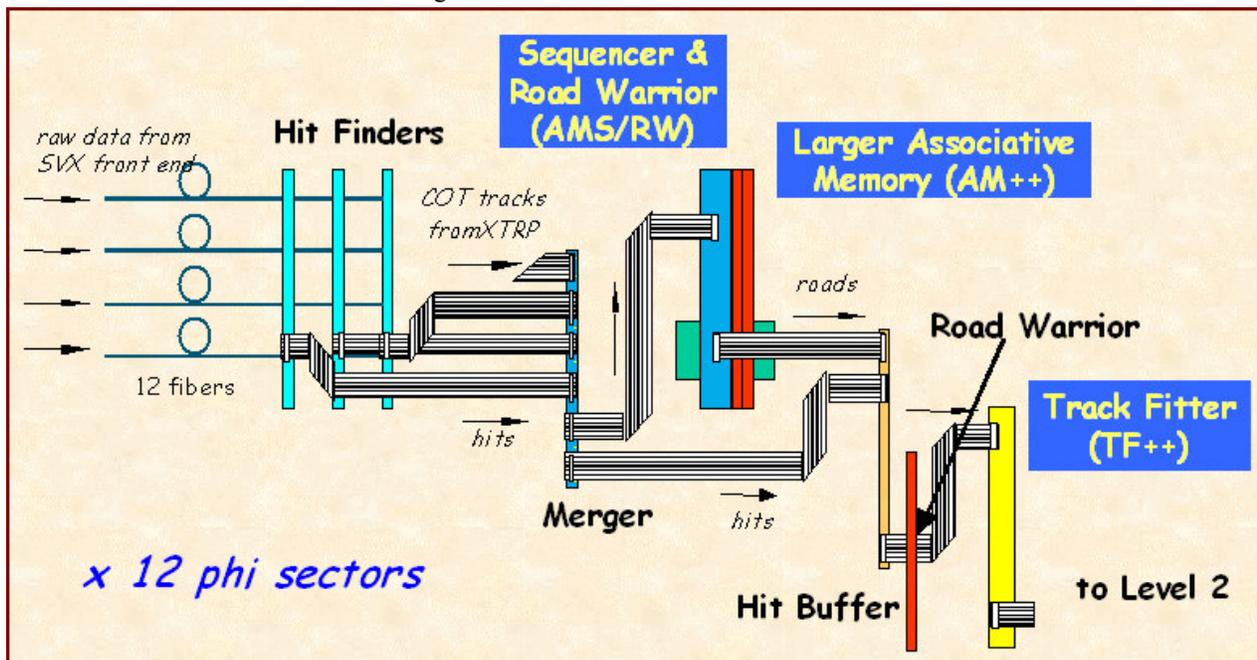


Figure 1: The upgraded SVT processor. The highlighted elements are the one that have been upgraded in this first step. The Road Warrior, which used to be located right before the TF has been removed in the upgrade and its functionality has been implemented in the new Sequencer and Road Warrior (AMS/RW). It is still shown in the picture.

memories can be read via VME so that the pipeline's content, and the offending data, can be analyzed. These memories are called Spy Buffers [12], and have been used extensively during the commissioning of the upgraded SVT.

III. SOFTWARE UPGRADE

The SVT control, downloading and monitoring software needed to configure SVT and communicate with the DAQ system was upgraded to handle the new hardware. These packages were developed and used during commissioning of the upgraded hardware so that they were ready and fully tested before installation. At the same time, we increased the flexibility of the preexisting software so that it is possible to handle a hybrid system in which only a portion of the system contains new hardware, making it possible to check the performance of the upgraded SVT at every step of the installation.

With the increase in the number of patterns, many parameters needed retuning: the size of the roads, the efficiency of the road bank, and the constants used in the linear track fit [13,14]. These parameters had to be clearly defined before the hardware was ready. We chose to emulate the upgraded SVT on real data already acquired with the CDF detector. The SVT bit-level simulation was therefore upgraded as well, implementing the emulation of each new hardware element before the actual installation. This was an essential tool that allowed us to develop and refine the new algorithms as well as test the prototype and production boards.

The parameters used to generate the 128k-pattern bank were chosen to make the new bank at least as efficient as the previous 32k patterns. We chose to keep the efficiency constant and use the increased bank space to reduce the width of the roads as much as possible to minimize the number of fits to be executed and thus the SVT processing time. This is a simple choice that allows for an easy comparison between the new and old system. Optimized tuning will be done when the upgrade is complete.

IV. TEST DESCRIPTION

The commissioning of the upgraded SVT took place in the summer of 2005 while the experiment was taking data. A very careful test procedure was devised to reduce to a negligible level the impact of the commissioning on detector operation and functionality.

The test steps are described below:

1. The first validation was performed on a stand-alone test stand, where random events were sent through the new boards and the output was compared to board level simulation.
2. A second level of validation was performed in the Vertical Slice test (VS). We used a test crate, placed near the real SVT, to reproduce the configuration of a single SVT wedge. Initially, the VS crate included the AMS/RW and AM++ prototypes and the old HB, RW and TF boards. We kept the old RW in the crate so that it was possible to turn on and off the RW functionality of the AMS/RW to test the AMS algorithm by itself. The TF++ was inserted later, as soon as it had passed the first step of validation. We used a MRG to provide the input hits to the VS in two different modes of operation. In the first mode, we

downloaded previous data read out from the Spy Buffers. In the second mode, we split the inputs to one SVT wedge and sent a copy to the VS. In this second configuration the output of the VS could be compared directly with that of the SVT wedge receiving identical input data. We proceeded in the installation only after all discrepancies between the two systems were completely understood. The same inputs were also fed to a board-level simulation. Comparisons between the hardware and the simulation were used to validate both the board and the upgraded simulation itself.

3. Finally, after a successful phase 2 the new board was inserted into a single SVT wedge for a short, low-luminosity test, and, if successful, for 100 hours of data taking at any luminosity. This final test was important because it provided higher statistics than the previous tests, allowing for detection and debugging of lower rate errors. Furthermore, it was an extra check that the control signals used by the DAQ system do not interfere with the board functionality. During these tests the standard SVT monitoring, which runs on crate processors, was used for validation. We monitored the impact parameter, the azimuthally angle and track transverse momentum distributions.
4. At this stage, SVT tracks were compared with simulation and tracks reconstructed in the Level-3 trigger. The efficiency and failure rates were monitored. The efficiency is defined as the fraction of SVT tracks matching Level-3 trigger tracks and is measured to be ~80%. The failure rate is defined as the fraction of reconstructed tracks that do not match to simulated tracks. The failure rate must be as low as possible for two reasons: failures can be a symptom of hardware problems, and we need the simulation to reproduce the hardware in a very detailed way for purposes of data analysis, where the simulation is used to understand various efficiencies. We allow for a failure rate of the order of 10^{-3} in the whole SVT.

Each prototype had to work correctly in data taking both at low and high instantaneous luminosity before it was considered ready to be installed into all 12 wedges.

V. INSTALLATION PROCEDURE

The AMS/RW and AM++ commissioning required the biggest changes to the crate layout. We removed 24 AM boards and 12 AMS boards and replaced them with 12 AMS/RW and 12 AM++ boards. Small boards were installed on the P3 backplanes to provide extra communication lines between the AMS/RW and the AM++ in each wedge.

A phased installation was followed for the commissioning of the first step of the upgrade.

We first installed the new AMS/RW and AM++ with 128k patterns per wedge. However, we did not change the functionality of the system; we started off using only 32k patterns, since a larger bank size was not supported by the old TF. However this step of installation was important because it was possible to change a single crate at a time in short intervals (just 1-2 hours) between two stores. The old hardware was easily mixed with the new hardware since at this stage they were fully equivalent. The RW function was implemented in the AMS/RW and the Pulsar board used for the old RW was removed for the next step, the TF++ installation.

We programmed the former RW Pulsars using the TF++ firmware and used them to replace the old TF boards.

With the TF++ in place, we utilized the full 128k-pattern bank.

We were able to install the new boards and configure them in the few hours available between two stores. This conservative procedure allowed for quick recovery if there were a failure, since each small change was immediately checked.

VI. UPGRADE PERFORMANCE

The main goal of the upgrade was to increase the Level-1 Accept rate (L1A rate) by reducing the tail of the SVT processing time distribution. Figure 2 shows the effect of the SVT upgrade. The fraction of events requiring more than 50 μ s inside SVT is plotted as a function of the instantaneous luminosity. Red points (old SVT) show the tail increasing rapidly with luminosity, while the black points (final 1st step upgrade result) show a much softer dependence. The blue and pink points show the effect of the intermediate steps: the AMS/RW and AM++ installation (blue) and the TF++ insertion (pink).

The best demonstration of the improvement due to the upgrades is provided by the percentage of dead time plotted as a function of the L1A rate (figure 3). The maximum L1A rate possible for a fixed deadtime can be seen in the plot. The blue squares are for the SVT after the AM++ and AMS/RW installation, but using 32k patterns. The Black triangles are from a run in which the new TF++ was installed but we still used only 32k patterns per wedge. The pink circles include the improvements from using 128k patterns per wedge. Squares (and circles) are connected by two adjacent lines: the upper line corresponds to high instantaneous luminosities (over $100 \times 10^{30} \text{ cm}^{-2} \text{ sec}^{-1}$) while the lower line corresponds to low instantaneous luminosities (below $100 \times 10^{30} \text{ cm}^{-2} \text{ sec}^{-1}$). At the lower luminosities, there is a different trigger mix that includes an increasingly greater fraction of lower P_T (soft) events that are characteristic of b-physics. The deadtime is maximum at the start of a store and decreases as the luminosity decreases. Deadtimes below 10% are allowed for high instantaneous

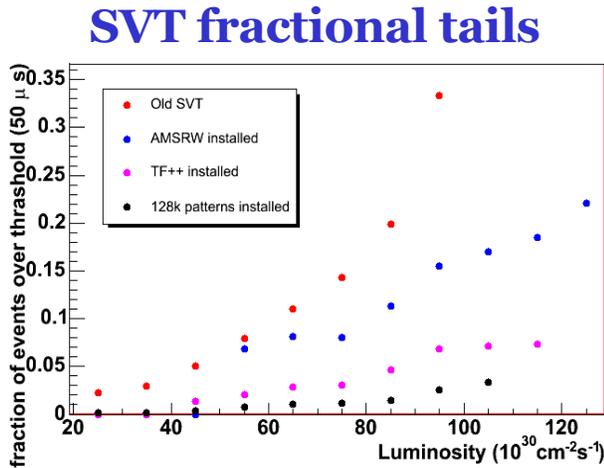


Figure 2: The fraction of events that took larger than 50 μ s to process as a function of instantaneous luminosity.

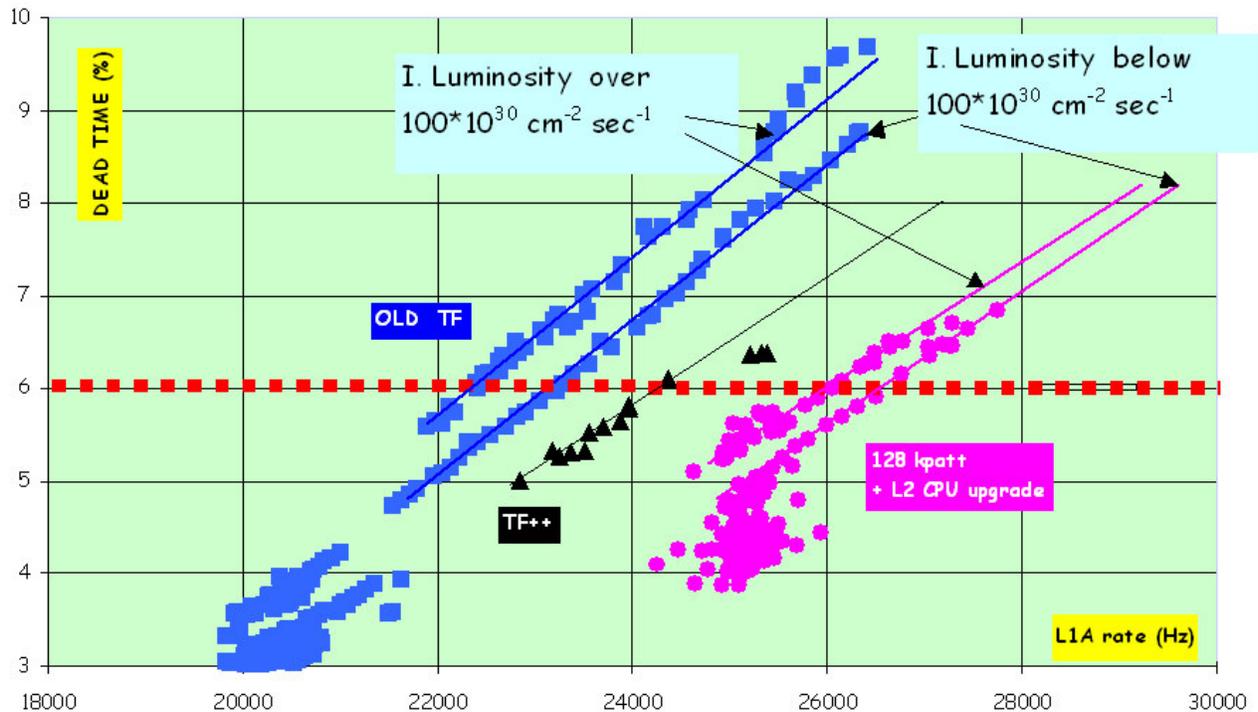


Figure 3: Deadtime (%) versus the L1A (Hz). The square points refer to SVT after the AM++ and AMS/RW installation, and with a small AM bank (32k/wedge). Black triangles refer to a run where the new TF++ was installed but the old bank of 32k patterns was used. Circles show the same system with a 128k pattern bank. A comparison with SVT before AMS/RW installation is not possible because the trigger table changed at that time. Squares (and circles) are connected by two adjacent lines: points on the left line correspond to the high instantaneous luminosity trigger table (over $100 \times 10^{30} \text{ cm}^{-2} \text{ sec}^{-1}$) and points on the right line correspond to low instantaneous luminosities (below $100 \times 10^{30} \text{ cm}^{-2} \text{ sec}^{-1}$)

luminosities, while dead times below 5% are required at low luminosities. The dead time for a softer trigger selection is slightly lower at the same L1A rate because these events have a lower occupancy in the detector and require less processing time. A comparison with SVT before the AMS/RW installation is not possible because of a significant change in the trigger table at that time.

The installation of the 128k pattern bank and the use of a faster L2 CPU allows the L1A rate to move up to 26 kHz with a dead time of 6%, to be compared to a rate of ~22 kHz available before the TF++ installation. Half of this gain is due to the new TF++ (black points) and the increased speed to fit each tracks. These advantages are expected to increase at even higher luminosities, where the events are more complex due to multiple interactions.

When the luminosity and the dead time are low enough, the trigger rate prescales are progressively released, keeping the L1A rate constant. This can be observed in the plot as the crowd of square points at 20-21 kHz and a second crowd of round points at 25-26 kHz. The extra bandwidth is filled in this case with B-triggers: the gain is ~5 kHz, corresponding to a 50% increase in signal acceptance at $1E32 \text{ cm}^{-2}\text{sec}^{-1}$, with dead time kept below 5%.

The improvement due to adding the AM++ and AMS/RW is clearly shown in Figure 2, even if it is not documented in Figure 3. The time gain is mainly due to moving the RW function before the HB board. The old HB is the slowest board in the system (24 MHz), so the removal of duplicate roads before the HB is a large advantage. The increased clock frequency of the AM++ and AMSRW (40 MHz vs. 30 MHz) is a secondary advantage.

One further and important advantage in the new AM++ system will be activated soon: roads with hits in all silicon layers (5/5: 5 matches for 5 layers) will be transmitted from the associative memory chip pipelines [9] as soon as they are available. They do not need to wait until all of the hits have entered the board, as was done in the old AM system. Roads are thus partially processed even before the event readout is complete. Roads characterized by a small inefficiency (4/5: 4 matches for 5 layers) will still be transmitted when the event has been totally read out. This road ordering readout (5/5 first, 4/5 second) allows for better removal of duplicate roads by the RW. The use of this ordering to further improve the SVT speed is part of the system optimization that will happen in the near future.

VII. FUTURE IMPROVEMENTS

A large gain in the future will be provided by a further factor of 4 increase in the size of the associative memory, which will be possible once the new HB (HB++) are installed.

With the final AM++ size of 512k patterns per wedge it will be possible to further reduce the combinatorics inside roads, which will be very narrow. The increased bank size can be used not only to reduce road width, but also for other purposes, for example:

- To increase the track reconstruction efficiency;
- To improve the purity of some triggers, we can expand the use of SVT in L2 track triggers for which cuts on impact parameter are not required. This is particularly important for triggers whose rates increase rapidly as a function of the instantaneous luminosity. An important example would

be the use of silicon-only tracks provided by SVT to improve the muon trigger in the forward region, where the XFT coverage is poor.

VIII. CONCLUSIONS AND PERSPECTIVES

The first part of the SVT upgrade has been successfully built, installed and successfully tested using CDF data. The installation was completed during data taking with minimum impact on detector operations.

As a result of the upgrade, the SVT processing time has been significantly reduced for complex events. As a consequence, the dead time has also been substantially reduced. It is now possible to increase the L1A rate. The speed and flexibility of the new SVT allows for the implementation of better algorithms to further speed up the system and improve tracking performance.

The second and last upgrade step will produce even further gains that will prepare SVT for the highest Tevatron luminosities.

REFERENCES

- [1] The CDF Collaboration, "The CDF II Detector Technical Design Report", FERMILAB-Pub-96/390-E, 1996.
- [2] W.Ashmanskas et al, "The CDF Silicon Vertex Tracker," Nucl. Instr. and Meth., vol. A518, 2004 pp. 532-536.
- [3] E.J. Thomson, et al., "Online track processor for the cdf upgrade", IEEE Trans. on Nucl. Sci., Vol. 49, 2003, pp. 1063-1070.
- [4] The CDF Collaboration, "The CDF IIb detector: Technical design report", FERMILAB-TM-2198.
- [5] R. Amendolia et al., "The AMchip: a Full-custom MOS VLSI Associative memory for Pattern Recognition," IEEE Trans. on Nucl. Sci., Vol. 39, 1992 pp. 795-797.
- [6] S. Belforte et al., "The SVT Hit Buffer," Nucl. Instr. and Meth., vol. A409, 1998 pp. 658-661.
- [7] J.Adelman et al. "The Road Warrior for the CDF Online Silicon Vertex Tracker," Proceedings of 2004 IEEE NSS/MIC, Rome, Italy, 16-22 Oct. 2004.
- [8] K.Anikeev et al, "CDF Level 2 Trigger Upgrade", accepted for publication on Sept. 2005, IEEE TNS No.: TNS-00650-2004.R2. <http://hep.uchicago.edu/~thliu/projects/Pulsar/>
- [9] A.Bardi et al, "The AM++ Board for the Silicon Vertex Tracker Upgrade," presented at this Conference.
- [10] A.Bardi et al, "The AMS/RW Board for the Silicon Vertex Tracker Upgrade," presented at this Conference.
- [11] W.Ashmanskas et al., "Performance of the CDF Online Silicon Vertex Tracker", IEEE Trans. On Nucl. Sci., Vol 49, 2002, pp. 1177-1184
- [12] M.Rescigno, "Online Monitoring for the Silicon Vertex Tracker at CDF II", proceedings of IEEE-RT2003 conference, Montreal 18-23 2003; http://www-cdf.fnal.gov/cdfnotes/cdf7188_spymon.pdf
- [13] R.Carosi and G.Punzi, "An algorithm for a real time track fitter", Conference Records of the 1998 IEEE Nuclear Science Symposium, Toronto, Canada, Nov. 1998.
- [14] H.Wind, "Principal component analysis and its applications to track finding" in *Formulae and methods in experimental data evaluation*, vol. III, R. Bock, K. Bos, S. Brandt, J. Myrheim, M. Regler, Eds. European Physical Society, 1984, pp. k1-k16.