Fermi National Accelerator Laboratory

FERMILAB-Conf-97/244-E
CDF

SVT: An Online Silicon Vertex Tracker for the CDF Upgrade

A. Bardi et al.
For the CDF Collaboration

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

July 1997

Published Proceedings of the *7th Pisa Meeting on Advanced Detectors*,
La Biodola, I.d'Elba, Italy, May 25-31, 1997

Disclaimer

Distribution

CDF/PUB/TRIGGER/PUBLIC/4246

SVT: an online Silicon Vertex Tracker for the CDF upgrade

A.Bardi[c], S.Belforte[c], J.Berryhill[a], A.Cerri[c], A.G.Clark[b], R.Culbertson[a], M.Dell'Orso[c], S.Donati[c], J.Dusatko[a], H.J.Frisch[a], S.Galeotti[c], P.Giannetti[c], A.Leger[b], E.Meschi[c], F.Morsani[c], T.Nakaya[a], G.Punzi[c], L.Ristori[c], H.Sanders[a], M.Shochet[a], T.Speer[b], F.Spinella[c], P.Wilson[a], X.Wu[b], A.M.Zanetti[d].

[a] *University of Chicago - U.S.A.*
[b] *University of Geneve - Switzerland*
[c] *University, Scuola Normale Superiore and INFN Pisa - Italy*
[d] *INFN Trieste - Italy*

The SVT is an online tracker for the CDF upgrade which will reconstruct 2D tracks using information from the Silicon VerteX detector (SVXII) and Central Outer Tracker (COT). The precision measurement of the track impact parameter will then be used to select and record large samples of B hadrons. We discuss the overall architecture, algorithms, and hardware implementation of the system

1 Introduction

The SVT will work in the level 2 of the CDF trigger chain, to refine the level 1 tracking information from the eXtra Fast Tracker (XFT), which uses data from the central drift chamber COT. It will combine XFT tracks with hit coordinates from the SVXII.
The level 2 latency time is about 20 μs, therefore the design of the SVT has concentrated on parallelizing the various tasks, from the reconstruction of the hit coordinates from the single strip pulse heights to the pattern recognition and final precision track fitting (Fig.1). The result is a data driven architecture in which many functions overlap in the internal processor pipeline, and which comprises several different modules, built on 9U Eurocard boards with VMEbus implementation for diagnostic and control. The SVT specifications require 30 MHz operation for each module and an asynchronous data transfer rate of 630 Mbit/s on custom data paths.

2 Tracking Strategy

In a typical two-jet event at CDFII about 60 tracks will traverse the SVXII, and about 6,000 strips will have a significant pulse height (we assume a 5 % occupancy dominated by noise). SVT will receive sparsified and digitized pulse heights from the front end via an optical link. After pedestal subtraction, the Hit Finder board (HF) finds hits and calculates the charge center-of-gravity.
Hits from HF and track parameters from the XFT are sent simultaneously to the Associative Memory bank (AM) and the Hit Buffer (HB). The COT track candidates found by XFT (about 3 per event on the average) are then matched to SVXII hits in the AM, which performs the first pattern recognition with limited resolution (the detector is subdivided into superstrips

[1] Published proceedings 7th Pisa meeting on Advanced Detectors, La Biodola, I. d´Elba (Italy), May 25-31, 1997 - presented by E.Meschi

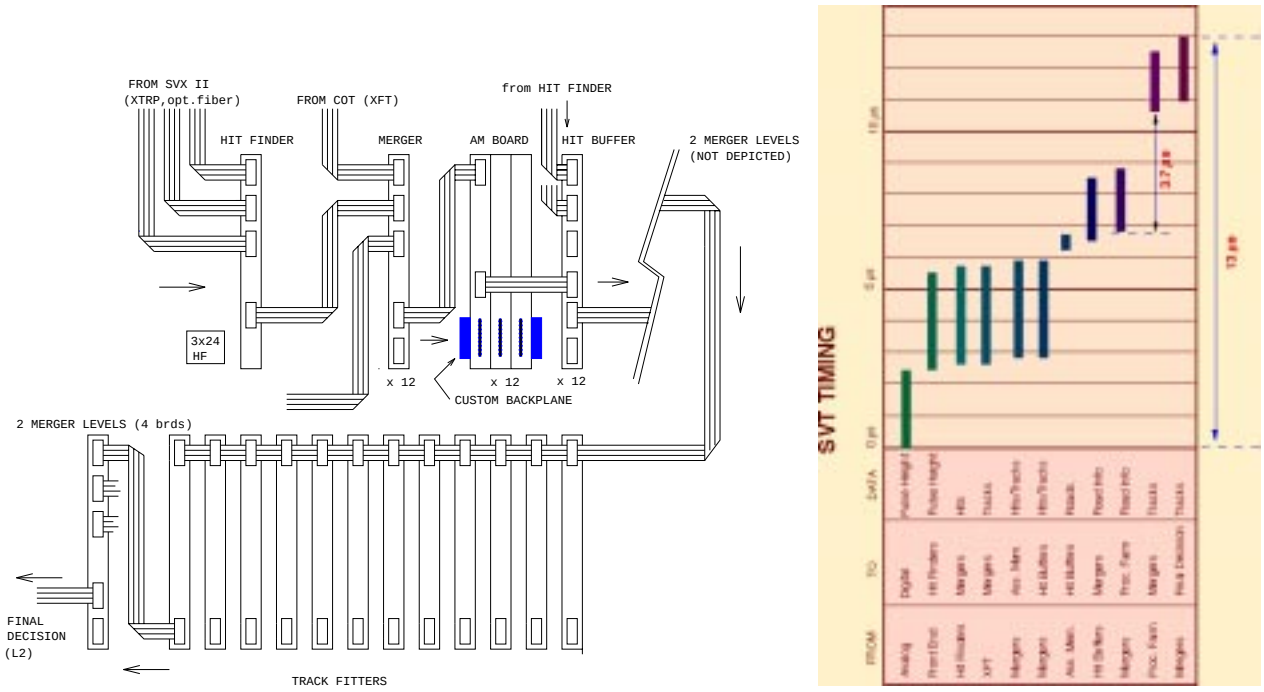

Fig. 1. The SVT architecture (left) and processing time (right)

of 250 μm transverse size). An average of three low-resolution tracks (roads) per event are output from the AM to the HB, which reassociates each road with a list of hits. The road+hits information is subsequently sent to the Track Fitters (TFs) which perform a full resolution ($\sim$ 15 μm) fit using a fast linearized algorithm. Each processor in the TF farm gets one road and reconstructs one or more tracks within that road. The TFs run in parallel on different roads and the number of TF processors is large enough that each TF will only work on one road per event for the majority of the events. The final output of the TFs is a set of track parameters (P_t, ϕ and impact parameter), to be used by the L2 trigger processors for the final decision. Fig.1(right) shows the overall timing of SVT operation on the average event, starting from L1 accept. Readout, hit finding and AM input overlap in time, and when data are sent to the TFs ($\sim$ 7 μs) a new event can be fed into the HF-AM chain, so that the SVT works basically as a 2-stage pipeline. Under the assumption that all the boards work on a 30 MHz clock, simulations show that the average event processing time for SVT is about 13 μs.

2.1 Associative Memory

In the AM, patterns corresponding to real tracks are stored as a combination of superstrip addresses in the various planes of the detector. The AM recognizes roads while receiving the stream of hits coordinates by successively comparing them with all stored patterns [1]. Roads which have fired in the current event are flagged, and output as a stream of road addresses at the end of the input stage (Fig.2). The AM is subdivided into banks, each AM bank working on one of the 12 30^o ϕ sectors of the SVXII (wedge). Since the size of the AM needed to store all possible roads increases more than exponentially with the number of channels, the algorithm is only applied with limited resolution: the size of the superstrips, and thus the size of each AM bank, is dictated by a compromise: reducing the resolution increases the number of fake roads, while reducing the number of patterns to store. It also increases the probability of multiple hits

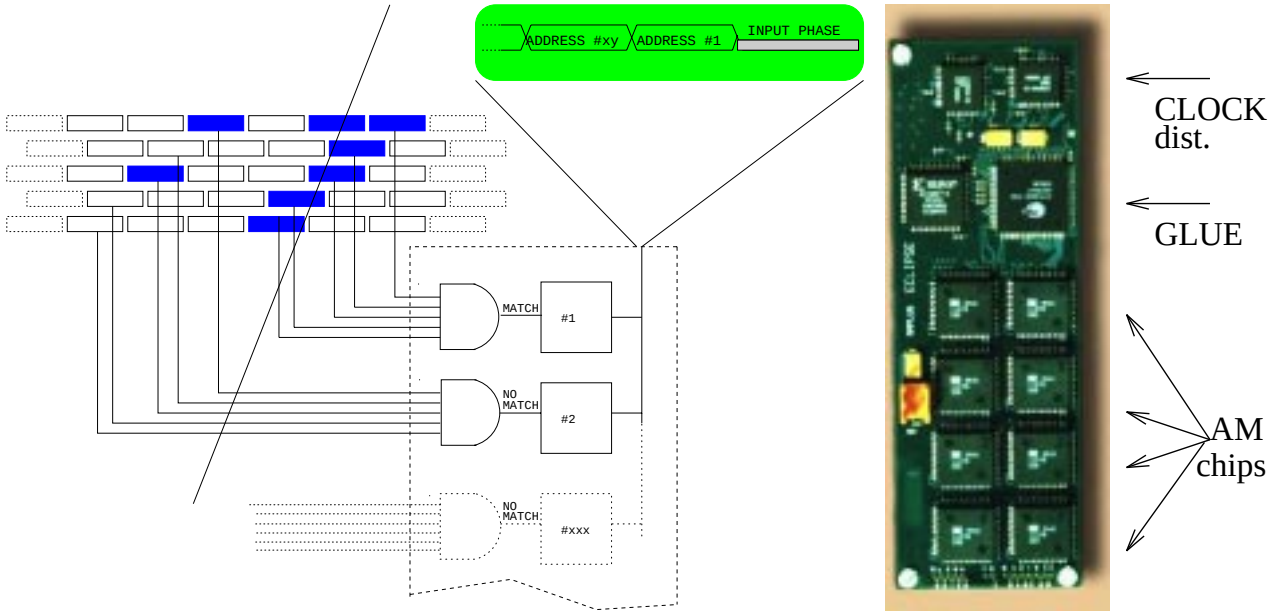

Fig. 2. Associative Memory working principle (left) and a AM8 prototype module (right)

inside a single road, thus increasing the processing time in the TFs. A simulation shows that with a 250 μm superstrip size and assuming an average of 3 real tracks per event, the average number of roads per event is three and a total of about 350K patterns is enough to cover all tracks with $P_t > 2$ GeV/c. To fullfill these requirements, each of the AM banks must have a 32K pattern capacity.

2.2 Linearized Track Fitting

Each TF processor receives a road in the form of a road address and a list of 'hit' coordinates. It must then process all the possible combinations of six coordinates (one for each of the 4 SVXII layers used plus a ϕ and P_t from the XFT track candidate). A combination corresponding to a real track must satisfy three geometrical constraints, which define a 3D hypersurface in the 6D 'hit' coordinate space. It has been shown [2] that it is possible to approximate the constraint hypersurface to an hyperplane, thus reducing the problem of track fitting and track parameters extraction to the performance of a small number of scalar products. The small number of parameters needed to define the constraint plane are calculated at the beginning of a run. Simulations show that a single set of parameters can be used without a significant loss of resolution for all the roads in a single wedge.
To perform the actual calculations two options have been considered. The first uses a commercial DSP to perform the scalar products, and custom logic for road distribution and interconnection with the rest of the system. The second, which is the current default, uses a specialized hardware based on lookup tables which are directly addressed by the hit coordinates [3].

3 Hardware implementation

The architecture of the system is data-driven. Various tasks are performed by different functional blocks, which exchange data via high-speed point-to-point links. Each block starts working upon receiving data and outputs results as they are ready, without any syncronization or handshaking other than the data flow.

A uniform protocol has been designed for data exchange, using unidirectional differential lines on a flat cable, with pipeline transfer driven by an asynchronous clock from the source and a FIFO on the receiving end, which is pushed at every clock cycle. A loose handshake is provided by an "almost full" signal issued by the FIFO and interpreted by the source as an "HOLD". The 25 lines cable has 21 data lines, a clock, HOLD, End Packet (EP) and End Event (EE) signals. An EP bit marks the last word of a packet, while EE marks the end of data for the current event: each module will issue an EE after it has received an EE at all its inputs.
Hits are received via optical fibers terminated by receivers mounted on transition modules. The data is transmitted through the backplane to the HF modules containing the hit clustering logic. Three HF's are required to process one wedge of the SVXII. In a HF, strip address and PH flow through a shift register: as they go, PH in three adjacent strips is used as address for a lookup table implementing the clustering algorithm (e.g. adjacent strips with low-high-low configuration) and containing the cluster center coordinate.
The Merger board can merge four independent input streams into a single output stream, and is standardized for use in all the system where a merging is needed. Merging is performed on a first come, first served basis, which preserves the packet structure and the event structure. EE bit is asserted on output after EE is received on all the input streams.
Each AMbank comprises three boards. The first, called the AM sequencer, receives hit coordinates from the HF and maps them to superstrips, using a lookup table, coordinates various AM operations, and outputs road info to the HB. The AM chips are housed on two twin boards (AM boards) connected to the sequencer via a custom backplane (see Fig.1). The current AM implementation [3] is a .7 μm full custom CMOS chip with 128 patterns (one pattern has 6 layers of 12 bits each), so each AM board houses 128 AMchips in a tree structure connected in groups of 8 by an intermediate logic (GLUE), realized on large FPGAs, to perform address decoding, pipelining and syncronization (Fig. 2 right).
Each HB board works on one wedge, receiving and storing hits and tracks in a Hit Memory, and the current hit count for each superstrip in a Hit Count Memory (Cache-Tag RAM). The superstrips associated with a given road are stored in a lookup table in the HB. When a road is received this information is used to address the Hit Count Memory, and the hits corresponding to that superstrip are output in a road-info packet.
The TF processor farm receives the road-info packets. Although all N processors receive all data, each one will only process $1/\mathrm{N}^{\mathrm{th}}$ of the packets according to a simple algorithm (e.g. processor i works on the i^{th},$(i+N)^{\mathrm{th}}$,$(i+2N)^{\mathrm{th}}$,... roads. All tracks from all the processors are merged into a single stream and delivered to a L2 processor.
The hardware just described requires eight standard CDF crates as used for the CDF DAQ upgrade. Six crates house the HFs, AMs, and HBs for the 12 ϕ sectors of SVXII. A crate houses the Merger modules needed to combine output from the 12 sectors and from the 12 Track Fitters into one stream. The last crate can house up to 20 TF modules.

References

1. M.Dell'Orso and L.Ristori, Nucl. Inst. Meth. **A278** (1989), 436-440;
2. SVT Technical Design Report internal note `CDF/DOC/TRIGGER/PUBLIC/3108`;
3. S.R.Amendolia et al. IEEE Trans. Nucl. Sci. **39** n.4 (1992);