The MERGER Board of the CDF Silicon Vertex Tracker

M. Bari, S. Belforte, A. Cerri, M. Dell'Orso, S. Donati, S. Galeotti, P. Giannetti, F. Morsani, G. Punzi, L. Ristori, F. Spinella, and A. M. Zanetti

***Abstract*—The Merger board is part of the Silicon Vertex Tracker (SVT), a device dedicated to perform real-time track reconstruction with offline-like resolution and high efficiency at the Level 2 trigger of the CDF experiment. The Merger is a custom 9U $\times$ 400 mm VME board, running at an internal clock frequency of 33 MHz. Its main functional task in SVT is to merge up to four independent data streams into a single one. The merging operation can be performed on a first come, first served basis or according to an ordered sequence. There are four input streams and two identical output streams, so that the Merger also serves as a data fanout function. The board implements detailed error handling and sophisticated data monitoring that make it possible to trace back misfunctioning both in the Merger and in other parts of SVT. Furthermore, the Merger has special modes of operation that can be selected for test and diagnostic purposes. In these working modes, the Merger is a powerful tool that allows one to test other SVT boards at their maximum operating frequency.**

***Index Terms*—Online track reconstruction, trigger, VME.**

I. Introduction: Silicon Vertex Tracker

THE Silicon Vertex Tracker (SVT) [1] is a complex electronic device that reconstructs online charged particle trajectories with offline-like resolution and high efficiency at the Level 2 trigger of the Collider Detector at Fermilab (CDF) experiment. SVT finds and fits tracks combining the digitized hits from the Silicon VerteX detector (SVXII) [2] with the tracks reconstructed in the central outer tracker drift chamber (COT) by the Level 1 track trigger processor (the extremely fast tracker—XFT [3]).

SVT is a parallel pipelined data-driven device. It is made of more than one hundred 9U $\times$ 400 mm VME boards housed in eight crates. SVT is segmented into 12 identical systems running in parallel, each handling one SVXII sector (wedge) covering a 30° azimuthal angle. In each SVT sector, track reconstruction is accomplished by six different VME boards (Fig. 1): Hit Finder, Merger, Sequencer, and Associative Memory boards (the Associative Memory system), Hit Buffer, and Track Fitter.

Raw SVXII data flow from the front-end to Hit Finder boards that find clusters of silicon strips with a significant energy deposit and compute the coordinate of the centroid (hits). The hits found by the Hit Finders and the P_t and ϕ of COT tracks as computed by XFT are then merged in the Merger board and sent to both the Associative Memory system and to the Hit Buffer [4]. The Associative Memory system performs pattern recognition: input data are compared with a stored set of patterns in a parallel way using a dedicated custom VLSI chip [5]. Track candidates (roads) are found using a coarse spatial resolution in the SVXII (250 μm) and sent to the Hit Buffer. The Hit Buffer stores all hits and tracks in a wedge for each event in an internal memory, then for each road received from the Associative Memory retrieves the silicon hits and the XFT tracks belonging to that road and sends them to Track Fitter board. The Track Fitter performs quality cuts on tracks and estimates track parameters using the full available spatial resolution in a linearized fit. Finally, a set of four Merger boards collects all the high-precision SVT tracks of one event from the 12 SVT sectors and sends them to the CDF Level 2 trigger processors for the final decision.

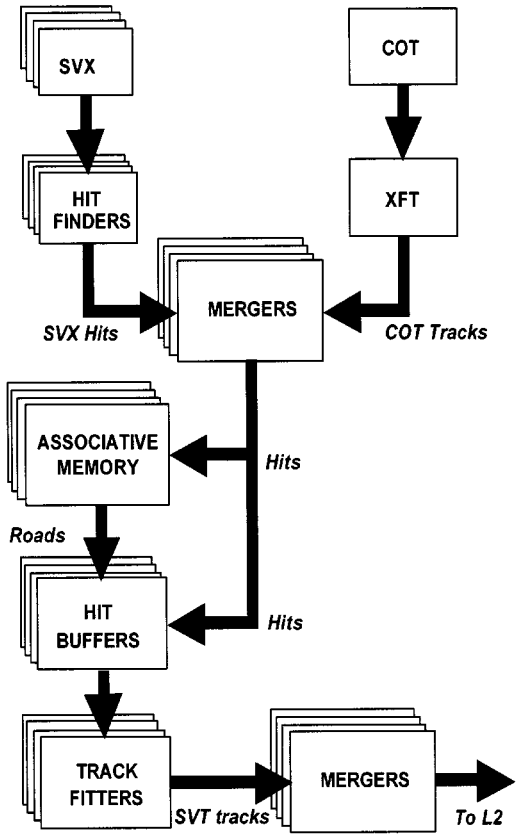

Fig. 1. SVT architecture. The main SVT functional blocks and data paths are shown.

Manuscript received December 10, 2001; revised April 11, 2002.

M. Bari and A. M. Zanetti are with INFN Trieste, 34012 Trieste, Italy.

S. Belforte, M. Dell'Orso, S. Donati, S. Galeotti, P. Giannetti, F. Morsani, G. Punzi, L. Ristori, and F. Spinella are with INFN Pisa, 56010 S.Piero A Grado (PI), Italy, Dipartimento di Fisica, Università di Pisa, 56100 Pisa, Italy, and Scuola Normale Superiore, 56100 Pisa, Italy.

A. Cerri was with Scuola Normale Superiore, 56100 Pisa, Italy. He is now with Lawrence Berkeley Laboratory, Berkeley, CA 94720 USA.

Digital Object Identifier 10.1109/TNS.2002.801663

Data flow through the SVT pipeline on uniform point-to-point connections running at about 700 Mbit/s each. The SVT latency on average is 15 μs, depending on event size.

SVT construction has been completed, and all hardware has been installed and fully operational since the beginning of 2001 [7].

II. Communication Protocol and Data Format

A common data communication protocol is used among all SVT boards, allowing one to combine them in several different ways to accommodate various system configurations and test needs.

Each data connection is implemented on a flat cable carrying 25 low-voltage differential signals that connects the front panel of two boards. On each cable, there are (Table I) 21 data bits, one end packet (EP) bit, one end event (EE) bit, a data strobe (DS), and one hold (HD) signal.

The hold and DS lines control the communication protocol among boards. A word is transferred on each positive-going DS edge in a simple asynchronous pipeline mode. The receiving board drives the hold line to signal a busy state while it still has some input buffer space, so that no data are lost even if the source takes a few clock cycles to suspend data flow in response. Data sending resumes as soon as hold is released. Since there is no word-by-word acknowledgement, the transfer can proceed at full speed even when cable transit times are long.

Data are organized as a sequence of variable-length packets of words. The EP bit marks the last word of each packet. The EE bit is used to mark the end of the data stream for the current event (EndEvent word).

The EndEvent word (a single word packet) contains no physics data and is used to keep track of possible errors occurred during data transmission or processing and to record global information on the event. It includes (Table II) an event identification 8-bit field (event tag), the event parity bit (PA), a summary of error flags for the event, a summary of the Level 1 trigger decision (L1T), and the number of Level 2 buffer (L2B) used to store the event data.

III. Merger Architecture

A diagram of the Merger logic blocks and of the main data and control paths is shown in Fig. 2.

The Merger board has four input data streams and two identical output streams. Each input can be enabled or disabled via the VME interface so that any combination of inputs can be selected. The board can be set to ignore the hold signal from any of the two outputs. The Merger clock frequency is 33 MHz, and one data word is processed each clock cycle.

Input data are received asynchronously: on each of the four board inputs, the incoming DS signal acts as a write clock to store data into a 4K words deep first-in, first-out (FIFO) buffer. The FIFO provides an Almost Full signal that drives the Hold line of the input cable. If the Merger does not keep up with the incoming data rate, this signal is asserted and the source responds by temporarily suspending the data flow.

The Merger combines the packets received from all active inputs and belonging to the same event. This operation preserves the integrity of each packet. The merging of the streams can be performed in two ways, selectable by VME:

1) on a first come, first serve basis: packets in the output stream will be a random mix of packets from the different input streams;
2) in ordered sequence: all the data from the first active input, then from the second active input, and so on.

The EndEvent word is asserted on the output stream after EndEvent is received on all the enabled inputs and all the data packets have been sent to output. The Event Tag number in the EndEvent words of all the inputs is checked for consistency. Thus the Merger acts as checkpoint for data synchronization in the SVT system.

The Merger has a VME interface that enables full access to internal registers, FIFOs, and RAMs and allows one to select different board operation modes. All VME data transfers are 32 bits, and both single word and block transfer are implemented for VME read or write operation. The board recognizes the VME geographical addressing as specified by the VME 64x standard.

A dedicated line on the P2 backplane, named INIT, is used to signal full initialization of the whole SVT system. When INIT is asserted, the Merger responds with a full reset operation of the board and is ready for data processing in a few clock cycles.

The Merger has additional modes of operations that can be selected for test and diagnostic purposes.

TABLE I
SVT Cable Signal Assignment

24	23	22	21	20...0
HD	DS	EE	EP	Data field

TABLE II
EndEvent Word Data Field Format

20...19	18...17	16...9	8	7...0
L2B	L1T	Error Flags	PA	Event Tag

IV. Data Monitoring

Debugging and monitoring the SVT data-driven architecture is a challenging task. SVT has many input streams and only one output stream, performs a very large data reduction (several tens of internal data paths are eventually merged to one single data stream), handles a large volume of data at high speed, and has background rejection power such that only 1/1000 of the input events to SVT is stored on tape. Therefore, it is very important to check the proper working of SVT in real time, especially on events that do not pass through data acquisition (DAQ). The Merger implements careful detection of error conditions and full data flow monitoring via the Spy Buffer system.

A. Error Handling

The Merger can detect the following error conditions.

1) *Parity Error:* The parity of input data does not match the Parity bit of the corresponding input EndEvent word. Data in input are probably corrupted. This function helps to isolate single bit problems in data transmission.

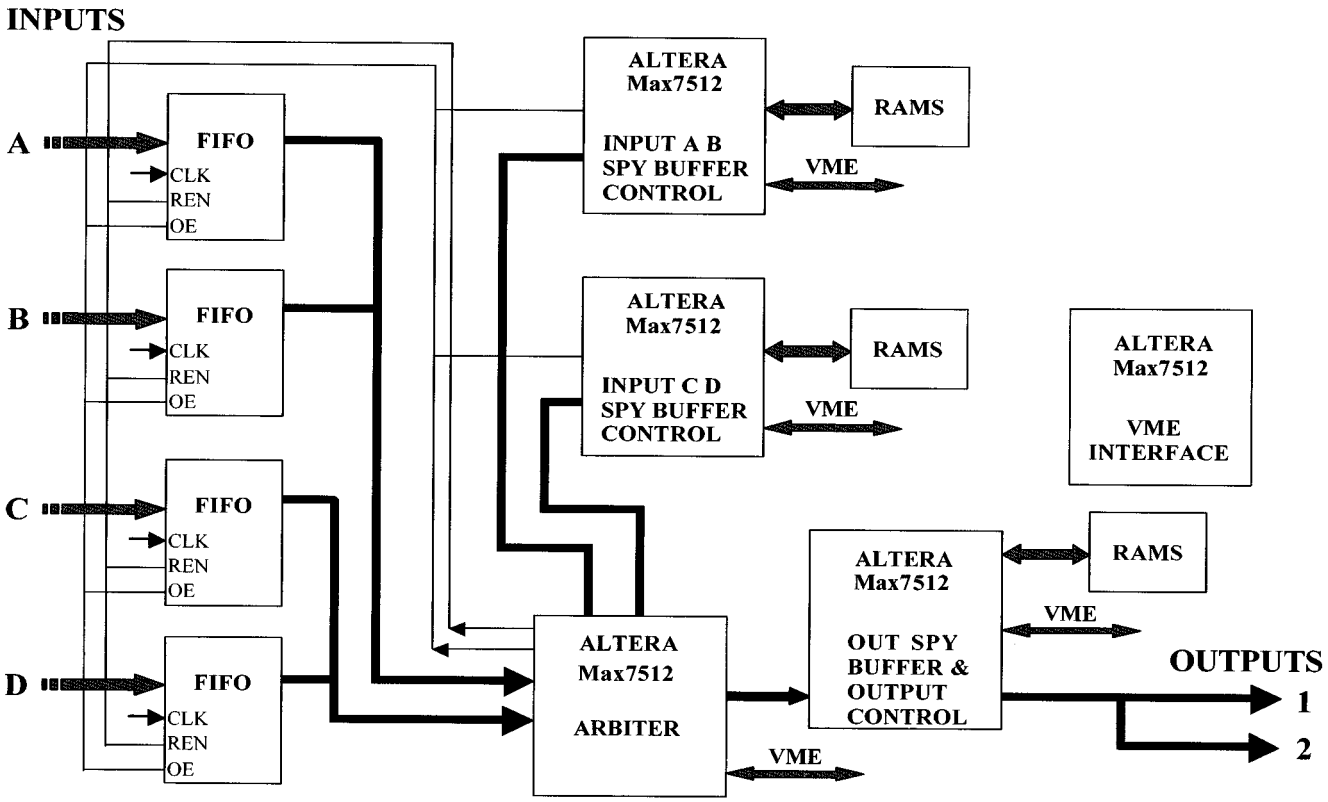

Fig. 2. Merger architecture. The main logic blocks, data, and control paths are shown. Read Enable (REN) and Output Enable (OE) are FIFO control signals.

2) *Lost Sync:* The event tags of the End Event words received in input do not match. The synchronization of the multiple-event data streams has been lost.
3) *FIFO Overflow:* One input FIFO became full. It detects failures of the Hold mechanism that would cause data loss.

Upon error occurrence, an error bit is set in one onboard register and also in the EndEvent word, so that the error condition is propagated in the SVT data stream. Furthermore, the Merger has the ability to drive two lines on the VME backplane: SVT_ERROR and CDF_ERROR. SVT_ERROR is an SVT reserved line and is used for internal SVT diagnostic (see Section IV-B) . CDF_ERROR line is a general CDF reserved line and signals a severe error condition, which can trigger a global DAQ reset. The Merger internal error bits are ANDed with two VME-programmable masks to define which of them can cause SVT_ERROR or CDF_ERROR to be asserted.

In addition, the Merger can drive the CDF_ERROR line in response to specific error bits being set in the incoming EndEvent word(s).

B. Spy Buffers

The input and the output data streams are monitored by continuously copying the data to circular memories, called Spy Buffers, which act as built-in logic state analyzers. These memories are 128K-word static RAMs and can contain several hundred events.

Spy Buffer mode of operation is controlled by an SVT dedicated line (FREEZE) on the VME backplane. When this line is asserted, writing into the Spy Buffers is suspended and their content can be read through the VME interface without any interference with the data flow. This mechanism is used to take a snapshot of all data flowing through each SVT board, for instance upon detection of an error condition or for data monitoring. The SVT_ERROR line can be used to assert the FREEZE in order to record the data that caused the error.

V. Test and Diagnostic Operation Modes

The Merger can be set into different operation modes besides the Run Mode in which it processes trigger data, as described in the previous sections. These additional modes include data sink, data source, and data checker functionalities that allow one to test other SVT boards at their maximum operating frequency over long periods of time, without being limited by the speed at which an external host computer can feed or read data.

The different Merger operation modes can be selected through the VME interface and are the following.

1) *Test Mode:* Data processing is halted; incoming data, if any, cumulate in the input FIFOs until full. Access is enabled from the VME interface to internal components and data paths: FIFOs can be read from VME, Spy Buffer RAMs can be written and read back, and control registers can be written and read.
2) *Data Sink:* Data from input channels are processed as in Run Mode and stored in input Spy Buffers for later readout and check via VME. No data are sent to output, and input data processing never stops.
3) *Data Source:* A set of test words can be stored in the Output Spy Buffer and then sent to the board output at

full speed, obeying the Hold protocol. This allows one to send test data to other boards in exactly the same way as normal SVT data processing. Test data can be sent to output once or in an infinite loop, under VME control.

4) *Data Checker:* The Merger can compare incoming data in real time against expectation and flag and record possible errors. Data arriving on input A can be tested against a set of words preloaded in Spy Buffer B (the same applies for channels C and D). Data on input channel A are processed normally and sent to output as well as copied to the A input Spy Buffer. The list of expected data (e.g., from a simulation of the upstream board) has to be prepared in advance and stored in the B Spy Buffer. Every time a word is written into Spy Buffer A, one is read from Spy Buffer B at the same address. When the list of words preloaded in B is exhausted, the address for A and B Spy Buffer is reset and the process keeps going. Words from A and B are compared, and the Merger can be requested to perform various actions if they are found different: setting an internal error bit, activating a dedicated signal on the Merger front panel to trigger a scope or a logic state analyzer, and halt data processing. Each of these actions can be independently enabled or disabled via VME. In this way, up to 128K words of data output from any SVT board can be tested in real time. This procedure is particularly powerful when the board under test is driven by a Merger, which sends test data in infinite loop; in this case, that board can be exercised and tested at full speed for an extended time with very complex data patterns.

VI. Merger Implementation and Performance

The Merger is implemented as a nine-unit Eurocard board 400 mm deep (9U $\times$ 400 mm), using an eight-layer printed circuit board with two power-supply planes (5 and 3.3 V). Almost all board logic resides in five fast in-system programmable chips (Altera MAX7512AE) with a transit time of 7 ns in a 208-pin package. The five Merger Spy Buffers are implemented on 15 asynchronous RAMs, 128K $\times$ 8 bit each, with 10 ns access time, while eight asynchronous FIFOs 4K $\times$ 8 bit (10 ns access time) handle the four-input streams.

Clock lines have been routed with controlled impedances and series terminations and have the same lengths; all components of the Merger internal pipeline receive the clock at the same time. Clock signal distribution is handled by two Roboclock chips that implement a zero delay fanout. Roboclock's adjustable delays are used to control placement of critical data strobe signals within the clock cycle.

Since Altera chips run at 3.3 V, while other components run at at 5 V, dedicated circuits provide failsafe operation of the Merger board in case of failure of one single voltage, or simply if one power supply turns on (off) before the other, providing protection against driving inputs of unpowered chips beyond tolerances or asserting VME control or data lines or output Data Strobe.

In total, 30 Merger boards have been produced and used in test stands and in the actual operation of the SVT trigger. SVT runs using 16 Merger boards, but a few more are usually seated in SVT crates and used for test and diagnostic functions.

All boards have been in operation for about one-and-a-half years, and no device failure has been observed so far.

References

[1] S. Belforte, M. Dell'Orso, S. Donati, G. Gagliardi, S. Galeotti, and P. Giannetti *et al.*, "The CDF trigger SVT," *IEEE Trans. Nucl. Sci.*, vol. 42, pp. 860–864, Aug. 1995.

[2] J. A. Valls, "The SVX II silicon vertex detector at CDF," *Nucl. Phys. B*, vol. 78, pp. 311–314, 1999.

[3] C. Ciobanu, "Online track processor for the CDF upgrade," *IEEE Trans. Nucl. Sci.*, vol. 46, pp. 933–937, Aug. 1999.

[4] S. Belforte, M. Dell'Orso, S. Donati, G. Gagliardi, S. Galeotti, and P. Giannetti *et al.*, "The SVT hit buffer," *IEEE Trans. Nucl. Sci.*, vol. 43, pp. 1810–1813, June 1996.

[5] S. R. Amendolia, S. Galeotti, F. Morsani, D. Passuello, L. Ristori, and N. Turini, "The AMchip: A full-custom CMOS VLSI Associative Memory for pattern recognition," *IEEE Trans. Nucl. Sci.*, vol. 39, pp. 795–797, Aug. 1992.

[6] M. Bari, S. Belforte, A. Cerri, M. Dell'Orso, S. Donati, and S. Galeotti *et al.*, "Error handling for the CDF Silicon Vertex Tracker," *IEEE Trans. Nucl. Sci.*, vol. 48, pp. 1286–1289, Aug. 2001.

[7] W. Ashmanskas, A. Barchiesi, A. Bardi, M. Baumgart, M. Bari, and S. Belforte *et al.*, "Performance of the CDF online Silicon Vertex Tracker," *IEEE Trans. Nucl. Sci.*, pp. 1177–1184, June 2002.