The Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System

D. Acosta, A. Madorsky (Madorsky@phys.ufl.edu), B. Scurlock, S.M. Wang

University of Florida

A. Atamanchuk, V. Golovtsov, B. Razmyslovich, L. Uvarov

St. Petersburg Nuclear Physics Institute

Abstract

We report on the development and test of a prototype track-finding processor for the Level-1 trigger of the CMS endcap muon system. The processor links track segments identified in the cathode strip chambers of the endcap muon system into complete three-dimensional tracks, and measures the transverse momentum of the best track candidates from the sagitta induced by the magnetic bending. The algorithms are implemented using SRAM and Xilinx Virtex FPGAs, and the measured latency is 15 clocks. We also report on the design of the pre-production prototype, which achieves further latency and size reduction using state-of-the-art technology.

I. INTRODUCTION

The endcap muon system of CMS consists of four stations of cathode strip chambers (CSCs) on each end of the experiment. The coverage in pseudo-rapidity (η) is from 0.9 to 2.4. A single station of the muon system is composed of six layers of CSC chambers, where a single layer has cathode strips aligned radially (from the beam axis) and anode wires aligned in the orthogonal direction. The CSC chambers are trapezoidal in shape with a 10° or 20° angular extent in azimuth (φ). The CSC chambers are fast (60 ns drift-time) and participate in the Level-1 trigger of CMS.

A "Local Charged Track" (LCT) forms the most primitive trigger object of the endcap muon system. Both cathode and anode front-end LCT trigger cards search for valid patterns from the six wire and strip planes of the CSC chamber. The anode data provide precise timing information as well as η information, and the cathode data provide precise φ information. A motherboard on the chamber collects the LCT information, associates the wire data to the cathode data, tags the bunch crossing time, and selects the two best candidates from each chamber. The end result is a three-dimensional vector, encoded as a bit pattern, which corresponds to a track segment in that muon station. It is transmitted via optical links to the counting house of CMS. To reduce the number of optical connections, only the three best track segments are sent from nine chambers (18 track segments).

The Track-Finder must reconstruct muons from track segments received from the endcap muon system, measure their momenta using the fringe field of the central 4 T solenoid, and report the results to the first level of the trigger system (Level-1). This objective is complicated by the non-uniform magnetic field in the CMS endcap and by the high background rates; consequently, the design must incorporate full 3-dimensional information into the track-finding and measurement procedures.

The experimental goal of the Track-Finder is to efficiently identify muons with as low a threshold in transverse momentum (P_T) as possible in order to meet the rate requirement of the Level-1 Trigger of CMS. This translates into a single muon trigger rate which does not exceed about 1 kHz per unit rapidity at the full luminosity of the LHC. The resolution on P_T, therefore, should be less than about 30% at least, which requires measurements of the φ and η coordinates of the track from at least three stations.

II. TRACK-FINDER LOGIC

The reconstruction of complete tracks from individual track segments is partitioned into several steps to minimize the logic and memory size of the Track-Finder [1]. The steps are pipelined and the trigger logic is deadtime-less.

First, nearly all possible pairwise combinations of track segments are tested for consistency with a single track. That is, each track segment is extrapolated to another station and compared to other track segments in that station. Successful extrapolations yield tracks composed of two segments, which is the minimum necessary to form a trigger. The process is not complete, however, since the Track-Finder must report the number of distinct muons to the Level-1 trigger. A muon that traverses all four muon stations and registers four track segments would yield six track "doublets." Thus, the next step is to assemble complete tracks from the extrapolation results and cancel redundant shorter tracks. Finally, the best three muons are selected, and the track parameters are measured.

A. Extrapolation

A single Extrapolation Unit forms the core of the Track-Finder trigger logic. It takes the three-dimensional spatial information from two track segments in different stations, and tests if those two segments are compatible

with a muon originating from the nominal collision vertex with a curvature consistent with the magnetic bending in that region. The test involves the following:

- Determine if each track segment is in the allowed trigger region in η

- Compare the η values of the two segments to determine if both lie along a straight line projection to the collision vertex

- Compute the difference in φ between the two track segments

- Check if that difference is consistent with the bending direction in φ measured at each station

- Compare the difference in φ to the maximum allowed at that η for several P_T thresholds

- Compare the quality of the two track segments

- Check that at least one of the track segments is not parallel to the beam axis

- Assign an overall quality to the extrapolation

All possible extrapolation pairs should be tested in parallel to minimize the trigger latency. This corresponds to 81 combinations for the 15 track segments of the endcap region. However, we have excluded direct extrapolations from the first to fourth muon station in order to reduce the number of combinations to 63. This prohibits triggers involving hits in only those stations, but saves logic and reduces some random coincidences (since those chambers are expected to have the highest rates). It also facilitates track assembly based on "key stations," which is explained in the next section.

B. Track Assembly

The track assembly stage of the Track-Finder logic examines the results of the extrapolations and determines if any track segment pairs belong to the same muon. If so, those segments are combined and a code is assigned to denote which muon stations are involved. The underlying feature of the track-assembly is the concept of a "key station." For this design, the second and third muon stations are key stations. A valid trigger in the endcap region must have a hit in one of those two stations. The second station is actually used twice: once for the endcap region and once for the region of overlap with the barrel muon system, so there are a total of three data streams. The track assembler units output a quality words for the best track for each hit in the key stations.

C. Final Selection

The final selection logic combines the nine best assembled tracks, cancels redundant tracks, and selects the three best distinct tracks. For example, a muon which leaves track segments in all four endcap stations will be identified in both track assembler streams of the endcap since it has a track segment in each key station. The Final Selection Unit must interrogate the track segment labels

from each combination of tracks from the two streams to determine whether one or more track segments are in common. If the number of common segments exceeds a preset threshold, the two tracks are considered identical and one is cancelled. Thus, the Final Section Unit is a sorter with cancellation logic.

D. Measurement

The final stage of processing in the Track-Finder is the measurement of the track parameters, which includes the φ and η coordinates of the muon, the magnitude of the transverse momentum P_T, the sign of the muon, and an overall quality which we interpret as the uncertainty of the momentum measurement. The most important quantity to calculate accurately is the muon P_T, as this quantity has a direct impact on the trigger rate and on the efficiency. Simulations have shown that the accuracy of the momentum measurement in the endcap using the displacement in φ measured between two stations is about 30% at low momenta, when the first station is included. (It is worse than 70% without the first station.) We would like to improve this so as to have better control on the overall muon trigger rate, and the most promising technique is to use the φ information from three stations when it is available. This should improve the resolution to at least 20% at low momenta, which is sufficient.

In order to achieve a 3-station P_T measurement, we have developed a scheme that uses the minimum number of bits necessary in the calculation. The first step is to do some pre-processing in FPGA logic: the difference in φ is calculated between the first two track segments of the muon, and between the second and third track segments when they exist. Only the essential bits are kept from the subtraction. For example, we do not need the same accuracy on the second subtraction because we are only trying to untangle the multiple scattering effect at low momenta. The subtraction results are combined with the η coordinate of the track and the track type, and then sent into a 2 MB memory for assignment of the signed P_T. Tracks composed of only two track segments are allowed also in certain cases.

III. FIRST PROTOTYPE SYSTEM ARCHITECTURE

The Track-Finder is implemented as 12 "Sector Processors" that identify up to the three best muons in $60°$ azimuthal sectors. Each Processor is a 9U VME card housed in a crate in the counting house of CMS. Three receiver cards collect the optical signals from the CSC chambers of that sector and transmit data to the Sector Processor via a custom point-to-point backplane. A maximum of six track segments are sent from the first muon station in that sector, and three each from the remaining three stations. In addition, up to eight track segments from chambers at the ends of the barrel muon

system are propagated to a transition board in the back of the crate and delivered to each Sector Processor as well.

A total of nearly 600 bits of information are delivered to each Sector Processor at the beam crossing frequency of 40 MHz (3 GB/s). To reduce the number of connections, LVDS Channel Link transmitters/receivers from National Semiconductor [2] were used to compress the data by about a factor of three through serialization/de-serialization. A custom point-to-point backplane operating at 280 MHz is used for passing data to Sector Processor.

Each Sector Processor measures the track parameters (P_T, φ, η, sign, and quality) of up to the three best muons and transmits 60 bits through a connector on the front panel. A sorting processor accepts the 36 muon candidates from the 12 Sector Processors and selects the best 4 for transmission to the Global Level-1 Trigger.

A prototype Sector Processor was built using 15 large Xilinx Virtex FPGAs, ranging from XCV50 to XCV400, to implement the track-finding algorithm, and one XCV50 as VME interface (Fig. [1]).

The configuration of the FPGAs, including the VME interface, was done via fast VME-to-JTAG module, implemented on the same board. This module takes advantage of the VME parallel data transmission, and reduces the configuration time down to 6 seconds, instead of ~6 minutes if we use a standard Xilinx Parallel III cable.

The following software modules were written to support testing and debugging:

- Standalone version of the C++ model for Windows

- Module for the comparison of the C++ model with the board's output

- JTAG configuration routine, controlling the fast VME-to-JTAG module of the board

- Lookup configuration routine, used to write and check the on-board lookup memory

- Board configuration database with Graphic User Interface (GUI), which keeps track of many configuration variants and provides a one-click selection of one of them. Each variant contains the complete information for FPGA and lookup memory configuration.

All software was written in portable C++ or C, to simplify porting into another operating systems. The Board configuration database is written in JAVA, since this is the simplest way to write a portable GUI. All software can and will be used for the second (pre-production) prototype debugging and testing.

The first prototype was completely debugged and tested. Simulated input data or random numbers were transmitted over the custom backplane to this prototype, and the results were read from the output FIFO. These results were compared with a C++ model, and 100% matching was demonstrated. The latency from the input of the Sector Receivers [3] (not including the optical link latency) to the output of the Sector Processor is 21 clock, 15 of which are used by Sector Processor logic.

Fig. [2] shows the stand, used for testing and debugging of the first prototype.

IV. SECOND (PRE-PRODUCTION) PROTOTYPE SYSTEM ARCHITECTURE

Recent dramatic improvements in the programmable logic density [4] allow implementing all Sector Processor logic onto one FPGA. Additionally, the optical link components have become smaller and faster. All this allows combining three Sector Receivers and one Sector Processor of the first prototype onto one board. This board will accept 15 optical links from the Muon Port Cards, where each link carries the information about one muon track segment. Additionally, the board receives up to 8 muon track segments from the Barrel system via a custom backplane.

Since the track segment information arrives from 15 different optical links, it has to be synchronized to the common clock phase. Also, because the optical link's deserialization time can vary from link to link, the input data must be aligned to the proper bunch crossing number.

Next, the track segment information received from the optical links is processed using the lookup tables to convert the CLCT pattern number, sign, quality and wire-group number into the angular values describing this track segment. The angular information about all track segments is fed to the FPGA, which contains the entire 3-dimentional Sector Processor algorithm. On the first prototype this algorithm occupied 15 FPGAs.

The output of the Sector Processor FPGA is sent to the P_T assignment lookup tables, and the results of the P_T assignment for the three best muons are sent via the custom backplane to the Muon Sorter.

In the second (pre-production) prototype Track-Finder system we stopped using Channel Links for the backplane transmission because of their long latency, and moved to the GTLP backplane technology. This allows transmitting the data point-to-point (from Sector Processor to Muon Sorter) at 80 MHz, with no time penalty for serialization since the most relevant portions of data are sent in the first frame. The data in the second frame are not needed for immediate calculation, so they do not delay the Muon Sorter processing.

The entire second (pre-production) prototype Track-Finder system will fit into one 9U VME crate (Fig. [3]).

V. SECTOR PROCESSOR ALGORITHM AND C++ MODEL MODIFICATIONS

The Sector Processor algorithm was significantly modified to fit into one chip and reduce latency. The comparison of the old and new algorithms is shown on Fig. [4]. In particular, the following modifications were made:

- The algorithms of the extrapolation and final selection units are reworked, and now each of them is completed in only one clock.

- The Track Assembler Units in the first prototype were implemented as external lookup tables (static memory). For the second prototype, they are implemented as FPGA logic. This saved I/O pins on the FPGA and one clock of the latency.

- The preliminary calculations for the P_T assignment are done in parallel with final selection for all 9 muons, so when three best out of nine muons are selected, the pre-calculated values are immediately sent to the external P_T assignment lookup tables.

All this allowed reducing the latency of the Sector Processor algorithm (FPGA plus P_T assignment memory) down to 5 clocks (125 ns) from 15 clocks in the first prototype.

The current version of the Sector Processor FPGA is written entirely in Verilog HDL. The core code is portable; it does not contain any architecture-specific library elements. It is completely debugged with Xilinx simulator in timing mode, and its functionality exactly matches the C++ model.

During the construction and debugging of the first prototype, we have encountered many problems related to the correspondence between hardware and C++ model. In particular, sometimes it is very problematic to provide the exact matching, especially if the model uses the C++ built-in library modules, such as lists and list management routines, etc.

To eliminate these problems in the future, the C++ model was completely rewritten in strict line-by-line correspondence to the Verilog HDL code. All future modifications will be done simultaneously in the model and Verilog HDL code, keeping the correspondence intact.

VI.   SUMMARY

The conceptual design of a Track-Finder for the Level-1 trigger of the CMS endcap muon system is complete. The design is implemented as 12 identical processors, which cover the pseudo-rapidity interval $0.9 < \eta < 2.4$. The track-finding algorithms are three-dimensional, which improves the background suppression. The P_T measurement uses data from 3 endcap stations, when available, to improve the resolution to 20%. The input to the Track-Finder can be held for more than one bunch crossing to accommodate timing errors. The latency is expected to be 7 bunch crossings (not including the optical link and timing errors

accommodation). The design is implemented using Xilinx Virtex FPGAs and SRAM look-up tables and is fully programmable. The first prototype was successfully built and tested; the pre-production prototype is under construction now.

VII.   REFERENCES

[1]. D. Acosta *et al*. "The Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System." Proceedings of the LEB 1999 Workshop.

[2] National Semiconductor, DS90CR285/286 datasheet.

[3] CMS Level-1 Trigger Technical Design Report, section 12.4.
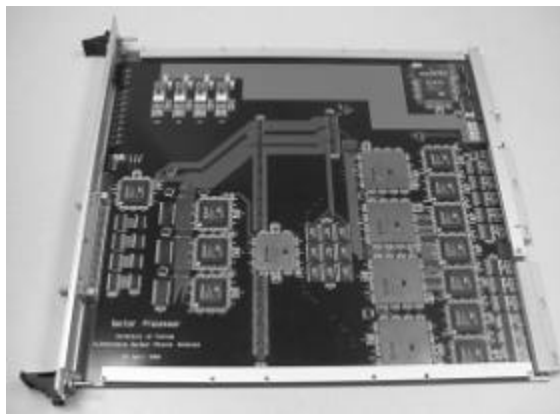
[4] Xilinx Inc., www.xilinx.com
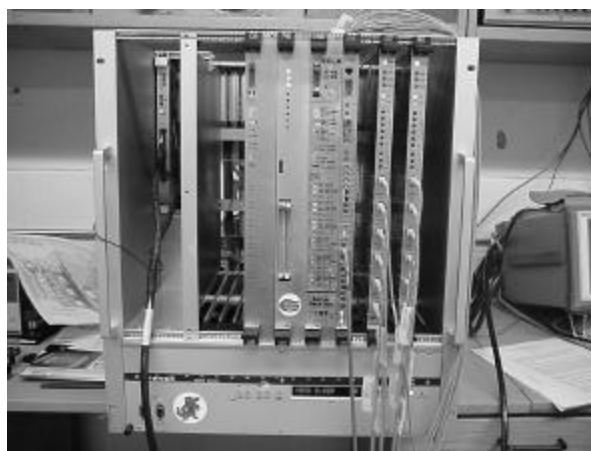
Figure 1: The first prototype of Sector Processor
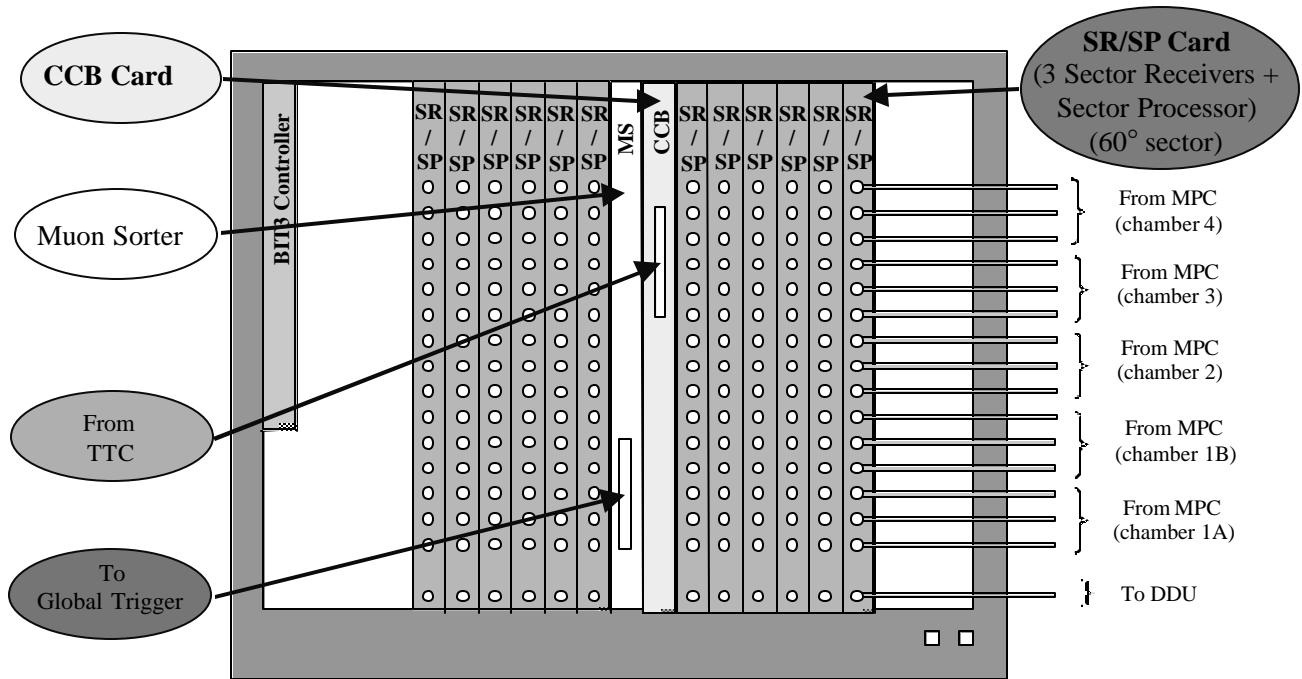

Figure 2: Prototype test crate.

CCB Card

SR/SP Card
(3 Sector Receivers +
Sector Processor)
(60° sector)

Muon Sorter

From
TTC

To
Global Trigger

BTB Controller

SR / SP   SR / SP   SR / SP   SR / SP   SR / SP   SR / SP   MS   CCB   SR / SP   SR / SP   SR / SP   SR / SP   SR / SP   SR / SP

From MPC
(chamber 4)

From MPC
(chamber 3)

From MPC
(chamber 2)

From MPC
(chamber 1B)

From MPC
(chamber 1A)

To DDU

Figure 3: Second Prototype Track-Finder Crate.

First prototype data flow

From Muon Port Cards

	Optical receivers	Sector Receiver st.1	Sector Receiver st.2,3	Sector Receiver st.4
1	Front FPGAs			
1	Lookup tables			
4	Channel link transmitters			

	Channel link receivers	Sector Processor
2	Bunch crossing analyzer (not implemented)	
3	Extrapolation units	
2	9 Track Assembler units (memory)	
3	Final selection unit 3 best out of 9	
3	Pt precalculation for best 3 muons	
2	Pt assignment (memory)	

Latency

Total: 21

To Muon Sorter

Pre-production prototype data flow

From Muon Port Cards

1	Optical receivers
1	Front FPGAs
1	Lookup tables

SP/SR board

0	Bunch crossing analyzer (not implemented)
1	Extrapolation units
1	9 Track Assembler units

Sector Processor FPGA

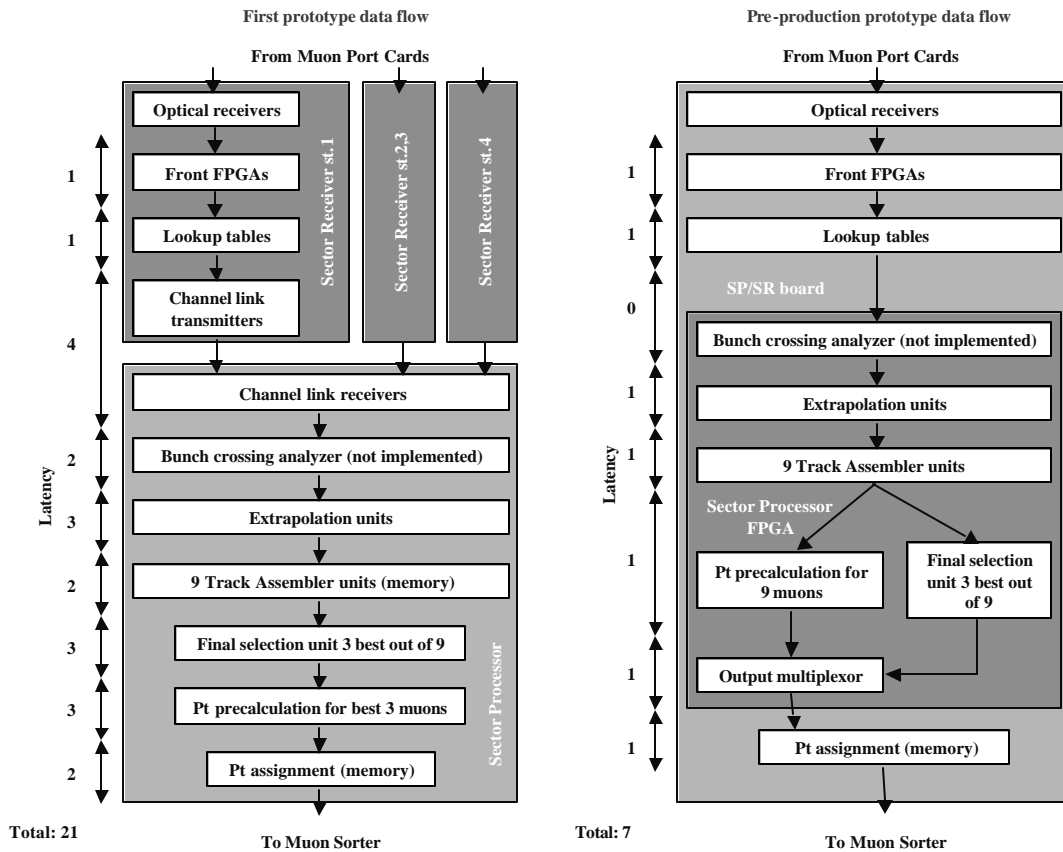| 1 | Pt precalculation for 9 muons | Final selection unit 3 best out of 9 |
| 1 | Output multiplexor | |
| 1 | Pt assignment (memory) | |

Latency

Total: 7

To Muon Sorter

Figure 4: Comparison of the first and pre-production prototypes.