Performance of a Pre-Production Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System

D.Acosta, L.Gray, A.Madorsky*, N.Park, B.Scurlock, H.Stoeck

Department of Physics, University of Florida, Gainesville, FL 32611, USA

V.Golovtsov, M.Kan, L.Uvarov

High Energy Physics Division, Petersburg Nuclear Physics Institute, Gatchina, Leningradskaya oblast, 188300, Russia

*presented by Alex Madorsky (madorsky@phys.ufl.edu)

Abstract

We report on the performance of a track-finding processor for the CMS Level-1 trigger that performs 3-D track reconstruction and measurement from data recorded by the cathode strip chambers of the endcap muon system. A pre-production prototype has been constructed and successfully tested. The processor design is implemented using Xilinx Virtex-2 FPGAs and SRAM look-up tables and is fully programmable. Track segment data are received by the processor using optical links operating at 1.6 Gb/s. Detailed results on the electronics tests and the performance of the processor are reported. The processor was included into the trigger path of a beam test of several cathode strip chambers, and provided a track trigger for the experiment. The beam was delivered with a 25 ns LHC bunch structure at CERN, and the processor maintained synchronization with the delivered machine clock.

I. INTRODUCTION

The endcap regions of the Compact Muon Solenoid (CMS) experiment will consist of four stations of Cathode Strip Chambers (CSC) [1]. These chambers will provide CMS complete azimuth coverage (in ϕ), as well as 0.9 to 2.4 in pseudo-rapidity (η). Six cathode strip chambers compose a single station in the endcap system. The chambers are trapezoidal in shape, extending 10^o or 20^o in ϕ, and are composed of cathode strips aligned radially from the beam axis, and anode wires aligned in the orthogonal direction.

In the endcap system, muon track-finding is electronically partitioned into six 60^o sectors in each endcap. A single Sector Processor (SP) receives trigger primitives from front-end electronics [2] which sit on or near the CSCs. The front-end electronics form Local Charged Tracks (LCTs) from the six detector layers of a station. A Muon Port Card (MPC) collects the LCTs for a given station and sector, sorts them, and sends the best three to an SP via optical fibers. A single SP collects LCTs sent via fifteen 1.6 Gbit/s optical links, and is responsible for linking LCTs in ϕ and η in order to form full tracks, and to report the transverse momentum (p_t), ϕ, and η for each track. The entire Track-Finding processor is composed of twelve such SPs housed in a single 9U VME crate. One Muon Sorter sorts up to 36 possible tracks and reports the 4 highest rank to the Global Trigger. The challenge for the Track-Finding Processor is to report muon candidates with the lowest possible Pt threshold, and yet maintain a single muon trigger rate below 1 kHz/η at full Large Hadron Collider (LHC) luminosity.

This paper is organized as follows: Section II will describe the Track-Finder (TF) algorithm, Section III will discuss pre-production prototype architecture, Section IV will discuss test results, and a summary can be found in SectionV.

II. TRACK-FINDER LOGIC

The principle of the TF logic [3] is illustrated in Figure 1. The Track-Finding process is partitioned into several steps. Each MPC may send as many as three LCTs reported to the SP. These LCTs are then converted into track-segments, which are described in terms of their ϕ, and η coordinates. Each track-segment in each station should be checked against the other segments in neighboring stations for consistency to share a single track. Thus, each track-segment is extrapolated through to other stations, and compared against existing segments. If an extrapolation is successful, these segments are "linked" to form a single track. Each possible pairwise combination is tested in parallel. After extrapolation, doublets are then linked to assemble full tracks. Redundant tracks are cancelled, the best three tracks are selected, and the track parameters are then measured.

The SP has the ability to handle LCTs received out of step from the actual bunch crossing time in which they originated. This is accomplished by the Bunch Crossing Analyzer, which allows the Sector Processor to form tracks from LCTs received up to one bunch crossing later than the earliest LCT.

The first step in the track-finding process is to extrapolate pairwise combinations of track-segments. This is accomplished by requiring the two segments to be consistent with a muon originating from the collision vertex and with an appropriate curvature induced by the non-uniform magnetic field. A successful extrapolation is assigned when two stubs lie within allowed windows of ϕ and η - neither LCT should be parallel to the beam axis, and both should appear to originate from the interaction region.

The Track Assembler Units (TAUs) examine successfully extrapolated track-segment pairs to see if a larger track can be formed. If so, those segments are combined and a code is assigned to denote which muon stations are involved.

A list of nine possible tracks is sent to the Final Selection Unit (FSU). Since different data streams may contain data

Figure 1: Sector Processor logic

from the same muons, the FSU must cancel redundant tracks, and select the best three distinct candidates.

The final stage of processing in the TF 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 selections based on this quantity have a direct impact on the trigger rate and on the efficiency. The technique used incorporates the ϕ information from up to 3 stations when it is available. This yields a resolution of approximately 22% at low momenta, which is sufficient. In order to achieve a three-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 precision on the second subtraction because we are only trying to untangle the multiple scattering effects at low momenta. The subtraction results are combined with the η coordinate of the track and the track type, and then sent into a 4 MB memory for assignment of the signed P_T. Tracks composed of only two track segments are allowed also in certain cases.

III. Pre-Production Prototype System Architecture

The Sector Processor board accepts fifteen optical links from the Muon Port Cards, where each link carries information corresponding to one muon track segment, which is described by a 32-bit word (two 16-bit frames). Each link transmits a 32-bit word every 25 ns. Finisar's FTRJ-8519-1-2.5 transceiver is used by the MPC for transmission, and by the SP for reception of LCT data. Texas Instrument's TLK2501 transceiver is used to deserialize data received by the optical receivers. Additionally, the board receives up to 4 muon track segments, sent using LVDS synchronous with the 40 MHz LHC clock, from the Barrel Muon system via a transition card behind the custom backplane. Because the track segment information arrives from 15 different optical links as well as the barrel system, it is aligned to the proper bunch crossing number.

The Sector Receiver algorithm is implemented in a series of cascaded look-up memories, using 45 GSI SRAMs [4], in order to minimize the size of SRAM chips required: chamber specific LCT data is sent to the first memory, this memory then sends a local ϕ measurement to two more memories which then use this local ϕ measurement along with the

original LCT data to form a global ϕ and η measurement for the given sector.

The angular information for all track segments is then passed to the main XC2V4000 FPGA, which executes the entire 3-dimensional tracking algorithm of the Sector Processor. This FPGA sits on a mezzanine card on the SP, thus allowing for maximum design flexibility for future improvements. In the first prototype this algorithm required 15 FPGAs.

The output of the Sector Processor FPGA is sent to the P_T assignment lookup tables, also GSI SRAMS [5], and the results of the P_T assignment for the three best muons are sent via the custom backplane to the Muon Sorter using GTLP 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 time-critical portions of data are sent first. The SP will also send data to a data acquisition readout board over the SP's sixteenth optical link. This board receives the LCT data received by the SP from the MPC, and final results completed by the SP. The entire second prototype Track-Finder system fits into one 9U VME crate.

In order to ensure an exact match between the SP firmware functionality and C++ simulation, a class library has been developed that allows one to write both the simulation code and firmware in C++, and then translate this code into Verilog HDL. Thus, our code serves a dual purpose: it can be compiled either for simulation purposes, or as the generator of the Verilog code used for FPGA programming. This guarantees a bit-for-bit compatible simulation. This Verilog code can then be synthesized by our FPGA vendor tools and is used as our SP Firmware. This allows us to verify the SP logic through standard C++ debugging tools. We can also run this code as a part of the CMS simulation and reconstruction; therefore, a line-by-line correspondence is maintained between simulation logic and Firmware logic.

IV. TESTS

A. *On-Board Functionality*

A pre-production prototype of the Sector Processor has been fabricated, and extensive tests to verify its functionality have been performed. The VME interface has been verified; firmware as well as LUT contents have been successfully downloaded and read back. The LUT data is derived from track-finder simulations; however, random number sequences have also been successfully loaded. Successful tests have also been completed to verify operation of the TLK2501 chips on the SP. These chips have the ability to generate pseudo random bit streams (PRBS) to facilitate testing of a single chip (internal loop-back test), or pairs of chips. In transmitting these bit streams over optical fibres, one optical link is used to send data, and another is used to receive data. Because the SP is equipped with optical transceivers, one can transmit and receive these bit streams on the same SP by utilizing multiple links. Such tests were also completed between a Muon Port Card and a Sector Processor using an external clock source. No errors were observed during PRBS tests. In addition to PRBS tests, routines were used which load test LCT patterns into the MPC input buffer, and transmit these patterns over the optical links into the SP. The output LCTs from the MPC were checked against the SP input LCTs, and were found to

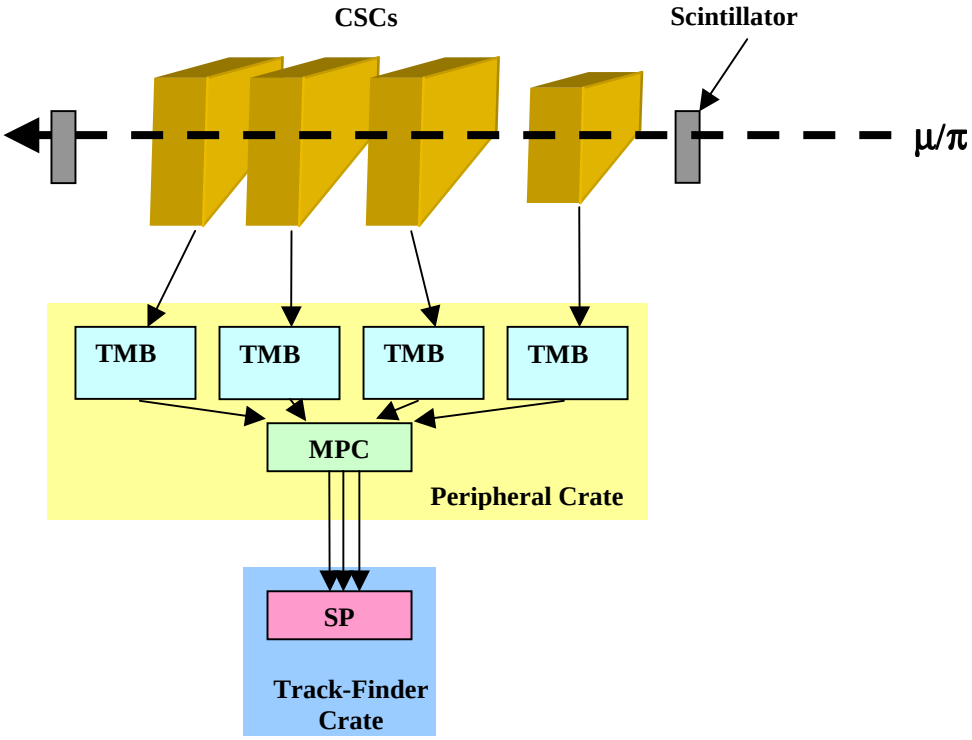

Figure 4: Conceptual layout of test beam stand. Arrows depict data flow.

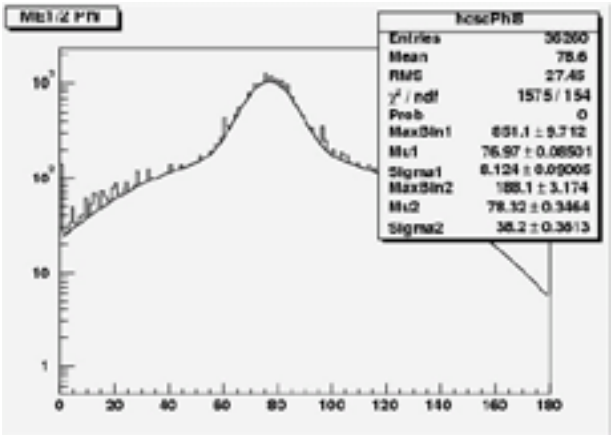

Figure 5: Distribution of LCT ϕ position reported by CSC front-end electronics and recorded by the SP.

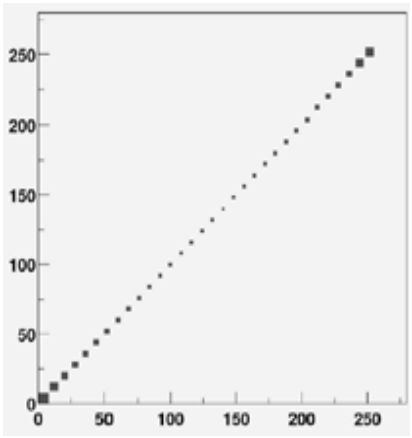

Figure 6: Difference in ϕ position for two CSCs as reported by hardware versus that reported by software.

be in agreement.

The SP also has the ability to load test LCT patterns into input FIFOs in the front FPGAs. These patterns can then be sent to the front LUTs and through the full SP logic. Output FIFOs can store data at each step of the track-finding process. Using this technique, we are able to fully monitor and simulate data flow using our simulation package to compare data against SP hardware functionality. The software tools used for these tests can also be used to load actual LCT data from the CSCs into the test FIFOs. This method has allowed us to verify the correct operation of the track-finder.

B. System Tests

The full electronic chain, up to the Sector Processor, underwent tests at CERN in May 2003. Both muon and pion beams were provided by the SPS, and illuminated CSCs in both time-structured and unstructured runs. During the time-structured runs, the SPS delivered particle bunches separated by approximately 25ns; the delivery of these bunches was synchronized to the SPS machine clock and distributed to all modules in the trigger electronics chain in order to simulated actual LHC running conditions.

Optical data transmission between the MPC to SP was hindered because of jitter in the clock signal distributed by the TTC system [6]. As a result, optical transmission errors were observed by the SP. This problem was resolved by adding a VCXO-based PLL on board the SP. The PLL allowed for a cleaner 80 MHz clock to be delivered to the SP's optical receivers. Thus, during subsequent testbeams in September 2003 and May 2004, no optical transmission errors were observed.

During the 2004 testbeam, all electronic modules from the SP down to the front-end electronics were used; however, the experiment only approximated the full electronic data flow during normal CMS running (Figure 4). Each chamber represented an entire CMS sector, but all chamber data was collected by either one or two MPCs, depending on the run mode, which then sent data to the SP. Since a single MPC is supposed to send data for a given station, the optical link on which a particular chamber's data was delivered to the SP was uncertain. Despite this complication, track-finding was possible by re-programming the SP's front LUTs (see Figure 5). This allowed the regional coordinates to be adjusted such that the track-finding logic received track-stubs between chamber pairs which satisfied conditions on ϕ and η limits, as well as $\Delta\phi$ and $\Delta\eta$ windows, thus allowing successful extrapolations. Because the SP was able to link tracks, the full system was triggered using the SP rather than scintillators.

LCT data sent through the Data Acquisition path was compared to LCT received over optical links by the SP. Approximately 99.7% of all data found by the SP was also recorded in the DAQ path. Because the SP only received three out of a possible six LCTs per event, LCT data from the DAQ event log had to be sent through the MPC emulation in order to compare the appropriate LCTs to those received by the SP optical links.

In addition to verifying data flow to the SP electronics, the SP's track-finding functionality was also tested against software emulation. The SP input data, as well as reconstructed track data, was logged to disk. The input data was then sent through the software model, and 100% agreement was found between hardware and software performance (see Figure 6).

Muon Sorter (MS) functionality was also tested during the testbeam period. The MS received reconstructed track data from the SP and sorts them according to quality. Winner bits were then sent back to the SP, and added to the event record in the output FIFO. These winner bits were checked against the software model, and found to agree with logged data.

V. SUMMARY

The design of a Track-Finder for the Level-1 trigger of the CMS endcap muon system is mature and has been successfully prototyped. 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 approximately 22%.

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 pre-production prototype has been built and has been successfully tested in multiple test beam experiments. The track-finder successfully provided a trigger for the May 2004 test beam.

VI. REFERENCES

[1] CMS, The Muon Project Technical Design Report, CERN/LHCC 97-32

[2] CMS Level-1 Trigger Technical Design Report, section 12.4. CERN/LHCC 2000-038

[3] 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.

[4] GSI G58320Z18T datasheet

[5] GSI GS881Z18AT datasheet

[6] TTC website: http://ttc.web.cern.ch/TTC/intro.html