

THE TRACK-FINDING PROCESSOR FOR THE LEVEL-1 TRIGGER OF THE CMS ENDCAP MUON SYSTEM

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

University of Florida

A. Atamanchuk, V. Golovtsov, B. Razmyslovich

St. Petersburg Nuclear Physics Institute

Abstract

We present the design of a fast three-dimensional track-finding processor for the Level-1 trigger of the CMS Endcap Muon System. Each track is assembled from track segments identified in the Endcap Muon chambers. The processor measures the transverse momentum of the best muon candidates from the sagitta induced by the magnetic bending. The trigger algorithms are programmable since the processor is based on FPGA and RAM logic, allowing the experiment to adapt to different background conditions and colliding rates.

1. 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 [1]. 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.

A version of the muon Track-Finder designed for the barrel muon system of CMS is under design [2]. Since it is intrinsically two-dimensional, it does not satisfy the requirements of the endcap system. Moreover, the large number of signals shared between trigger regions (because of the staggering of the barrel chambers) is not appropriate for the endcap. Therefore, we have rethought the entire Track-Finder architecture given that we are compelled to design a new trigger processor.

2. SYSTEM ARCHITECTURE

The Track-Finder is foreseen be 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 [3] are 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 under design.

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 is currently under development, and a preliminary sketch of the functional layout on a 9U VME board is shown in Fig. 1.

3. 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. The steps are pipelined and the trigger logic is deadline-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*.

3.1 Data Input

The input data to the Sector Processor is de-serialized by Channel Link receivers and synchronized with the local 40 MHz clock before being sent into the Extrapolation Units. The Sector Processor is designed to have some ability to analyze track segments received on different clocks (different bunch crossings) because of the intrinsic timing spread of the trigger primitive formation.

To incorporate a multi-bunch mode, we take advantage of the sparseness of the data. Track segments from other bunch crossings are considered only if there are empty track segments in the current crossing; otherwise, the size of the extrapolation logic would grow enormously. We assume that track segments may be late (but not early), and that the earliest track segment defines the beam crossing of the muon. The simplest case is a window that is open for two bunch crossings,

which should be sufficient for the CSC trigger system. Although the window is left open for more than bunch crossing, the Sector Processor must report triggers at the correct bunch crossing every crossing. In other words, overlapping time buckets are used. A flag is set to record which bunch crossing the track segment came from relative to the current crossing. This flag is used in the final track selection to inhibit double triggers.

3.2 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

Each test is accomplished by small look-up tables or arithmetic logic, as illustrated in Fig. 2 The resulting quality word is either 1 or 2 bits, depending on the stations involved. Its definition is programmable, but we use it to assign a coarse P_T (low, medium, and high) to extrapolations involving the first muon station. Otherwise, the quality just represents whether the extrapolation was successful or not. This quality information is useful when track candidates must be ranked.

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.

The extrapolation logic should be programmable, and is presently implemented in the Xilinx Virtex series [4] field-programmable gate-arrays (FPGAs). The latency to accomplish this decision logic is expected to take 3 clocks.

3.3 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.

Each track segment of a key station, of which there are three for one 60° sector, is tested for extrapolations to the other stations. There are six track segments allowed from the first CSC muon station, and the extrapolation quality is 2 bits. There are three track segments allowed in each of the other non-key stations, and the extrapolation quality to those stations is 1 bit. Thus, 18 bits are interrogated. (The same is true for extrapolations to the barrel muon system.) Since the number of input bits is small, each of these “Link” units is implemented as a static RAM look-up memory as shown in Fig. 3. The latency, therefore, is just one clock. The output code is a 6-bit quality word giving the type and rank of the *best* assembled track, and a 9-bit word labelling the track segments used in each station. A memory of the same size is used for extrapolations to the barrel muon system as well.

3.4 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 Selection Unit is a sorter with cancellation logic.

A VHDL description of the Final Selection Unit has been synthesized using the Exemplar Logic engine of OrCAD Express [5]. The sorter part of the logic compares the qualities of all pairwise combinations of

tracks from the Track Assembler data streams. The cancellation part of the logic does the same for the hit labels. Not all track segments need to be identical for two tracks to be considered identical. Bremsstrahlung, for example, might cause a single muon to deliver two track segments in one station, and this would lead to a fake di-muon trigger which should be suppressed. The actual criterion employed is programmable. The two comparison steps are done in parallel in one 25 ns clock. The next step of the logic examines the results of all these comparisons, cancels redundant tracks, and reports the identities of the three best and distinct muons. These two steps take two clocks. Finally, the track segment information of the selected muons is taken from a multiplexer and transmitted to the Assignment Units of the measurement system. Additional logic connected to the multiplexer determines if all track segments of a given muon come from a later bunch crossing, in which case the muon is suppressed to inhibit double triggers.

3.5 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. We take advantage of the large multiple scattering for low P_T muons. Although there is a small probability that a scattering will offset the large magnetic bending between the first two stations (and thus appear as a high momentum muon), it is much less likely to offset the bending between all three stations.

In order to achieve a 3-station P_T measurement, one must be careful not to include too much data; otherwise, the size of the look-up memories will be prohibitive. 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. This scheme is illustrated in Fig. 4.

4. 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, and the full latency is expected to be 14 bunch crossings. The design is implemented using Xilinx Virtex FPGAs and SRAM look-up tables and is fully programmable. The board layout and backplane design has begun.

REFERENCES

- [1] J. Hauser, these proceedings.
- [2] J. Ero, these proceedings.
- [3] National Semiconductor, DS90CR285/286 datasheet.
- [4] Xilinx Inc., <http://www.xilinx.com/products/virtex.htm>
- [5] OrCAD Inc., http://www.orcad.com/products/express/express_f.htm

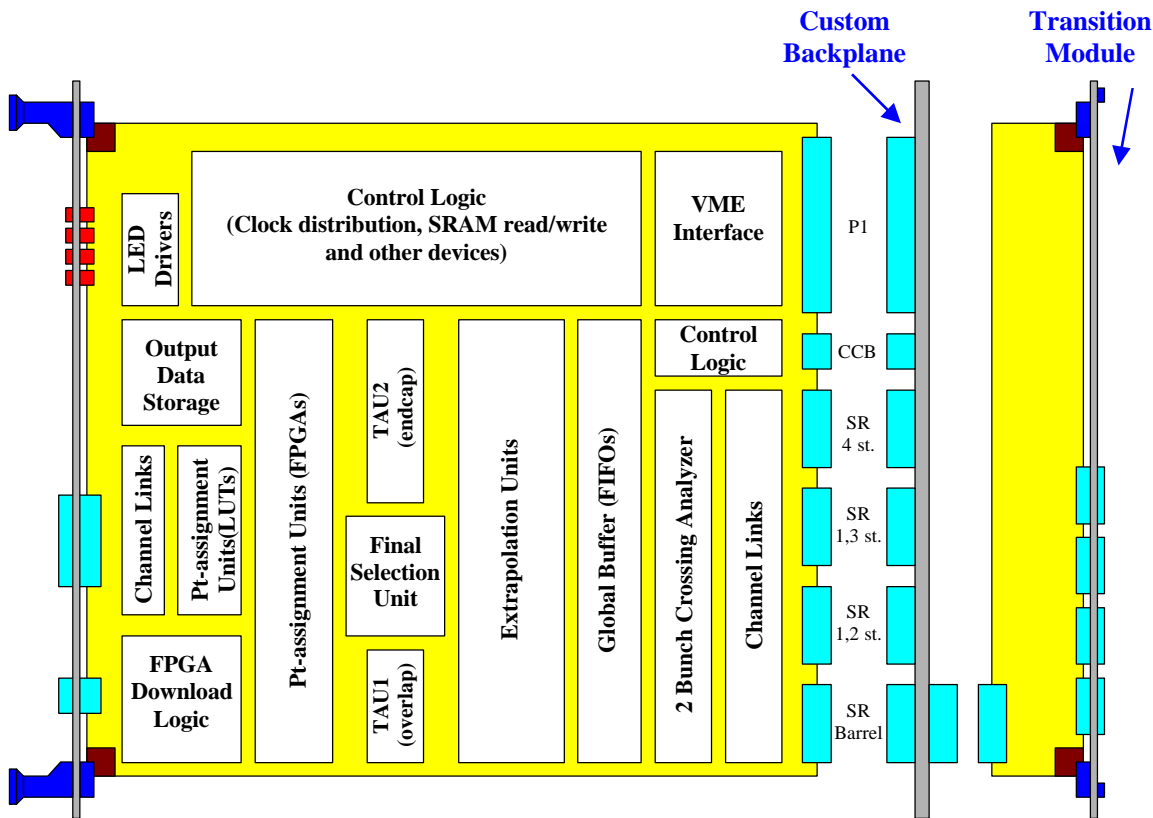


Figure 1: Functional layout of the Sector Processor on a 9U VME board showing the major components of the Track-Finder logic.

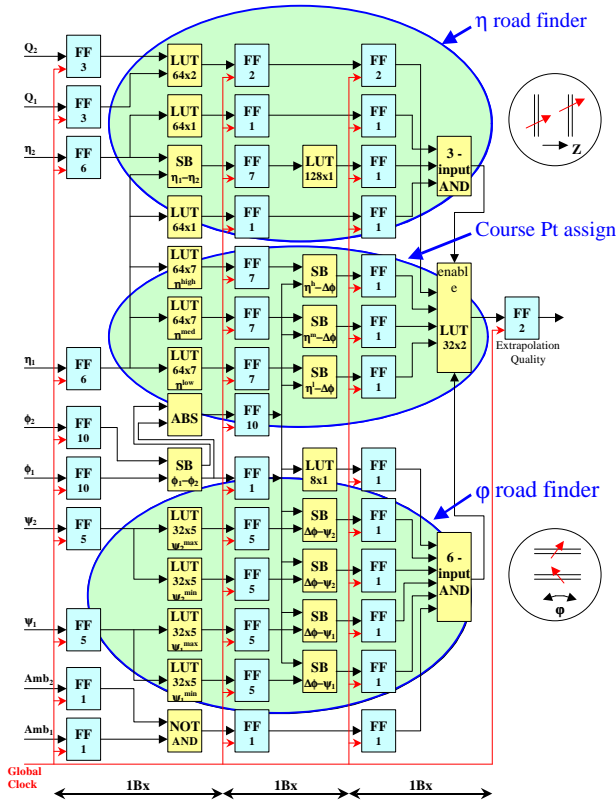


Figure 2: Logic diagram of a single extrapolation.

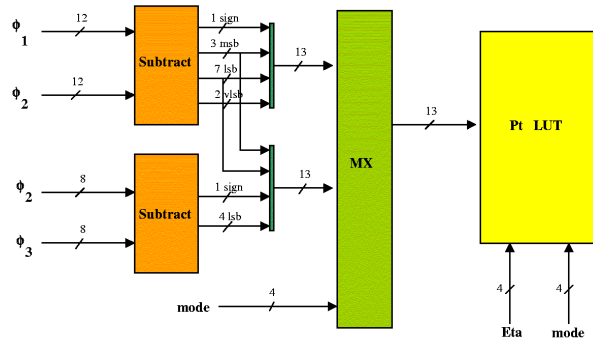


Figure 4: The assignment of P_T using FPGA preprocessing followed by an SRAM look-up table.

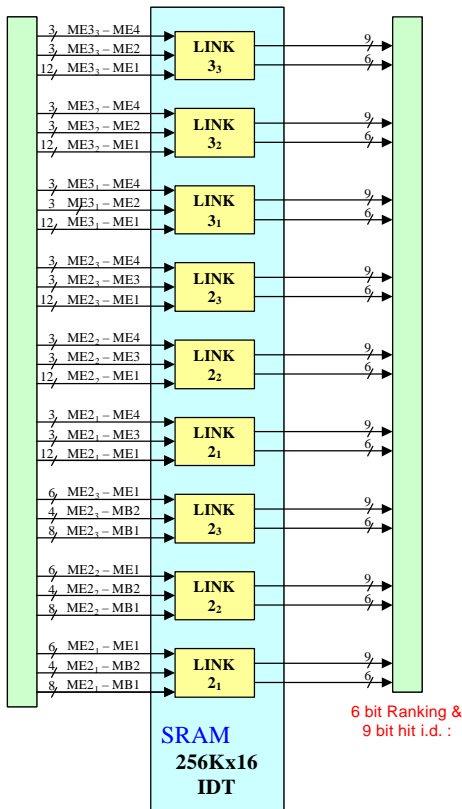


Figure 3: The Track Assembler Unit implemented as 9 static RAM memories.