


The Compact Muon Solenoid Experiment
Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



14 October 2009 (v2, 16 October 2009)

Design Considerations for an Upgraded Track-Finding Processor in the Level-1 Endcap Muon Trigger of CMS for SLHC Operations

Alexander Madorsky for the CMS Collaboration

Abstract

D. Acosta, M. Fisher, I. Furic, J. Gartner, G.P. Di Giovanni, A. Hammar, K. Kotov, A. Madorsky, D. Wang

University of Florida/Physics, POB 118440, Gainesville, FL, USA, 32611

L. Uvarov

Petersburg Nuclear Physics Institute, Gatchina, Russia

M. Matveev, P. Padley

Rice University, MS 61, 6100 Main Street, Houston, TX, USA, 77005

The conceptual design for a Level-1 muon track-finder trigger for the CMS endcap muon system is proposed that can accommodate the increased particle occupancy and system constraints of the proposed SLHC accelerator upgrade and the CMS detector upgrades. A brief review of the architecture of the current track-finder for LHC trigger operation is given, with potential bottlenecks indicated for SLHC operation. The upgraded track-finding processors described here would receive as many as two track segments detected from every cathode strip chamber comprising the endcap muon system, up to a total of 18 per 60 degree azimuthal sector. This would dramatically improve the efficiency of the track reconstruction in a high occupancy environment over the current design, since some of the track segments are filtered out in order to reduce transmission bandwidth and track processing logic. However, such an improvement would require significantly higher bandwidth and logic resources over the current design. We propose to use fastest available serial links, running asynchronously to the machine clock. Another enhancement critical for the overall Level-1 trigger capability for physics studies in phase 2 of the SLHC is to include the inner silicon tracking systems into the design of the Level-1 trigger. This requires matching muons identified in the endcap muon system and matching them to hits in the inner tracking system and refining the momentum measurement to improved precision for better rate reduction capabilities. Some preliminary ideas on the precision of information available from the endcap track-finder trigger will be presented along with possible algorithms for the matching.

Design Considerations for an Upgraded Track-Finding Processor in the Level-1 Endcap Muon Trigger of CMS for SLHC operations

D. Acosta^a, M. Fisher^a, I. Furic^a, J. Gartner^a, G.P. Di Giovanni^a, A. Hammar^a, K. Kotov^a, A. Madorsky^a, M. Matveev^c, P. Padley^c, L. Uvarov^b, D. Wang^a

^aUniversity of Florida/Physics, POB 118440, Gainesville, FL, USA, 32611

^bPetersburg Nuclear Physics Institute, Gatchina, Russia

^cRice University, MS 315, 6100 Main Street, Houston, TX, USA, 77005

madorsky@phys.ufl.edu

Abstract

The conceptual design for a Level-1 muon track-finder trigger for the CMS endcap muon system is proposed that can accommodate the increased particle occupancy and system constraints of the proposed SLHC accelerator upgrade and the CMS detector upgrades. A brief review of the architecture of the current track-finder for LHC trigger operation is given, with potential bottlenecks indicated for SLHC operation. The upgraded track-finding processors described here would receive as many as two track segments detected from every cathode strip chamber comprising the endcap muon system, up to a total of 18 per 60° azimuthal sector. This would dramatically improve the efficiency of the track reconstruction in a high occupancy environment over the current design. However, such an improvement would require significantly higher bandwidth and logic resources. We propose to use the fastest available serial links, running asynchronously to the machine clock to use their full bandwidth. The work of creating a firmware model for the upgraded Sector Processor is in progress; details of its implementation will be discussed. Another enhancement critical for the overall Level-1 trigger capability for physics studies in phase 2 of the SLHC is to include the inner silicon tracking systems into the design of the Level-1 trigger.

I. CMS ENDCAP MUON LEVEL-1 TRIGGER SYSTEM OVERVIEW

The CMS Endcap Muon system consists of 540 six-plane cathode strip chambers¹. Strips, milled on the cathode panels, run radially in the endcap geometry and thus provide a precise measurement of the ϕ -coordinate. Wires are stretched across strips and define the radial coordinate of muon hits.

A. Generation of Trigger Primitives

Electronic components responsible for the generation of trigger primitives include:

- Cathode Front End Board (CFEB), 5 per chamber
- Anode Local Charged Track board (ALCT), 1 per chamber
- Trigger Mother Board (TMB), 1 per chamber

The CMS Endcap Muon system is comprised of two endcaps. Each endcap consists of 4 layers of Cathode Strip

Chambers (CSCs); these layers are commonly called “stations”. Station ME1 is the closest to the Interaction Point (IP), station ME4 is the farthest.

For the purposes of Trigger system, each endcap is subdivided into six 60° sectors. Each sector is served by one Sector Processor (SP) board; there are 12 SPs in the Endcap Muon Trigger system. Each SP is implemented as a 9U VME board; all SPs are housed in one VME crate that is located in the CMS Underground Support Cavern (USC55).

The TMB associated with each chamber can provide up to two trigger primitives on any bunch crossing. Each trigger primitive contains the following information:

- Cathode hit coordinate (half-strip number)
- Cathode pattern type (measure of the track bend angle)
- Anode hit coordinate (wiregroup number)
- Anode pattern type (collision or halo track)
- Trigger primitive quality

The trigger primitives generated by TMBs are delivered to Muon Port Cards (MPC), also located in the Peripheral Crates. There is one MPC per station (9 chambers), except station 1 that has 2 MPCs because there are 18 chambers in it. Each MPC receives up to 18 trigger primitives per bunch-crossing (BX). The MPC selects the best three trigger primitives out of 18, and sends them via 1.6 Gbps optical links to the Sector Processor.

B. Track reconstruction in Sector Processor

The Sector Processor (SP) receives trigger primitives from MPCs associated with all stations in a specific sector, for a total of up to 15 primitives per BX. In addition to that, the Barrel Muon system (Drift Tube Chambers, or DT) delivers up to two trigger primitives from the region where it overlaps with the Endcap Muon system. If one or two more DT trigger primitives are available at the same BX, they can be delivered with a delay of one clock cycle.

Track reconstruction involves the following hardware modules:

1) Conversion of raw trigger primitives into geometrical parameters.

In the current design, the conversion of raw trigger primitives into ϕ and η (pseudorapidity) is performed using

¹ 468 chambers installed and operational and 72 additional chambers (ME4/2) to be fabricated and installed.

large 2-stage look-up tables (LUTs). The amount of memory required to convert a single trigger primitive is around 4MB.

2) Multiple Bunch Crossing Analysis (BXA)

Cathode Strip Chambers may not report all the trigger primitives related to a certain track at the same precise BX; some trigger primitives are delivered with a delay of one or even two BXs because of charged particles drift time inside the chamber or imperfect synchronization. In order to build a track that has such delayed trigger primitives, the SP needs to analyze up to 2 BXs in addition to the current one. The BXA keeps the history of trigger primitives belonging to two previous BXs. All trigger primitives (current and delayed, total of 9) from each station are sorted on each BX, and best three primitives are sent for further processing. This ensures that the tracks are built taking the highest quality primitives into account.

3) Extrapolation Units (EUs)

Each EU checks that ϕ and η parameters of two trigger primitives from two different stations (A and B) are within certain limits (windows) from each other.

In the current Track-Finder design, almost all possible combinations of stations have to be extrapolated; this brings the total number of extrapolations² to 210. In addition, the EUs for the ME1-ME2 and ME1-ME3 extrapolations provide a 2-bit extrapolation quality based on the ϕ difference between the trigger primitives.

4) Track Assembly Units (TAUs)

Each TAU takes one particular trigger primitive from ME2, ME3, and ME4, and tries to find as many valid extrapolations as possible to other stations. If the search is successful, TAU reports a possible track candidate. There are 12 TAUs for collision tracks and 6 for halo tracks (accelerator produced muons outside the beam pipe). Each track candidate receives a rank that encodes stations and extrapolation qualities used to construct it. The rank reflects the “quality” of the track candidate – the higher that number is, the more stations have participated in the track.

5) Transverse Momentum (P_t) Assignment Units (PAU)

The tracks assembly results from available primitives are delivered to PAUs. There is one P_t Assignment Unit per TAU. These units identify the track segments used to build each track candidate, assign ϕ and η parameters (taken from the best available track segments) to track candidates, and calculate the ϕ difference for the best available 2 or 3 stations. On the output, they provide the address for the P_t Assignment Lookup Table (P_t LUT).

6) Final Selection Unit (FSU)

There are two FSUs: one for collision and one for halo tracks. Each FSU receives the ranks of all track candidates (12 collision or 6 halo candidates). FSU keeps a history of track

candidates 2 BXs in the past, and selects the best three collision tracks or 1 best halo track out of all available candidates. Simultaneously, it checks for tracks that have η and ϕ parameters close to each other. If such tracks are found, only one of them having the highest rank is left; all others are removed. This $\eta+\phi$ track cancellation is necessary because TAUs sometimes may produce different track candidates that correspond to a single physical track. One more reason for the cancellation is chamber drift time (see BXA unit description above). This leads to multiple track candidates created over a duration of up to 3 BXs, so taking the track candidate history into account becomes necessary to find the best tracks.

7) Output Multiplexer (OM)

The results of the final selection are delivered to the OM. This module passes the track parameters of the best tracks selected by FSUs to its outputs. Priority is given to collision tracks. A halo track (if found) is multiplexed to the first unused output.

8) BX Correction Unit (BXC)

The final step in the Track-Finder logic is Bunch-Crossing number correction. For the best performance the timing for a track should be set to the BX when the second trigger primitive for it was received. The BXC is applying variable delay to the output tracks to make sure this timing requirement is satisfied.

9) P_t Assignment Lookup Table (P_t LUT)

The P_t LUT is a separate hardware module implemented as memory IC. The address of this memory is provided by the SP logic and is formed by P_t Assignment Units (see PAU description above). The output includes track P_t encoded into 5-bit value, track quality (2-bit value), and “valid” flag.

II. TARGETING SLHC

The current design of the CMS CSC Endcap Track-Finder is totally adequate up to the current LHC design luminosity. However, for the SLHC operation, there are a number of problems that have to be addressed. This section lists these problems and proposed solutions.

A. MPC filtering

Currently, the MPC selects the best three trigger primitives out of 18 available. However, with a luminosity upgrade to $L=10^{35} \text{ cm}^{-2}\text{s}^{-1}$, we can expect at least 7 trigger primitives per BX in every MPC. This number is based on simulations [1], and in reality could be higher.

Our current intention is to design an upgraded Track-Finder that can process all available trigger primitives (2 per chamber, or 18 per MPC). This would allow us to reduce significantly the dependence on background hits in the CSCs, the rate of which is unknown at this time for both LHC and SLHC.

B. Optical link bandwidth

Trigger primitives are delivered from MPCs to SPs using 1.6 Gbps optical links. To deliver 18 trigger primitives instead

² ϕ and η extrapolations are counted separately. The number shown is for SP with mezzanine card upgraded in 2008, and does not include halo extrapolations.

of 3, we will need 6 times more bandwidth than we have now. To accommodate that, data links with larger throughput have to be used.

We are considering two options: faster optical links working at a higher bit rate (10 Gbps), or multi-channel links running at a moderate bit rate (1.6 to 2.4 Gbps). Both options seem to be suitable for our purposes. The 10 Gbps links require fewer fibers but have to be run asynchronously to the machine clock to reach full bandwidth. The parallel links can be run in “traditional” mode (synchronous to the machine clock), but require special multi-core fibers and more serializer-deserializer pairs.

Removing the MPC trigger primitive filtering and upgrading the optical links will require a complete MPC redesign and a system-wide replacement (60 boards).

C. Trigger primitive conversion to angular coordinates.

Currently, this conversion requires 4MB of memory per primitive, which is unacceptable for the upgraded design. We plan to use FPGA logic combined with much smaller LUTs implemented inside the FPGA. The fact that we plan to receive trigger primitives from all chambers means that chamber numbers do not have to be explicitly analyzed during the conversion, which leads to savings in logic and LUT size.

1) Coordinate systems

The angular coordinates that were used in the current SP design are not very convenient. For example, the ϕ coordinate uses 4096 values per 62° sector, which is $\sim 0.015^\circ$ per ϕ unit. The corresponding angular coordinate in trigger primitives is the half-strip number, with unit value of 0.06665° for the majority of chambers. If the ϕ scale is selected that has the unit value of $0.06665/4 = 0.0166625^\circ$, the half-strip to ϕ conversion for most of the chambers becomes as simple as adding or subtracting one value and then adding two least significant bits.

The wiregroup number arriving with trigger primitives is currently converted into an η coordinate. This is also not the optimal coordinate for further SP logic processing, since the η unit value is not constant relative to angular value of that coordinate (known as θ). Ideally, to compensate for that would require the extrapolation windows for η EUs to depend on the absolute value of η ; in other words, the closer the track is to the beam axis, the wider extrapolation windows should be used. This compensation cannot be implemented in the current SP design because of insufficient logic size, so some average η extrapolation windows are selected that allow for track reconstruction of sufficient quality.

For the SLHC SP design, we intend to convert the wiregroup to θ directly. This would allow for uniform extrapolation windows with no dependence on θ .

At the end of the pipelined logic, when the best 3 tracks are identified, the SP will still assign ϕ and η values to them as required along with any alignment corrections of the chamber positions for the best accuracy. However, this assignment for just 3 tracks consumes a very small amount of logic resources.

2) Half-strip to ϕ conversion

The track-finding algorithm can operate using a ϕ coordinate limited in precision to one strip in ME1/2, ME2/2, ME3/2, and ME4/2 chambers (0.1333°). This significantly reduces logic resources without compromising the performance.

The half-strip coordinate is first multiplied by a certain factor. For most chambers this factor is $\frac{1}{2}$, which is equivalent to removing the least significant bit (LSB). For some chambers, this factor is exactly 1 (no operation). Finally, for a relatively small number of chambers, this factor is a certain “inconvenient” number, so an internal FPGA multiplier or LUT has to be used. The list of chamber types and corresponding factors is shown in Table 1.

Table 1: Multiplication factors for ϕ conversion.³

Chamber type	Strip angle	F
ME1/2, ME2/2, ME3/2, ME4/2	0.1333°	$\frac{1}{2}$ (remove LSB)
ME2/1, ME3/1, ME4/1	0.2666°	1 (no operation)
ME1/1a	0.2222°	0.8335
ME1/1b	0.1695°	0.636
ME1/3	0.1233°	0.4625

When the best 3 tracks are identified, the Track Finder will still need to assign the precise ϕ values to them. However, the conversion to full-precision ϕ has to be done for only 3 trigger primitives, which leads to logic size reduction.

3) Wiregroup to θ conversion

For the majority of chambers, this conversion can be done by a small LUT. It takes the wiregroup number as input, and provides a 7-bit θ value on the output.

The exception is ME1/1 chambers, because of their unique tilted-wire design [2]. The SP may receive two half-strip numbers and two wiregroup numbers on each BX from such chambers, and it is impossible to match each of these half-strip numbers to one particular wiregroup number, so all combinations have to be taken into account. This requires each wiregroup parameter to be converted into two distinct θ outputs, or “duplicated”. Figure 1 shows a graphical representation of the problem.




Figure 1: θ duplication in ME1/1 chambers

The current SP design does not implement this logic. To allow for using ME1/1 trigger primitives in the SP track

³ This table shows strip angle for each chamber type. Half-strip angle can be calculated by dividing strip angle by 2.

reconstruction, η extrapolation windows are made wide enough to be insensitive to ME1/1 wire tilt. This should work fine for LHC, but with increased SLHC background tighter extrapolation windows may become necessary.

The proposed wiregroup to θ conversion schematics is shown in Figure 2.




Figure 2: ME1/1 wiregroup to θ conversion

The 6-bit wiregroup number is converted into a base θ value by an LUT. Simultaneously, two other LUTs that take strip numbers and 2 most significant bits of wiregroup as inputs produce 4-bit correction values, which are added to the base θ value and form the duplicated θ outputs.

D. Geometry constraints for track building

In the current SP design, we have to consider almost all combinations of trigger primitives since each of them may come from any chamber in the station. In the proposed upgraded design, since we receive all primitives from all chambers without filtering, it is possible to implement logic only for the physically allowed chamber combinations.

There are two considerations that must be taken into account:

- Track bending in magnetic field is limited. The ϕ difference between primitives created by a single track in any two stations cannot be more than $\sim 10^\circ$.
- Track projection in θ direction is a straight line; bending happens only in ϕ projection. Therefore, a chamber coverage map in θ must be used to select valid chamber combinations.

Figure 3 shows such map. As an example, one can clearly see that extrapolations between chambers ME1/2 and ME3/1 are not necessary because any single track originating in Interaction Point (IP) cannot cross both of these chamber types. There are many other chamber type combinations that don't have to be considered. Note that for halo tracks, the chamber combinations would be different.

Using the above constraints, the track building maps were generated. Examples of such maps for collision tracks are shown in Figure 4.

E. Upgraded design – implementation of modules

1) Extrapolation Units

Since the CSC is not a pixel-type detector, when two trigger primitives are available from a certain chamber it is impossible to tell which half-strip coordinate corresponds to which wiregroup. This leads to additional complexities in the design of the track-finder because all combinations of half-strip and wiregroup coordinates should be analyzed. The current SP design takes this into account only for ME1 trigger primitives; ME2, ME3, and ME4 trigger primitives are assumed to have perfect match between half-strip and wiregroup coordinates, which is a trade-off. In the upgraded design, we must take this into account for all stations.




Figure 3: θ coverage map.




Figure 4: Track building maps for ME1 \rightarrow ME2 and ME1 \rightarrow ME3 extrapolations and track assembly.

Even with the geometry constraints shown above, the number of extrapolation units in the upgraded design will grow significantly. As can be seen from Table 2, the total number of required extrapolations is 2252, which is ~ 11 times more than in the current design.

2) Final Selection Unit

Such a large number of track candidates (54 collision and 54 halo) leads to a huge growth in final selection and $\phi + \theta$ cancellation logic. Since we need to keep the latency as low as possible, the implementation of selection and cancellation logic is very straightforward – each candidate has to be

compared with each other simultaneously. The number of such comparisons is proportional to the square of the number of candidates. This means that the logic size for FSU will grow relative to the current design by a factor of ~ 20 . Taking into account that FSU is already occupying the largest part of logic in the current design, it may become problematic to select a suitable FPGA for such upgraded design. The present SP board is using and FPGA from Xilinx’s Virtex-5 family (XC5VLX155). The largest FPGA that should be soon available is XC6VLX760, is just 5 times bigger.

Extrapolation	ϕ EU	θ EU
ME1-ME2	208	248
ME1-ME3	232	336
ME1-ME4	168	272
ME2-ME3	132	132
ME2-ME4	132 </td <td>132</td>	132
ME3-ME4	132	132
ME1-MB1	48	0
ME2-MB1	48	0
Total	1100	1252

Table 2: Numbers of ϕ and θ extrapolations⁴

F. Other modules

Implementation of other modules should not lead to any problems with FPGA capacity because the amount of logic they occupy is small relative to EUs and FSUs, and the logic size grows in direct proportion (not square) to the number of track candidates.

G. Pattern-based track reconstruction

Taking into account possible implementation problems of the upgraded SP logic based on our current design, we have decided to evaluate an approach that can lead to significant logic size savings while providing all the functionality that is required for SLHC operation. It is very similar to pattern search logic used in front-end boards, such as the ALCT. In the case of the Sector processor, the pattern is created from the trigger primitives in chambers, so 4 “layers” of chambers are considered by pattern detectors. Besides logic size reduction, other benefits of this approach include:

- “Natural” ability to analyze multiple bunch-crossings.
- Virtually ghost-free track candidates, which improves the quality of final tracks reported to Global Trigger, reduces the size of selection logic and eliminates cancellation logic.
- Track timing is automatically set by the second trigger primitive. In the current SP design, we had to implement a special module and increase the latency to achieve that.

⁴ Does not include halo extrapolations. For ME1 extrapolations, there are more θ EUs than ϕ EUs because of ME1/1 θ duplication.

There are separate pattern detectors for ϕ and θ projections. The sector is split into 5 ϕ zones and 6 θ zones defined by the chamber coverage map; each zone has its own independent pattern detector.

The preliminary structure of the pattern used for ϕ zones is shown in Figure 5. Before ϕ pattern detectors can be applied, trigger primitives from each chamber are decoded as described in section II.C.2). Then, “raw hits” are recreated inside the FPGA logic. Each dot on the diagram represents a certain number of raw hits ORed together; this way, sufficient ϕ coverage is achieved while keeping the logic size of the pattern detector relatively small. The number of ORed hits for each dot is shown above ME1 station. Such structure allows for precise detection of high- P_t tracks; low- P_t tracks are detected with much lower precision, which is acceptable.




Figure 5: Possible pattern structure for ϕ zones

H. CSC+Tracker = Better Trigger

One more important direction which is being investigated is the challenge of matching CSC triggers with an inner silicon Tracker. By doing this, we should be able to reach better rate reduction using the Tracker to confirm CSC trigger candidates, and improve track fitting.

III. CONCLUSIONS

We are moving ahead quickly with the hardware-independent design and simulation of the logic blocks for the upgraded Track-Finder. So far, importing all available trigger primitives seems possible. If some serious obstacles are encountered that would prevent us from doing that, we will consider returning to trigger primitive filtering in MPC (7 primitives per BX from each MPC).

Additionally, simulations are being developed for matching CSC and Tracker trigger primitives to achieve better trigger system performance.

IV. REFERENCES

1. “US CMS SLHC Muon Trigger R&D”, presentation by Darin Acosta (University of Florida) <http://indico.cern.ch/getFile.py/access?contribId=47&sessionId=5&resId=1&materialId=slides&confId=48781>
2. I.A.Golutvin et al, “ME1/1 cathode strip chambers for CMS experiment”, Physics of Particles and Nuclei Letters, Volume 6, Number 4 / July, 2009