Available on CMS information server **CMS NOTE 1999/060**

The Compact Muon Solenoid Experiment

# CMS Note

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

**December 8, 1999**

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

D. Acosta, A. Madorsky, B. Scurlock, S.M. Wang

*University of Florida, Gainesville, USA*

A. Atamanchuk, V. Golovtsov, B. Razmyslovich

*St. Petersburg Nuclear Physics Institute*

**Abstract**

The algorithms and hardware implementation of the Track-Finder for the CMS endcap muon system are described. The Track-Finder is implemented as 12 Sector Processors which identify up to the three best muons in $60^\circ$ azimuthal sectors. The coverage in pseudorapidity includes the region of overlap with the barrel muon system. The track-finding algorithms are inherently 3-dimensional to achieve maximum background rejection. A 3-station sagitta measurement assigns $p_{\mathrm{T}}$ to a precision similar to that achievable in the barrel region despite reduced magnetic bending in the endcap. The input to the Track-Finder is capable of collecting track segments from multiple bunch crossings, and the overall latency is expected to be about 350 ns.

# 1 Introduction

The endcap muon system of CMS [1] consists of four stations (ME1–ME4) of cathode-strip-chambers (CSC) on each end of the experiment (see Fig.1), although insufficient funds exist for the fourth station at present. A single station is composed of six layers of CSC chambers, where a single layer has cathode strips aligned radially and anode wires aligned in the orthogonal direction. The CSC chambers are trapezoidal in shape with a $10^\circ$ or $20^\circ$ angular extent in azimuth ($\phi$). The arrangement of the chambers into trigger sectors which are aligned from one station to the next is described in [2]. The CSC chambers are fast (60 ns drift-time) and participate in the Level-1 trigger of CMS.

The purpose of the CSC Track-Finding processor is to link trigger primitives (track segments) from individual muon stations into complete tracks (see Fig. 2), measure the transverse momentum $p_\mathrm{T}$ from the sagitta induced by the magnetic bending, and report the number and quality of tracks to the Level-1 Global Muon Trigger. This objective is complicated by the non-uniform magnetic field in the CMS endcap and by the high background rates; consequently, the present design incorporates full three-dimensional information into the track-finding and measurement procedures. In contrast, the Track-Finding processor for the barrel drift-tube (DT) muon system is intrinsically two-dimensional [3]. Issues related to the separation of the Track-Finding processors for the barrel and endcap muon systems are addressed in [4].

This Note is organized as follows: Sec. 2 describes the overall scheme for the Level-1 Trigger formation of the CSC system, and Sec. 3 lists the requirements of the CSC Track-Finder. Sec. 4 describes the precision and range of the track quantities needed by the Track-Finder, while Sec. 5 describes the format of the output. The trigger boundaries and data flow are discussed in Secs. 6 and 7, respectively. The logic algorithms of the Track-Finder are discussed in Sec. 8.

# 2 Level-1 Trigger Formation

A block diagram of the overall trigger scheme for the CSC Level-1 Trigger is shown in Fig. 3.

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 [5]. The anode data provide precise timing information as well as $\eta$ information, and the cathode data provide precise $\phi$ information. Comparators on the cathode cards localize a hit cluster to within a half strip for a given layer.

The Trigger Motherboard (TMB) 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 [6]. At this stage a track-segment vector exists. An ambiguity exists if there are two anode and cathode LCT candidates. One possible solution [7] to resolve the ghost combinations has been proposed using information from the Resistive Plate Chamber (RPC) muon detection system.

The purpose of the Muon Port Card (MPC) is to reduce the number of connections from the front-end trigger system to the counting house electronics. Each MPC collects track segments from the TMBs of nine chambers and selects up to the three best candidates out of the 18 possible [8]. For stations ME2–ME4, these nine chambers correspond to a $60^\circ$ sector in $\phi$. For station ME1, these chambers correspond to a $30^\circ$ sector in $\phi$ assuming the inner half of each

ME1/1 chamber ($|\eta| > 2.1$) does not participate in the trigger. The MPC then sends the three best track segments per sector via optical cable to the counting house to be processed by the CSC Track-Finding Processor.

The track segments from two MPCs are received, synchronized, reformatted and distributed in the counting house by a Sector Receiver (SR) card [9], which forms the input to the CSC Track-Finding processor. One SR card is needed to cover a $60^\circ$ sector in ME1, one for ME2 and ME3, and one for ME4. Along with reformatting the track segments into global coordinates suitable for the Track-Finding algorithms, the SR card also applies alignment corrections to the data to account for placement errors of the CSC chambers.

The signals from the SR are sent across a point-to-point backplane to the Sector Processor (SP), which implements the CSC Track-Finding algorithms and which is the focus of this document. It is this processor which achieves the objectives stated in the Introduction. The output of the SP is the momentum and quality of the three best muon candidates found in a $60^\circ$ sector in the endcap muon system.

The final stage of the Level-1 CSC Trigger is the Muon Sorter (MS), which collects the output of all SPs and selects the four highest-rank muons from entire CMS endcap muon system. For the SP design presented here, the MS accepts 36 muons from 12 sectors.

The output of the MS is sent to the Global Muon Trigger (GMT), which associates barrel and endcap muon candidates with RPC muon candidates and with quiet regions in the calorimeter [10]. The best 4 muons of CMS are then sent to the Global Level-1 Trigger, where specific $p_{\rm T}$ thresholds are applied to the muon candidates.

# 3 Requirements

The physics goal of the Track-Finder is to efficiently identify muons with as low a threshold in $p_{\rm T}$ as possible in order to meet the rate requirements of the Level-1 Trigger of CMS. The maximum Level-1 rate is designed to be 30 kHz, with muon triggers accounting for about half of it. The single muon trigger should account for half of that again, giving a rate per unit rapidity ($\eta$) of order 1 kHz.

The $p_{\rm T}$ resolution of the Track-Finder should be 30% or better to satisfy the rate requirement of 1 kHz per unit rapidity at the LHC design luminosity of $10^{34}$ cm$^{-2}$s$^{-1}$, as illustrated in Fig. 4. However, a larger safety factor in the trigger rate would be achieved if the $p_{\rm T}$ resolution could be improved to about 20%. In any case, such resolution in the non-uniform magnetic field of the endcap is possible only if both $\phi$ and $\eta$ coordinates of the track segments are used in the $p_{\rm T}$ measurement. Track segments in two—and probably three—muon stations will be needed.

The Track-Finder should analyze track segments from most of the geometrical coverage of the CSC system, $0.9 < |\eta| < 2.4$, though it is not clear that the region $|\eta| > 2.1$ (the bottom half of ME1/1) will participate in the trigger. Care must be taken in the overlap region between the barrel and endcap muon systems ($0.9 < |\eta| < 1.2$) to avoid a duplication of triggers which would spoil a multi-muon trigger. An agreement has been reached to define a sharp boundary in $\eta$ between the barrel and endcap Track-Finders [4]. Efficient coverage of the overlap region, though, dictates that signals should be shared between the two muon systems. Track segments from MB1 and MB2 in the outer barrel wheels will be sent to the endcap Track-Finder, and track segments from ME1/3 and ME2 will be sent to the barrel Track-Finder.

The probability to associate anode and cathode LCTs with the correct beam crossing is not expected to be better than about 92%. Moreover, the strong fringe fields in overlap region between CSC and DT muon chambers is expected to vary the drift time across the DT chambers in that region, reducing the probability to associate a DT trigger primitive with the correct beam crossing. Therefore, it is desirable that the Track-Finder have some ability to analyze track segments from several bunch crossings for maximum efficiency.

All data in the Track-Finding processor must be pipelined at the LHC beam crossing frequency of 40 MHz. The allocated latency for the Sector Processor is only about 16 crossings, or 400 ns. Finally, the trigger algorithms of the Track-Finder should be programmable so that the experimenter may vary the algorithms depending on the running conditions.

# 4 Input Data and Precision

The size and complexity of the logic for the Track-Finder depends strongly on the number of signals that must be correlated. In this section, we document the bit count necessary for the track-finding algorithms and for the $p_{\mathrm{T}}$ assignment. The bit count for all data exchanged between the CSC trigger cards is documented in [11]. Bit counts for CSC *and* DT interfaces, including range and precision, is documented in [12].

The most precise measurement made by the CSC chambers is $\phi$, since the bending of a muon in the plane orthogonal to the solenoid axis is used to determine $p_{\mathrm{T}}$. The cathode LCT trigger of the CSC system is expected to deliver a resolution of $\sim 0.1$ strip for the $\phi$ measurement, which corresponds to $\sim 0.23$ mrad ($\sim 0.46$ mrad) for $10^{\circ}$ ($20^{\circ}$) chambers. A segmentation of 0.26 mrad can be achieved if we assign 12 bits of data for a $60^{\circ}$ sector, which will contribute less than 5% to the overall $\phi$ resolution. However, it is the difference in $\phi$ measured by chambers in two different stations which is used to determine $p_{\mathrm{T}}$. The uncertainty in $\Delta\phi$ is dominated at low momentum by the multiple scattering of the muon as it penetrates the endcap region. Even if 10 bits are used for $\phi$ (a binning which corresponds approximately to one half-strip), the position resolution is not significant for $p_{\mathrm{T}} < 50$ GeV. Moreover, the range in $\Delta\phi$ does not exceed $7.5^{\circ}$ for a muon which punches through to the endcap muon system; thus, the 3 most significant bits of $\Delta\phi$ can be dropped in the Track-Finder logic.

Related to $\phi$ is the bend angle $\phi_{\mathrm{b}}$ (sometimes denoted $\Psi$) in that plane. It is the direction angle of the track segment with respect to the radial direction, in the plane orthogonal to the beam axis. The measurement of $\phi_{\mathrm{b}}$ is useful in the extrapolation of one track segment to another in a different muon station. The smallest non-zero angle measured is given by a cathode LCT pattern that extends across two half-strips. The relation between this minimum angle, $\phi_{\mathrm{b}}^{\mathrm{min}}$, and the angle subtended by a radial strip, $d\phi$, is given by:

$$\tan \phi_{\mathrm{b}}^{\mathrm{min}} = \frac{z}{h} \tan \frac{d\phi}{2} \tag{1}$$

where $h$ is the chamber thickness (12.5 cm), and $z$ is the placement of the chamber along the $z$ axis (replacing $z$ with $r$ gives the bend angle in the plane parallel to the beam axis). The smallest strip pitch is 2.3 mrad, so $\phi_{\mathrm{b}}^{\mathrm{min}}$ corresponds to about 60 mrad. This bend angle is independent of $\phi$ for a given LCT pattern since the cathode strips are radial; thus, $\phi_{\mathrm{b}}$ basically labels the

LCT pattern. The maximum extent of $\phi_{\rm b}$ corresponds to about 8 strips for muons at the lowest possible $p_{\rm T}$, so 5 bits (including a sign bit) are needed to label $\phi_{\rm b}$.

The limited precision of $\phi_{\rm b}$ implies that the full precision of $\phi$ is not necessary for the extrapolation portion of the Track-Finder. The station-to-station separation in the endcap is at least nine times the width of a CSC chamber, so the minimum bend angle of a half-strip corresponds approximately to a 4.5 strip change in $\phi$, which is 10 mrad. Therefore, we can drop the 5 least significant bits from the full $\phi$ precision and maintain 8.2 mrad precision. When comparing $\phi_{\rm b}$ with $\Delta\phi$ for the track extrapolation, only 5 bits are needed for both quantities.

Three factors determine the necessary precision on $\eta$: the $p_{\rm T}$ measurement, the alignment corrections to the chamber positions, and the separation of beam halo muons from primary interaction muons. Of these, the last constraint is the tightest. If 6 bits are used for $\eta$ (64 bins from 0.9 to 2.4), then it is possible to distinguish a particle traveling parallel to the beam axis from one originating from the interaction point by comparing the difference in $\eta$ between track segments in any two stations. For example, the hardest case is distinguishing a through-going muon at $\eta = 2.4$ between ME3 and ME4. The $\eta$ difference is 0.09, which corresponds to a difference of 4 bins. Although finer resolution could improve the rejection power, it would significantly reduce the efficiency for primary muons. The spread in $\eta$ due to multiple scattering is approximately 1 bin for $p_{\rm T} = 10$ GeV.

The necessary precision in $\eta$ for the $p_{\rm T}$ assignment is set by the variation of the magnetic bending for a particular region in $\eta$. The relative uncertainty on the magnetic bending in this region should contribute negligibly to the overall uncertainty of the transverse momentum determination. If we aim for a Track-Finder with the capability to measure $p_{\rm T}$ to 20% accuracy, then the uncertainty on the magnetic bending in a particular region should be less than 10%. If an $\eta$ segmentation of 0.1 units is chosen, then the variation in the highest $\eta$ bin is $\pm 7.5\%$ ($\pm 6\%$ for $2.0 < \eta < 2.1$). Thus, 4 bits are necessary to represent $\eta$ in the entire CSC muon system from $\eta = 0.9$ to $\eta = 2.4$ to determine $p_{\rm T}$ with sufficient accuracy in the endcap.

For completeness, we note that the precision of $\eta$ necessary to perform any corrections on the misplacement of CSC chambers with respect to their nominal position and orientation is 5 bits. These corrections are applied in the Sector Receiver, and allow for a tolerance of 1.6 cm in the placement of a chamber with no degradation in the accuracy of $\phi$.

In short, significantly less than the full 11-bit precision of $\eta$ available from the anode LCT cards is needed.

The bend angle along the $\eta$ view is not anticipated to be used by the Track-Finder. Instead, an "accelerator muon" flag will be set if the anode LCT pattern corresponds to a track segment that is parallel to the beam axis. This 1-bit flag, along with the $\eta$ coordinates from different track segments, will be used to distinguish beam halo muons from primary interaction muons.

In addition to the coordinates of each track segment, a 3-bit quality word will be sent to the Track-Finder. This word may encapsulate, among other things, the quality of the timing match between the anode and cathode LCT patterns and the number of layers missing from the LCT patterns. A quality of "000" will denote that no track segment exists (a hit flag).

In total, 27 bits are needed to describe each track segment, as tabulated in Tab. 1. In addition to those bits, each Sector Receiver is expected to send a 1-bit Bunch Crossing 0 (BXN0) identifier and a synchronization error bit. It is *not* anticipated that the Track-Finder will utilize a chamber identification number (6 bits), which could be used to untangle ghost combinations of LCTs in the same chamber. The complication, expense, and extra latency do not seem to warrant the effort unless further study shows it to be necessary. Moreover, there is a proposal [7] to solve the problem at the chamber level using RPC information.

The format of the data delivered by the barrel system (see [12]) is similar to that of the endcap system, but lacks $\eta$ information and has a wider bit field for $\phi$ and $\phi_{\rm b}$. We propose to take only the most significant bits to match the format of the CSC data, so that each DT segment is represented by the 20 bits listed in Tab. 2. The biggest reduction oocurs for $\phi_{\rm b}$, which has 10-bit precision. However, the CSC Track-Finder does not use $\phi_{\rm b}$ for the assignment of $p_{\rm T}$, so the additional precision is not necessary. In addition to the track quantities, the barrel system should send a 4-bit Bunch Crossing Number (BXN) and a 1-bit synchronization error bit (set if any errors occur at or before the synchronization of the barrel data).

# 5 Output Data

The Track-Finder logic should result in a list of muon candidates with several track quantities measured: $p_{\rm T}$, $\phi$, $\eta$, sign, and quality. The ultimate format of the data when it is received by the Global Level-1 Trigger is specified in [12]. However, for sorting purposes, it is more convenient to combine the 5-bit $p_{\rm T}$ word and the 2-bit track quality into a 7-bit rank. This rank is then used for sorting muons in the CSC Muon Sorter and in the Global Level-1 Muon Trigger. It can be decoded by the Global Level-1 Trigger back into $p_{\rm T}$ and quality. In addition, only a local $\phi$ and a local $\eta$ need be sent from the Sector Processor to the CSC Muon Sorter. By virtue of the cable connection, the sign of $\eta$ and the sector number can be inferred, which save 4 bits per muon. The output of the Sector Processor, therefore, is listed Tab. 3 with the precision required by the Level-1 Trigger.

In addition to the track quantities, each Sector Processor should send a 4-bit BXN and a 1-bit synchronization error bit (set if any errors occur at or before the synchronization of the Sector Receiver data).

# 6 Trigger Boundaries

The Track-Finder is a regional trigger. Each Sector Processor operates on data in an azimuthal sector and in a certain pseudorapidity interval. The results are collected globally by the CSC Muon Sorter.

The Muon Port Cards of ME2, ME3, and ME4 define the azimuthal sector in the endcap muon system to be $60^\circ$ by collecting information from six $10^\circ$ chambers and three $20^\circ$ chambers. The CSC chambers are projective in $\phi$, so we are justified in ignoring muons which cross boundaries because the maximum deflection and scattering between the stations is less than $7.5^\circ$ for muons with $p_{\rm T} \sim 3$ GeV. The Muon Port Cards of ME1 actually define $30^\circ$ sub-sectors, so two ME1 sub-sectors are combined to match the $60^\circ$ sectors of the other stations.

The CSC track segments in each sector are collected in the pseudorapidity interval $0.9 < |\eta| < 2.4$. The region of overlap with the barrel muon system occurs in the interval $0.9 < |\eta| < 1.2$ (see Fig. 1). In this region, the $60^\circ$ endcap sectors must match the $30^\circ$ (and non-projective) barrel sectors. The optimum choice of alignment starts the first $60^\circ$ sector at $\phi = 15^\circ$, as discussed in [2] and shown in Fig. 5. Since each endcap must match onto the barrel segmentation, this choice necessarily implies that each endcap must be mirror symmetrical. The gaps between the barrel muon chambers and the imperfect alignment in azimuth imply 14% dead area for both MB1 and MB2 in each sector. We also note that the coverage of ME1/3 is only 75% of azimuth since the chambers do not overlap.

# 7 Data Flow

The number of CSC track segments collected by one Sector Processor is 15, assuming that ME4 participates. Six track segments are delivered from ME1; three each are delivered from ME2–ME4. It is planned that the track segments from the barrel chambers in the overlap region will be processed by the same Sector Processor, so 8 additional track segments are delivered ($2 + 2$ in both MB1 and MB2). The format of the data delivered by each system is listed in Tabs. 1 and 2. Each CSC segment is represented by 27 bits, and each DT segment by 20 bits.

The CSC track segments are delivered to the Sector Processor by a custom backplane from three Sector Receivers (one for ME1, one for ME2 and ME3, and one for ME4). Thus, each SR delivers 162 bits (81 bits for ME4) representing six track segments (three for ME4), not including the BXN0 and error bit. Each card has a 9U VME form factor and contains a VME interface. The custom backplane is actually a point-to-point connection operating at 280 MHz. LVDS Channel Link transmitters and receivers from National Semiconductor are used for communication, which serialize/de-serialize 28-bits to/from 4 differential signals and a clock. Thus, there is about a factor 3 compression in the number of signal lines on the backplane. This is necessary so that all the signals can fit onto connectors at the back of the card.

The DT track segments are expected to be delivered to the Sector Processor via a transition board on the back of the crate with cable connections to the appropriate fan-out units of the barrel trigger system. These signals also are expected to be brought to a connector on the Sector Processor using Channel Link transmitters. A total of 160 bits are sent for 8 track segments, not including the BXN and error bits.

A schematic view of the backplane connections for one sector is shown in Fig. 6. It is anticipated that one crate can handle two such sectors. Connections to and from the barrel trigger system occur through transition boards at the back of the crate. A Clock and Control Board (CCB) distributes the LHC clock and control signals to the Sector Processors and Sector Receivers in the crate.

# 8 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. 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 then 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 which 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.

The overall scheme for the CSC Track-Finder is illustrated in Fig. 7. Each of the important blocks is described in detail below.

## 8.1 Data Input

The input data to the Sector Processor is de-serialized by Channel Link receivers and synchronized with the local clock before being sent into the Extrapolation Units. However, the Sector Processor should have some ability to analyze track segments received in different bunch crossings because of the intrinsic timing spread of the trigger primitive formation and because of the usefulness of such a feature when commissioning the trigger system.

To incorporate a multi-bunch mode, we take advantage of the sparseness of the data. If the data is not sparse, track segments would be lost already at the Muon Port Card, which selects only the three best track segments from 9 chambers. For example, the Poisson average must be less than 1.5 track segments for there to be less than a 5% probability of losing a track segment in the Muon Port Card. Therefore, we consider track segments from other bunch crossings 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. We plan to design a window which is open for two bunch crossings, which should be sufficient for the CSC trigger system if the timing of the trigger primitives is tuned so that there is little probability of receiving a muon before the correct bunch crossing. Although the window is left open for more than one bunch crossing, the Sector Processor must report triggers at the correct bunch crossing every crossing. In other words, overlapping time buckets are used.

This capability is introduced before the track segments are stored in a FIFO (for later retrieval by the Assignment Unit) and before the extrapolation logic. For a given station, the best three track segments (the best six for ME1) are selected from 2 crossings based on the track segment quality and on the deviation from the current crossing. The same can be done for the best four track segments from MB1 or MB2. In the simplest scenario, the track segments in crossing $N$ have highest priority, followed by those in $N + 1$. To keep the sorting logic compact and fast, it is proposed that the Muon Port Card send the best three track segments in ranked order.

This scheme is shown in Fig. 8. The order of the track segments into the rest of the Sector Processor can be changed; but as this occurs before storage in the local FIFO, it does not influence the rest of the logic. The latency of the scheme is 2 bunch crossings.

A 1-bit flag is set to record whether a track segment comes from the current bunch crossing or the following one. This flag will be used in the Final Selection Unit of the Sector Processor to determine if a trigger should be inhibited. Currently, nothing prohibits a track segment from contributing in multiple bunch crossings, and it is possible that a single muon can result in more than one trigger in subsequent bunch crossings. This is illustrated in Fig. 9 for a muon which has track segments in two bunch crossings. The first occurrence of the trigger is the correct one. On the other hand, a trigger will be reported only once at the correct bunch crossing if no more than one track segment is out of time, as illustrated in Fig. 10. This is accomplished with special logic in the Final Selection Unit which inhibits reporting a trigger at an earlier crossing. This logic also can inhibit the double trigger of the first example as well.

## 8.2 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:

1. Determine if each track segment is in the allowed trigger region in $\eta$
2. Compare the $\eta$ values of the two segments to determine if both lie along a straight line projection to the collision vertex
3. Compute the difference in $\phi$ between the two track segments
4. Check if that difference is consistent with the bend angles $\phi_{\mathrm{b}}$ measured at each station
5. Compare the difference in $\phi$ to the maximum allowed at that $\eta$
6. Correlate the quality bits between the two track segments
7. Check that the Accelerator Muon bit is not set for at least one of the track segments
8. Assign an overall quality to the extrapolation

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 $p_{\mathrm{T}}$ (low, medium, and high) to extrapolations involving the first muon station (ME1 or MB1). Otherwise, the quality just represents whether the extrapolation was successful or not. The expected $p_{\mathrm{T}}$ resolution for a $\phi$ resolution of 10 bits is about 30% when ME1 is involved. The quality word is used later when muon candidates are sorted.

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 ME1 to ME4 in order to reduce the number of combinations to 63. This prohibits triggers involving only hits in ME1 and ME4, 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 it is expected to be implemented in field-programmable gate-arrays (FPGAs). A logic diagram for the extrapolation of one track segment pair in the endcap region involving ME1 is shown in Fig. 11. That for an extrapolation involving track segment pairs in the outer stations of the endcap is shown in Fig. 12. The latency to accomplish this decision logic is expected to be 3 beam crossings.

Extrapolations in the region of overlap between the two muon systems involve track segments from both DT and CSC chambers. In particular, MB2/1, MB2/2, ME1/3, and ME2/2 participate.

In this case, the number of combinations is different from the endcap region. Each DT chamber reports up to two muons for a chamber which extends $30^\circ$ in azimuth; thus, 4 track segments each from MB1 and MB2 participate in a $60^\circ$ trigger sector. Moreover, no $\eta$ information will be used from the DT chambers, so the extrapolation in $\eta$ will be removed from the logic (Item 2). However, the CSC Track-Finder will require at least one track segment to come from a CSC chamber in the overlap region, so the $\eta$ coordinate of that track segment can be used for the other tests listed above. A logic diagram for the extrapolation of one track segment pair in the overlap region involving MB1 is shown in Fig. 13. That for an extrapolation in the overlap region without MB1 is shown in Fig. 14. The latency to accomplish this decision logic also is expected to be 3 beam crossings.

The number of extrapolations to perform in the overlap region depends on which stations are required to participate in the track assembly stage, discussed below. We propose to require a track segment in ME2, which implies that the number of additional extrapolations to perform is 24.

## 8.3 Track Assembly

### 8.3.1 Endcap Region

The track assembly stage examines the output of the extrapolation units 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 identification of the participating track segments is registered also.

The underlying feature of a Track Assembly Unit is the concept of a "key station." For this Track-Finder design, ME2 and ME3 are key stations. A valid trigger in the endcap region must have a hit in one of those two stations. In this way, the output of the extrapolation units in the endcap region can be separated into two data streams: one for patterns keying off ME2, and one for patterns keying off ME3. This is illustrated in Fig. 15. Some muons will be found by both streams, so the Final Selection Unit described in the next section must resolve the double counting.

Each track segment of a key station, of which there are three each for ME2 and ME3, is tested for extrapolations to the other stations. Therefore, the extrapolation results appropriate for that key segment are interrogated. The Track Assembler logic checks if the key track segment has successful extrapolations to more than one station. (If $2 \to 1$ is successful, is $2 \to 3$ also? If so, the track is at least a "1–2–3" type.) The output of this logic is a code designating the best track pattern which contains the given key segment. Thus, up to three tracks may be found per data stream, six total for the two endcap streams.

There are six track segments allowed in ME1, and the extrapolation quality to ME1 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, a total of 18 bits are interrogated. Since the number of input bits is small, each of these "Link" units can be implemented as a static RAM look-up memory, as shown in Fig.16. The latency, therefore, is just one beam crossing. The output code is a 9-bit word labelling the track segments used in each station (*e.g.* 3 bits for ME1, 2 bits each for ME2–ME4), and a 6-bit quality word giving the type and rank of the assembled track. A tabulation of the possible track categories for the assembled tracks of the two endcap data streams and the one overlap stream (described below) is given in Tab. 4. Also shown is a preliminary ranking used for sorting purposes in both the Track Assembler and in the Final Selection Unit.

### 8.3.2 Overlap Region

The inclusion of the overlap region creates additional data streams in the track assembly stage. It is reasonable to propose that the CSC Track-Finder should require at least one track segment in the CSC system (otherwise it should be found by the DT Track-Finder). In that case, ME1 and ME2 are natural choices for the key stations. However, ME1 has six track segments in a $60^\circ$ sector, and this would require six Link chips and would create a larger number of muons to sort. Additionally, the ME1/3 chambers are non-overlapping, and have only 75% azimuthal coverage. Hence, ME1 is not an ideal key station.

Using ME2 as a key station in the overlap region is quite acceptable, however. The only drawback is that the $\eta$ coverage stops at $|\eta| = 1.0$, which corresponds to two-thirds of the overlap region. If ME2 is the only key station, then there is little or no redundancy and the efficiency will be somewhat reduced. All muons in this region would be required to have a track segment in ME2. This would prohibit triggers with a one track segment in MB1 or MB2 and one in ME1/3. However, the ME2 chambers overlap in $\phi$, and the LCT efficiency is expected to be greater than 95%.

The proposed solution to the overlapping coverage of the barrel and endcap Track-Finders is to divide the region of overlap in half at $\eta = 1.05$. This would allow the barrel Track-Finder to trigger on track segments in MB1 and MB2 without knowledge of $\eta$. It would allow the endcap Track-Finder to use only ME2 as a key station (provided that simulations show that the efficiency is adequate). In this case, we find it feasible to incorporate both the endcap and overlap region logic onto the same Sector Processor (extrapolation units and track assembler units). One data stream keying off ME2 is added to the two endcap-only data streams.

## 8.4 Final Selection

The final selection logic combines the information from the Track Assembler streams, 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 should be canceled (presumably the lower rank combination, if the two tracks are not completely identical). Thus, the Final Section Unit is a sorter with cancellation logic.

A block diagram of the Final Selection Unit is shown in Fig. 17. The sorter part of the logic compares the qualities of all pairwise combinations of tracks from the Track Assembler 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 should be programmable.

The two comparison steps are done in parallel in one beam crossing. The next step of the logic, the Final Decision Unit, examines the results of all these comparisons and reports the identities of the three best and distinct muons. It also takes one beam crossing. Finally, the track segment information of the selected muons is taken from a multiplexer and transmitted in the next beam crossing 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 before going to the Assignment Unit. This inhibits one class of double triggers mentioned in Sec. 8.1.

The Final Selection Units sorts and cancels 9 tracks down to 3 since there are two endcap data stream streams and one overlap data stream.

## 8.5 Measurement

The Sector Processor measures the momentum of the identified muons in the final stage of processing. This includes the $\phi$ and $\eta$ coordinates of the muon, the magnitude of the transverse momentum $p_{\mathrm{T}}$, the sign of the muon, and an overall quality which we interpret as the uncertainty of the momentum measurement. The format of the data is specified in Tab. 3. In particular, $p_{\mathrm{T}}$ and the track quality are combined into an overall rank before transmission to the CSC Muon Sorter. Reference [12] specifies that the coordinates are to be reported at the second station, since this is convenient for later association with RPC trigger data in the Global Muon Trigger. This is also convenient for the Track-Finder because the muon track parameters do not need to be extrapolated back to the interaction point, which would be prone to errors.

The first step in the assignment of these quantities uses the track segment information delivered by the Final Selection Unit as input to a multiplexer to select the appropriate data from a FIFO, as shown in Fig. 18. Also, those output quantities which do not require high accuracy and which have small bit fields are determined at this step (which is basically everything except for the transverse momentum measurement).

The most important quantity to calculate accurately is the muon $p_{\mathrm{T}}$, as this quantity has a direct impact on the trigger rate and on the efficiency (see Fig. 4). Simulations have shown that the accuracy of the momentum measurement in the endcap using the displacement in $\phi$ measured between two stations is about 30% at low momenta, when ME1 is included. 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 $\phi$ information from three stations when it is available. This should improve the resolution to at least 20% at low momenta, which is sufficient. Basically, we take advantage of the large multiple scattering for low $p_{\mathrm{T}}$ muons. Although the scattering has some probability to 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 $p_{\mathrm{T}}$ 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 which uses the minimum number of bits necessary in the calculation, as shown in Figs. 19 and 20. The first step is to do some preprocessing in FPGA logic: the difference in $\phi$ is calculated between the first two track segments of the muon, and between the second and third track segments when they exist. The first subtraction is done using the full 12-bit precision of $\phi$, which corresponds to about 0.1 strips (0.23 mrad). The second subtraction is done using only

8-bit resolution, which corresponds to two strips. It does not need the same accuracy because we are only trying to untangle the multiple scattering effect at low momenta.

The $12 + 1$ bit precision of the first subtraction is sent to the look-up memory if a muon is composed of only two track segments. If a muon has at least three track segments, then the 2 least significant bits of the first subtraction are dropped (10-bit precision of $\phi$ is sufficient for $p_{\mathrm{t}} < 50$ GeV). In addition, the 3 most significant bits are tested to see if the difference is smaller than the value when one of those bits is set (*i.e*, $\Delta\phi < 7.5°$), and then those 3 bits are dropped also. A physical muon should always satisfy the requirement. A similar thing is done for the second subtraction, except that the 4 most significant bits are dropped (*i.e*, $\Delta\phi < 3.75°$) because the magnetic bending is smaller between chambers in ME2–ME4. The end result is a 13-bit "$\Delta\phi$" word whose meaning depends on the track type.

The subtraction results are combined with the $\eta$ coordinate of the track, and the track type, then sent into a memory for assignment of the signed $p_{\mathrm{T}}$, as shown in Fig. 20. The number of $\eta$ bits necessary is 4, and 4 additional bits are necessary to describe the track category. Thus, a 2 M$\times$8 bit memory is needed for each of the three muons.

# 9 Implementation

A preliminary sketch of the functional layout of the Sector Processor on a 9U VME board is shown in Fig. 21. Processing of track segments from the endcap region as well as the region of overlap between the DT and CSC systems occurs on the same board. A transition module brings the DT signals from the rear of the crate to the custom backplane.

A preliminary board layout showing the relative sizes of the various components is shown in Fig. 22. It appears feasible to fit all components on one 9U VME board, and to fit all input connectors on the backplane.

Large FPGAs from the Xilinx Virtex series [13] are proposed for most of the Sector Processor logic. These chips come in ball-grid-array (BGA) packages to handle the large I/O count. Static RAM memories are used for the large look-up tables of the Track Assembler Units and the $p_{\mathrm{T}}$ Assignment Units. Data is brought into and out of the board via Channel Link LVDS transmitters and receivers.

# References


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

[2] D.Acosta, S.M.Wang, G.Wrochna, "Muon Trigger Sectors and CSC Layout," CMS-IN 1999/022

[3] G.M. Dallavalle *et al.*, "A Simplified Track Assembler I/O for the Muon Trigger Track Finder," CMS Note 1998/042;
A.Kluge and T.Wildschek, "The Hardware Muon Trigger Track Finder Processor in CMS—Specification and Method, Architecture and Algorithm, Prototype and Final Implementation," CMS Notes 1997/091, 1997/092, 1997/093

[4] G.M. Dallavalle *et al.*, "Issues Related to the Separation of the Barrel and Endcap Muon Trigger Track-Finders," CMS-IN 1999/015

[5] J.Hauser, "Baseline Design for the CSC-based Endcap Muon Trigger," CMS TN/95-013

[6] M.Matveev, B.P.Padley, "CSC Trigger Motherboard Specification," `http://bonner-ntserver.rice.edu/motherboard/tmb.htm`

[7] G.Wrochna *et al.*, "CMS Endcap Muon RPC–CSC Trigger Connection," CMS-IN 1998/030

[8] N.Adams, M.Matveev, B.P.Padley, "CSC Trigger Muon Port Card Specification," `http://bonner-ntserver.rice.edu/motherboard/pcspec1.htm`

[9] S.Chandramouly, J.Hauser, "Sector Receiver Design Document," `ftp://hepsun0.physics.ucla.edu/pub/cms/sr/sr_doc.ps`

[10] N.Neumeister, P.Porth, and H.Rohringer, "Simulation of the Global Muon Trigger," CMS-IN 1997/023

[11] "CMS Endcap Muon Trigger System Inter-Module Signals," `ftp://hepsun0.physics.ucla.edu/pub/cms/trigger/triggerbits.ps`

[12] G.Wrochna, "Muon Trigger Objects," `http://cmsdoc.cern.ch/~wrochna/tmp/trobj.ps`

[13] Xilinx Inc., `http://www.xilinx.com/products/virtex.htm`

Table 1: Information delivered from Sector Receiver to Sector Processor.

| variable | function | bits / muon | bits / 6 muons |
|---|---|---|---|
| $\phi$ | Phi Coordinate | 12 | 72 |
| $\phi_b$ | Local Bend Angle | 5 | 30 |
| $\eta$ | Eta Coordinate | 6 | 36 |
| Accelerator Muon | Eta Bend Angle | 1 | 6 |
| Quality | Quality | 3 | 18 |
| BXN0 | Bunch 0 flag | — | 1 |
| Synch. Error | | — | 1 |
| Total | | 27 | 164 |

Table 2: Information delivered from Drift Tube Interface to Sector Processor.

| variable | function | bits / muon | bits / 8 muons |
|---|---|---|---|
| $\phi$ | Phi Coordinate | 12 | 72 |
| $\phi_b$ | Local Bend Angle | 5 | 30 |
| Quality | Quality | 3 | 18 |
| BXN | Bunch i.d. | — | 4 |
| Synch. Error | | — | 1 |
| Total | | 20 | 165 |

Table 3: Information delivered to the Muon Sorter.

| variable | unit/precision | range | bits / muon | bits / 3 muons |
|---|---|---|---|---|
| $\phi$ | $2.5^\circ$ | 0–60° | 5 | 15 |
| $\eta$ | $0.075$ $\eta$ unit | $0.9$–$2.4$ | 5 | 15 |
| Rank | $p_T$ and Quality | 2–140 GeV | $5+2$ | 21 |
| muon sign | — | — | 1 | 3 |
| BXN | — | — | — | 4 |
| Synch. Error | — | — | — | 1 |
| Total | | | 18 | 59 |

Table 4: Track categories and a possible ranking from the output of the Track Assemblers. The "Coarse $p_{\mathrm{T}}$" column is the result of extrapolations involving station 1.

| Track Type | Coarse $p_{\mathrm{T}}$ | Priority of Track |
|---|---|---|
| Nothing | 0 | 0 |
| $ME3 \to ME4$ | – | 1 |
| $ME2 \to ME4$ | – | 2 |
| $ME2 \to ME3$ | – | 3 |
| $MB2 \to ME2$ | – | 4 |
| $ME2 \to ME3 \to ME4$ | – | 5 |
| $MB2 \to ME1 \to ME2$ | – | 6 |
| $ME1 \to ME3$ | low | 7 |
| $ME1 \to ME2$ | low | 8 |
| $MB1 \to ME2$ | low | 9 |
| $ME1 \to ME3$ | medium | 10 |
| $ME1 \to ME2$ | medium | 11 |
| $MB1 \to ME2$ | medium | 12 |
| $ME1 \to ME3$ | high | 13 |
| $ME1 \to ME2$ | high | 14 |
| $MB1 \to ME2$ | high | 15 |
| $ME1 \to ME3 \to ME4$ | low | 16 |
| $ME1 \to ME2 \to ME4$ | low | 17 |
| $ME1 \to ME2 \to ME3$ | low | 18 |
| $MB1 \to ME1 \to ME2$ | low | 19 |
| $MB1 \to MB2 \to ME2$ | low | 20 |
| $ME1 \to ME3 \to ME4$ | medium | 21 |
| $ME1 \to ME2 \to ME4$ | medium | 22 |
| $ME1 \to ME2 \to ME3$ | medium | 23 |
| $MB1 \to ME1 \to ME2$ | medium | 24 |
| $MB1 \to MB2 \to ME2$ | medium | 25 |
| $ME1 \to ME3 \to ME4$ | high | 26 |
| $ME1 \to ME2 \to ME4$ | high | 27 |
| $ME1 \to ME2 \to ME3$ | high | 28 |
| $MB1 \to ME1 \to ME2$ | high | 29 |
| $MB1 \to MB2 \to ME2$ | high | 30 |
| $ME1 \to ME2 \to ME3 \to ME4$ | low | 31 |
| $MB1 \to MB2 \to ME1 \to ME2$ | low | 32 |
| $ME1 \to ME2 \to ME3 \to ME4$ | medium | 33 |
| $MB1 \to MB2 \to ME1 \to ME2$ | medium | 34 |
| $ME1 \to ME2 \to ME3 \to ME4$ | high | 35 |
| $MB1 \to MB2 \to ME1 \to ME2$ | high | 36 |

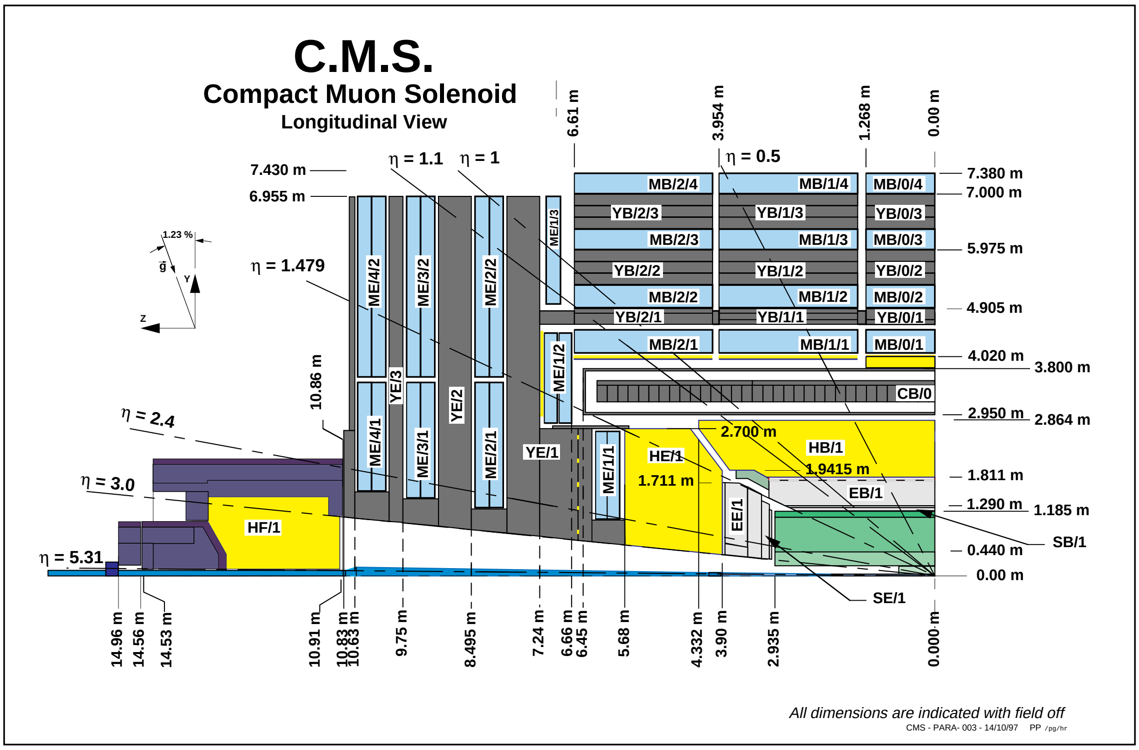


Figure 1: Cross-sectional view of the CMS detector.

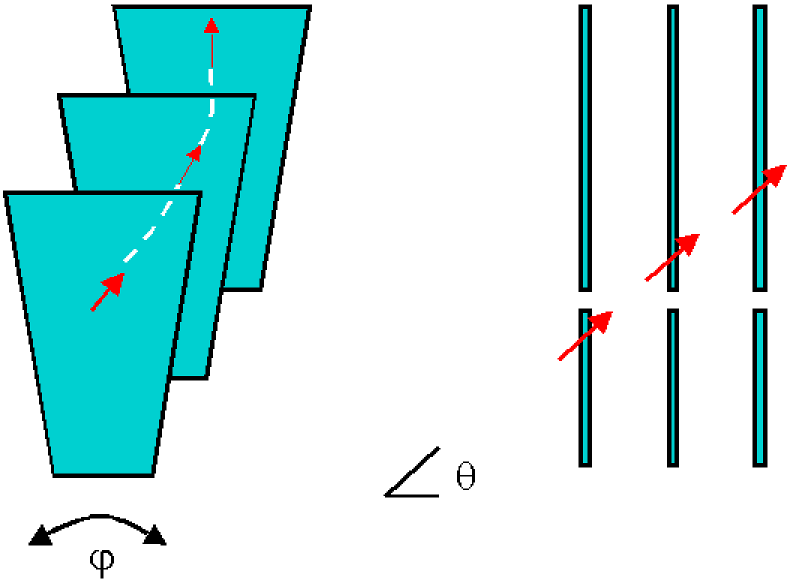


Figure 2: Principle of linking track segments into complete tracks in three dimensions by the CSC Track-Finder.

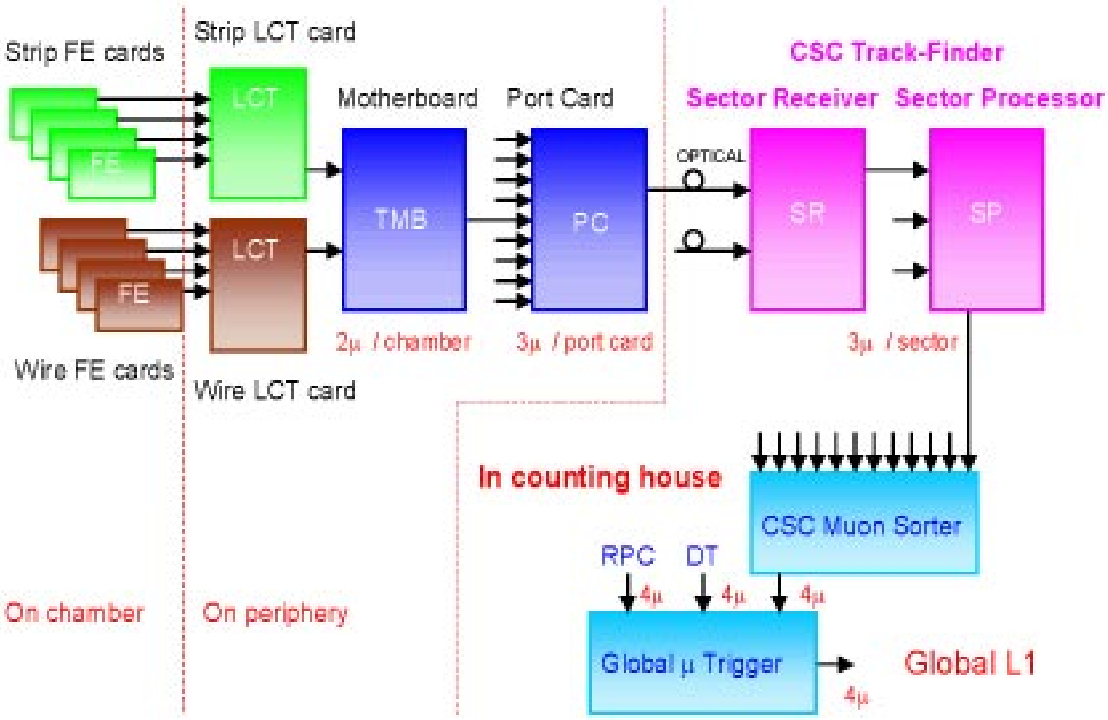


Figure 3: The Level-1 Trigger scheme for the endcap muon system.

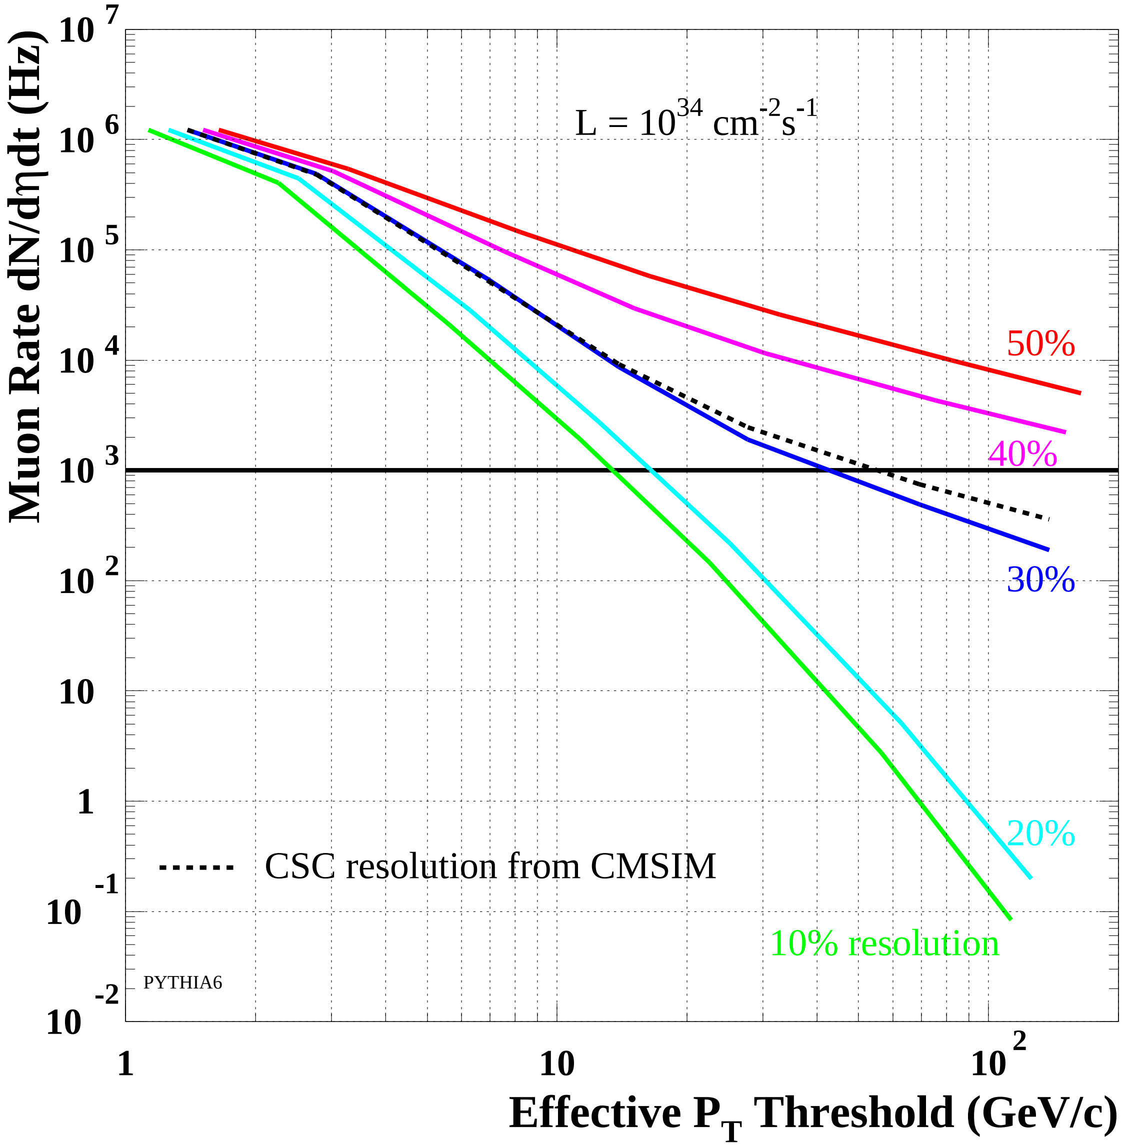


Figure 4: Expected Level-1 trigger rate per unit pseudorapidity from the CSC muon system as a function of the $p_T$ threshold. Various choices of the $p_T$ resolution are shown. The threshold is defined at 90% efficiency. The muon flux estimate is taken from PYTHIA, and the probability to penetrate the calorimeter is folded in.

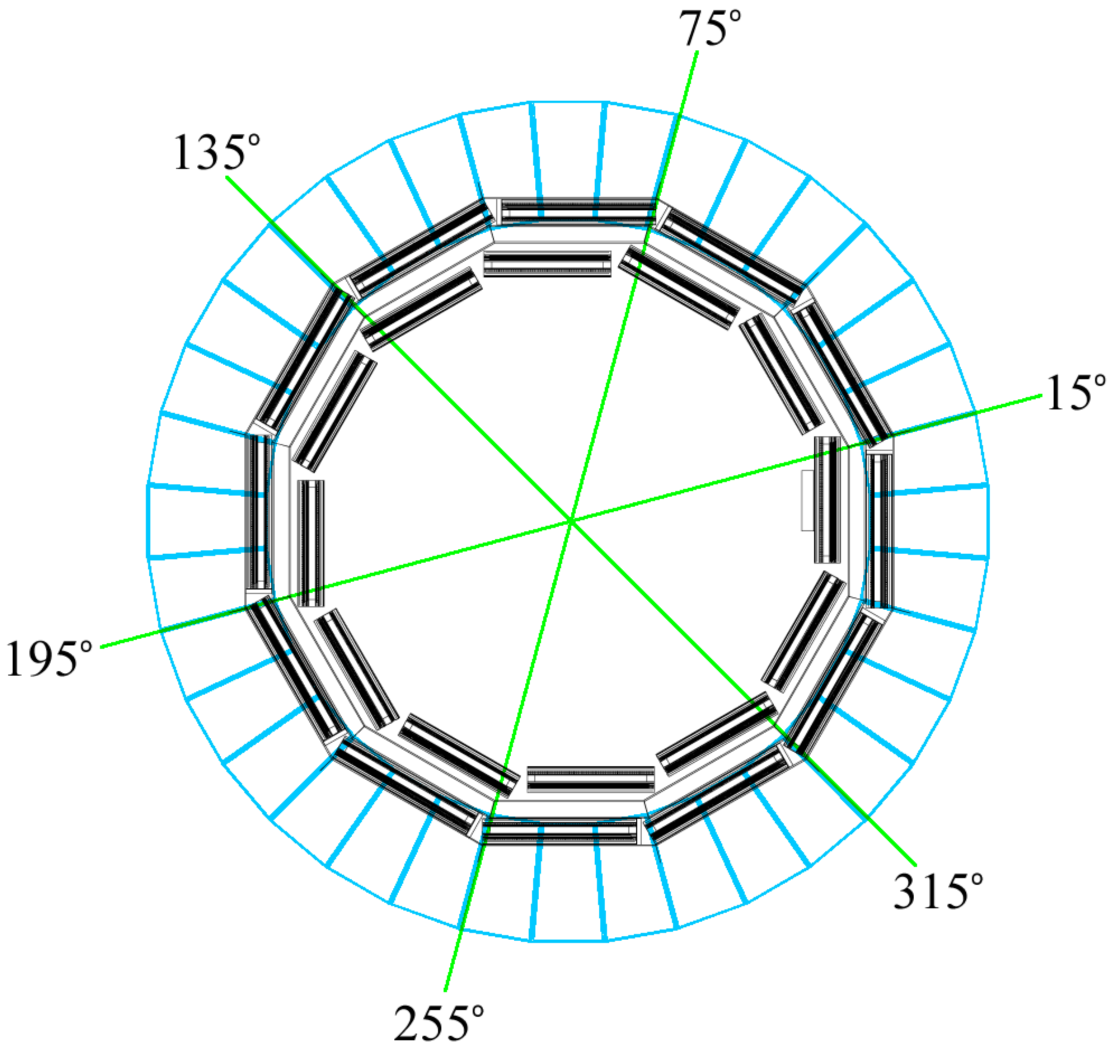


Figure 5: Layout of the $60^{\circ}$ sectors in the overlap region between the barrel and endcap muon systems. Shown are the MB1, MB2, and ME1/3 chambers. Information from neighboring sectors is not shared across boundaries.

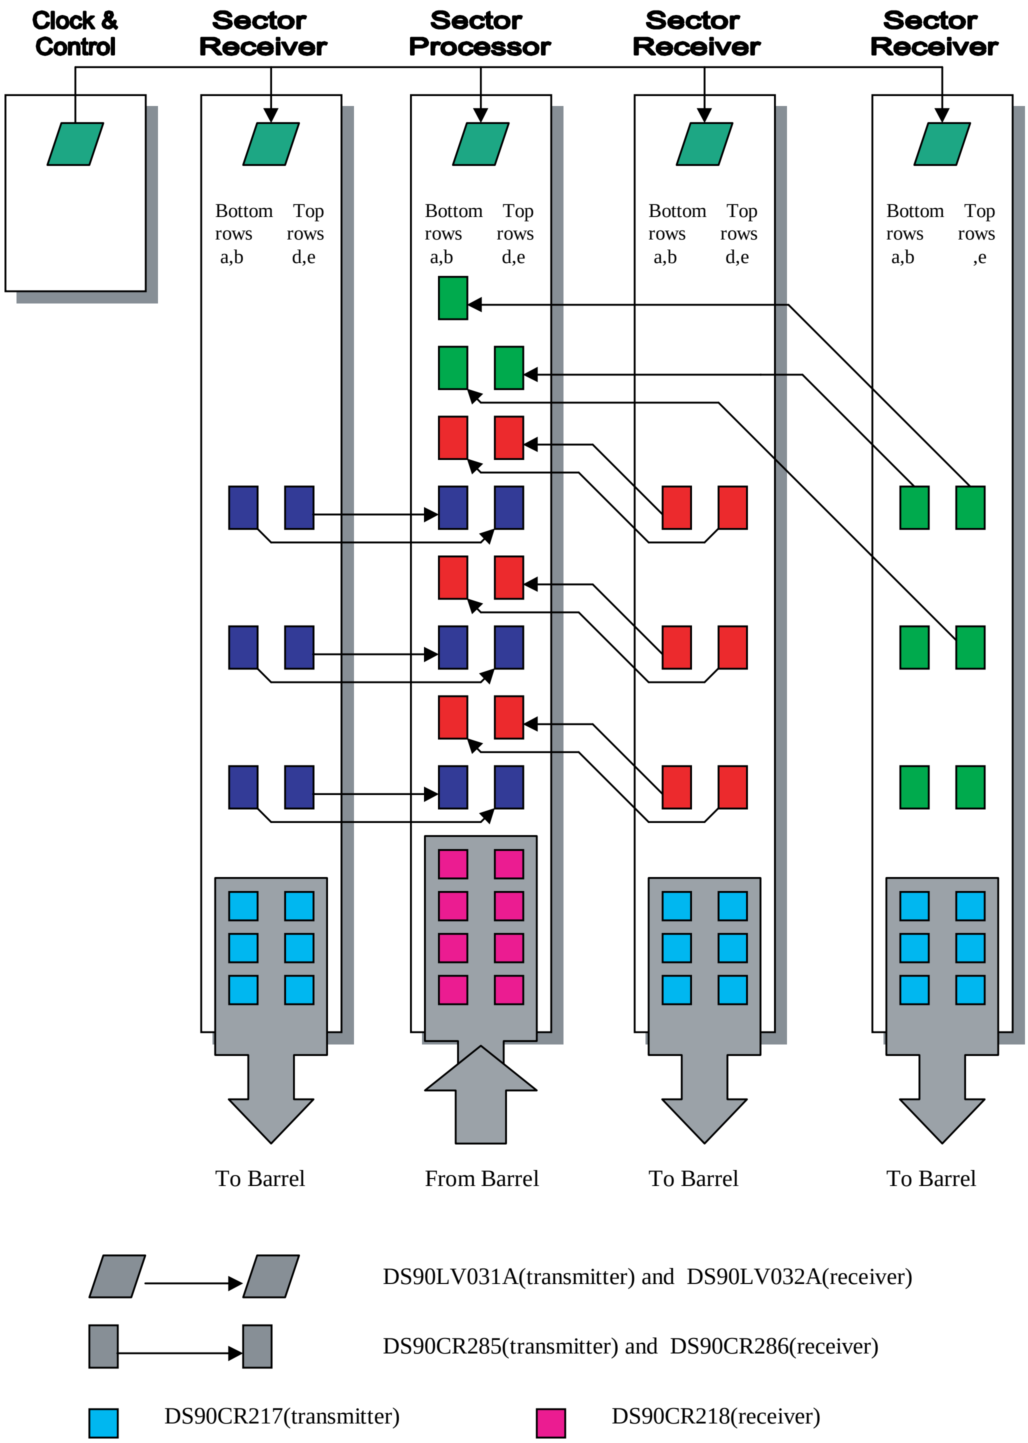


Figure 6: Illustration of the card placement and backplane connections for one sector in the Track-Finder crate. There will be two sectors per crate.

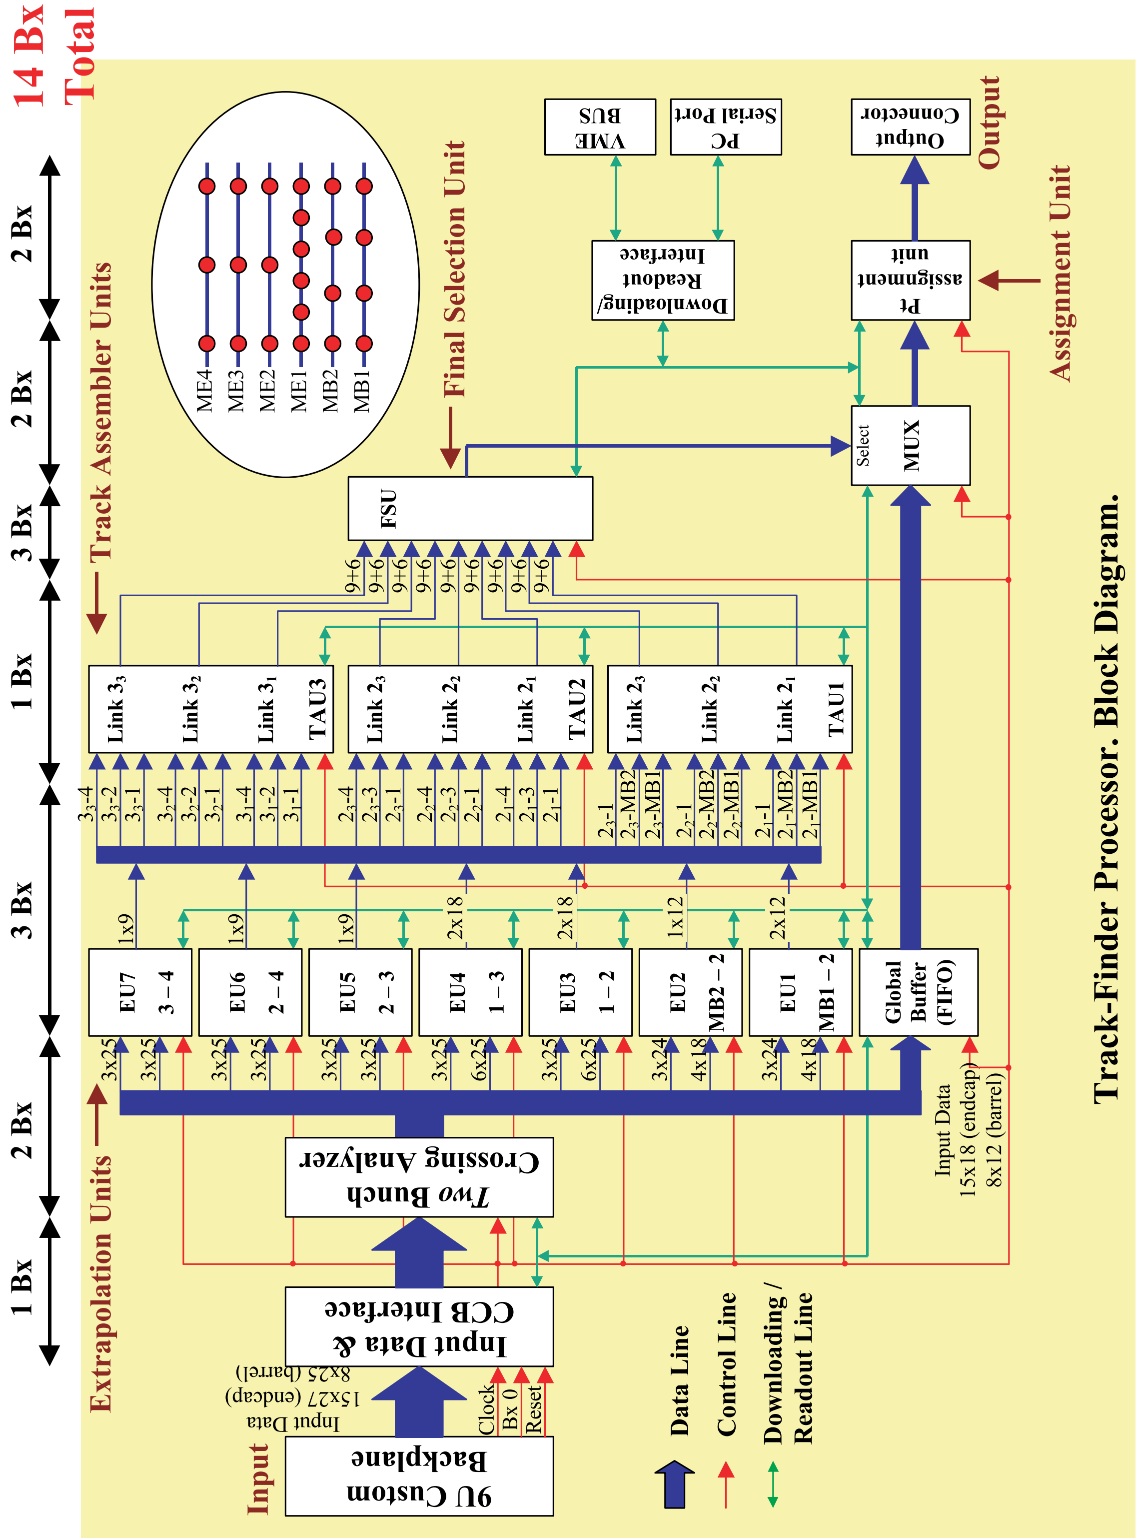


Figure 7: Block diagram of the Sector Processor architecture.

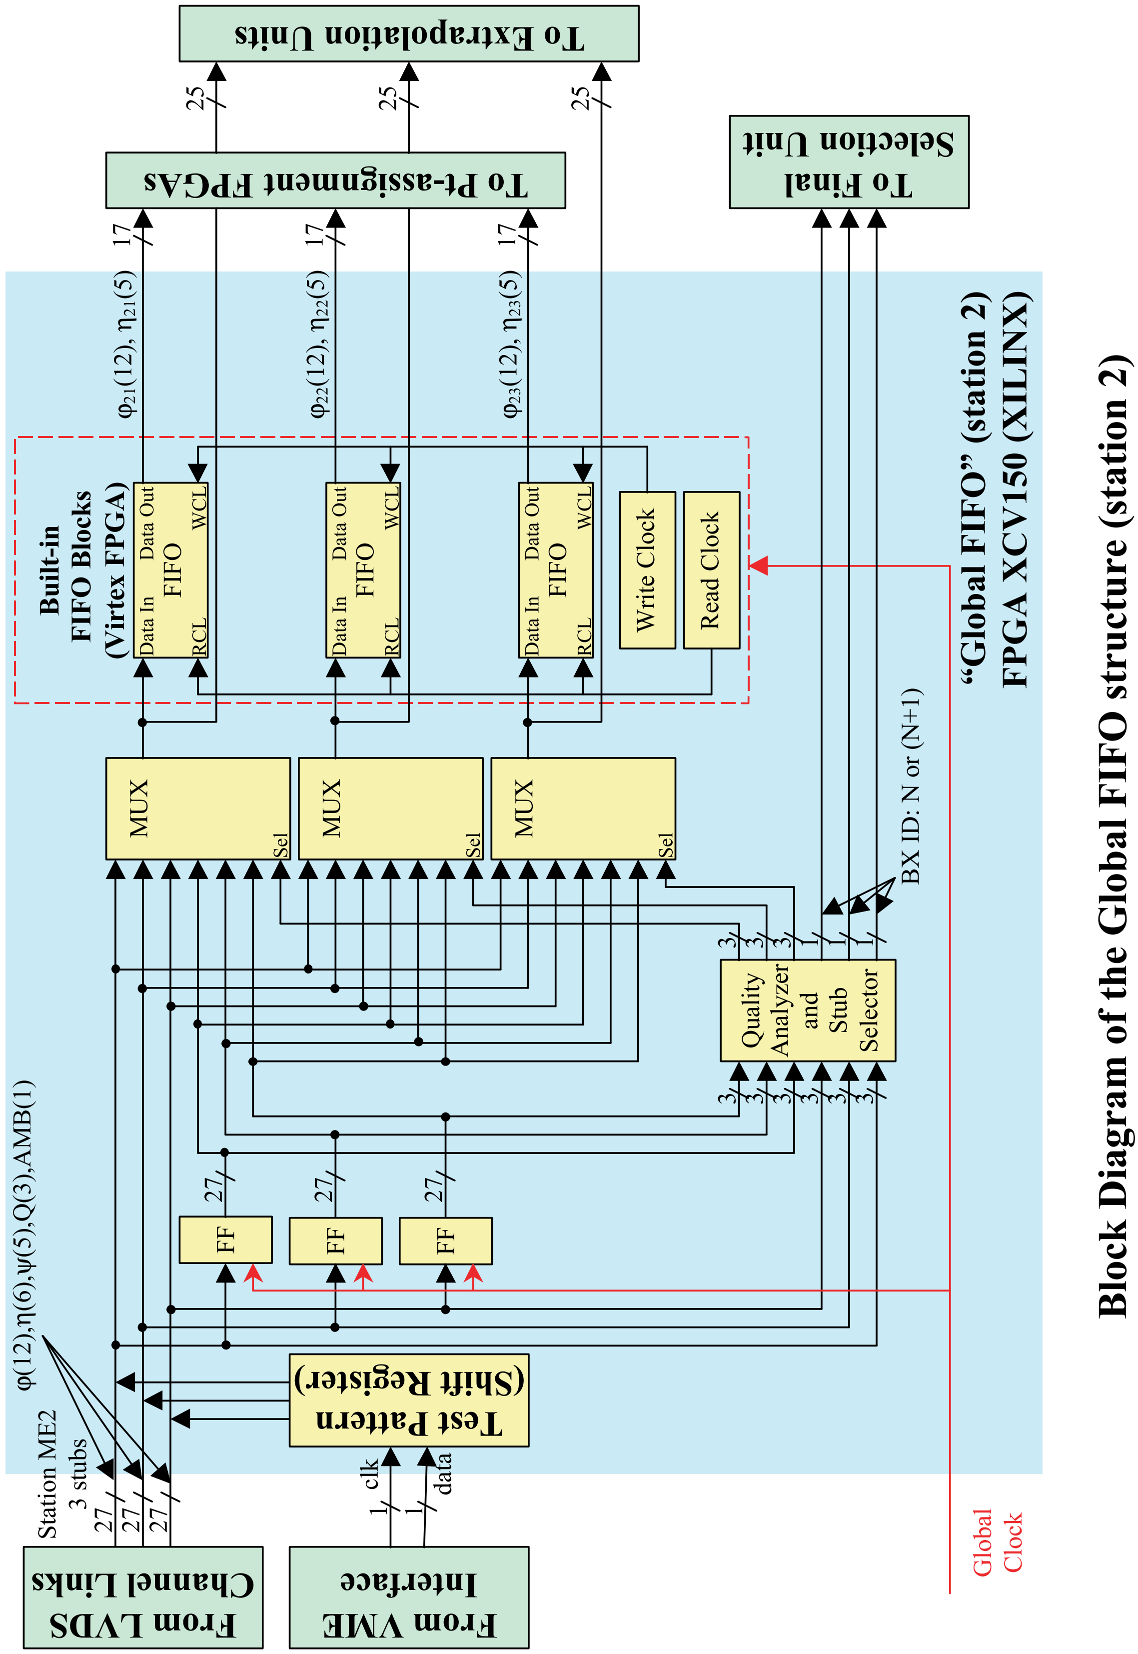


Figure 8: Block diagram of the data input section of the Sector Processor with the capability to analyze track segments from multiple bunch crossings.

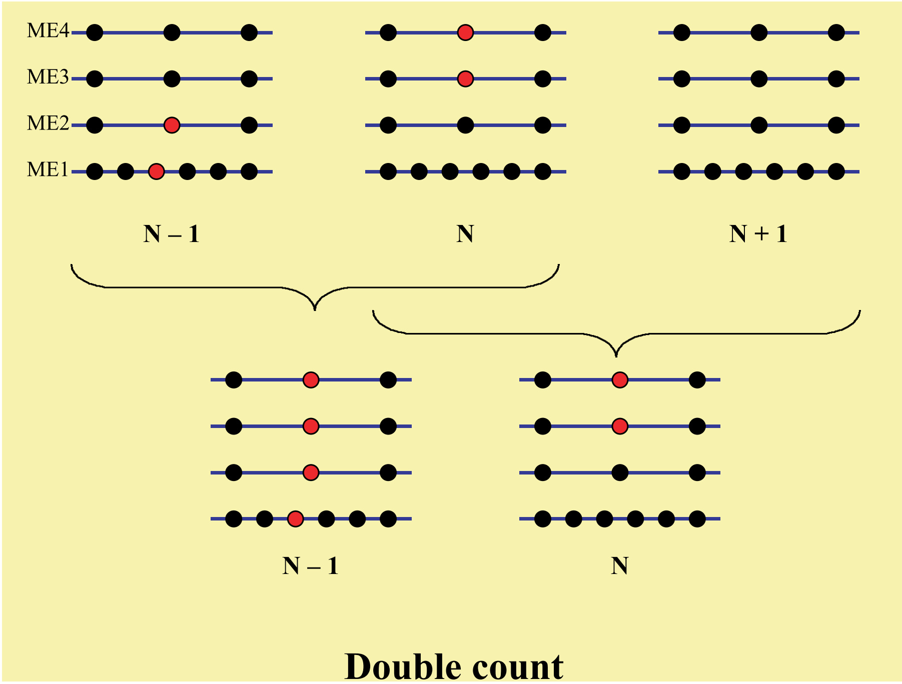


Figure 9: Illustration of the occurrence of a double trigger caused by a muon with track segments in two bunch crossings, when the input window to the Track-Finder is two bunch crossings wide.

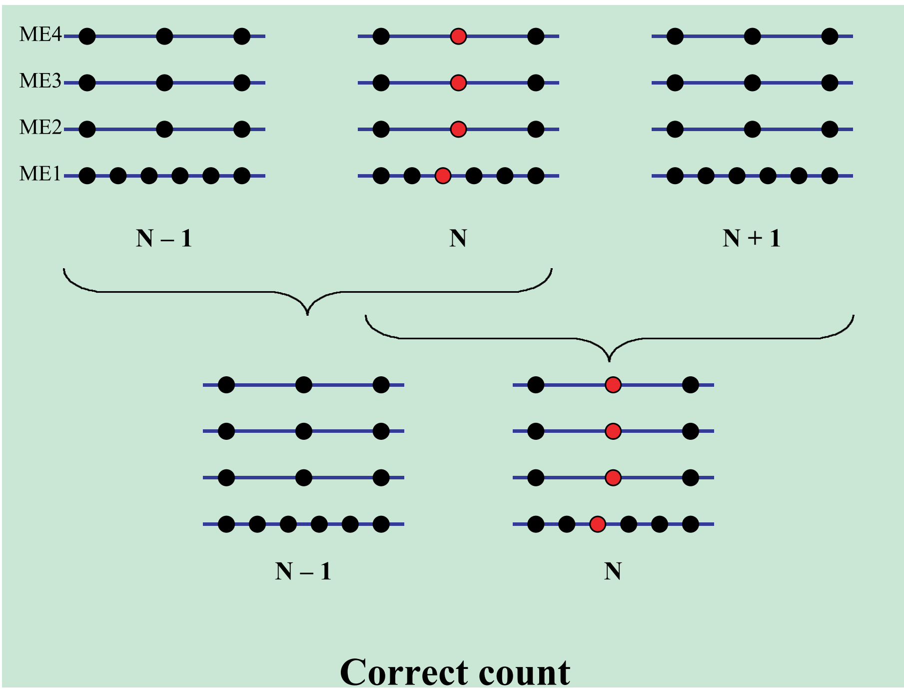


Figure 10: Illustration showing that a double trigger can be avoided if all track segments lie in the same bunch crossing, but only if extra logic is added to inhibit the trigger at the (wrong) earlier crossing.

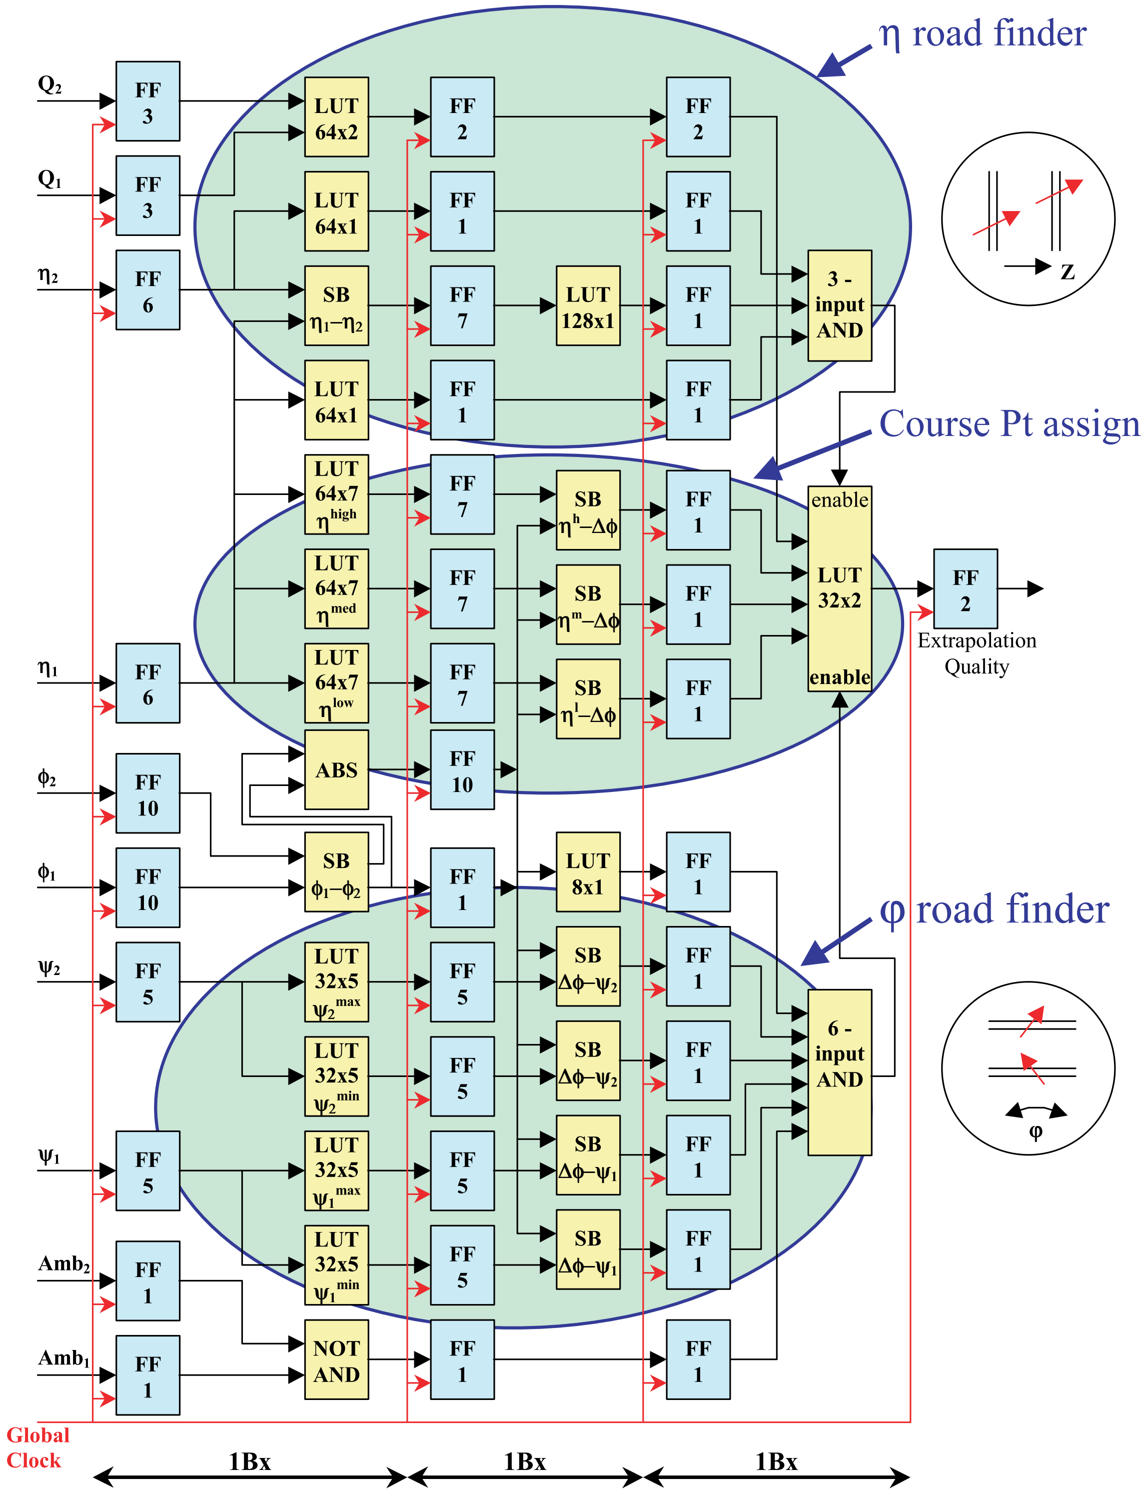


**Extrapolation Unit for ME1-ME2, ME1-ME3.**

Figure 11: Logic diagram for an extrapolation unit in the endcap region which includes a track segment from ME1.

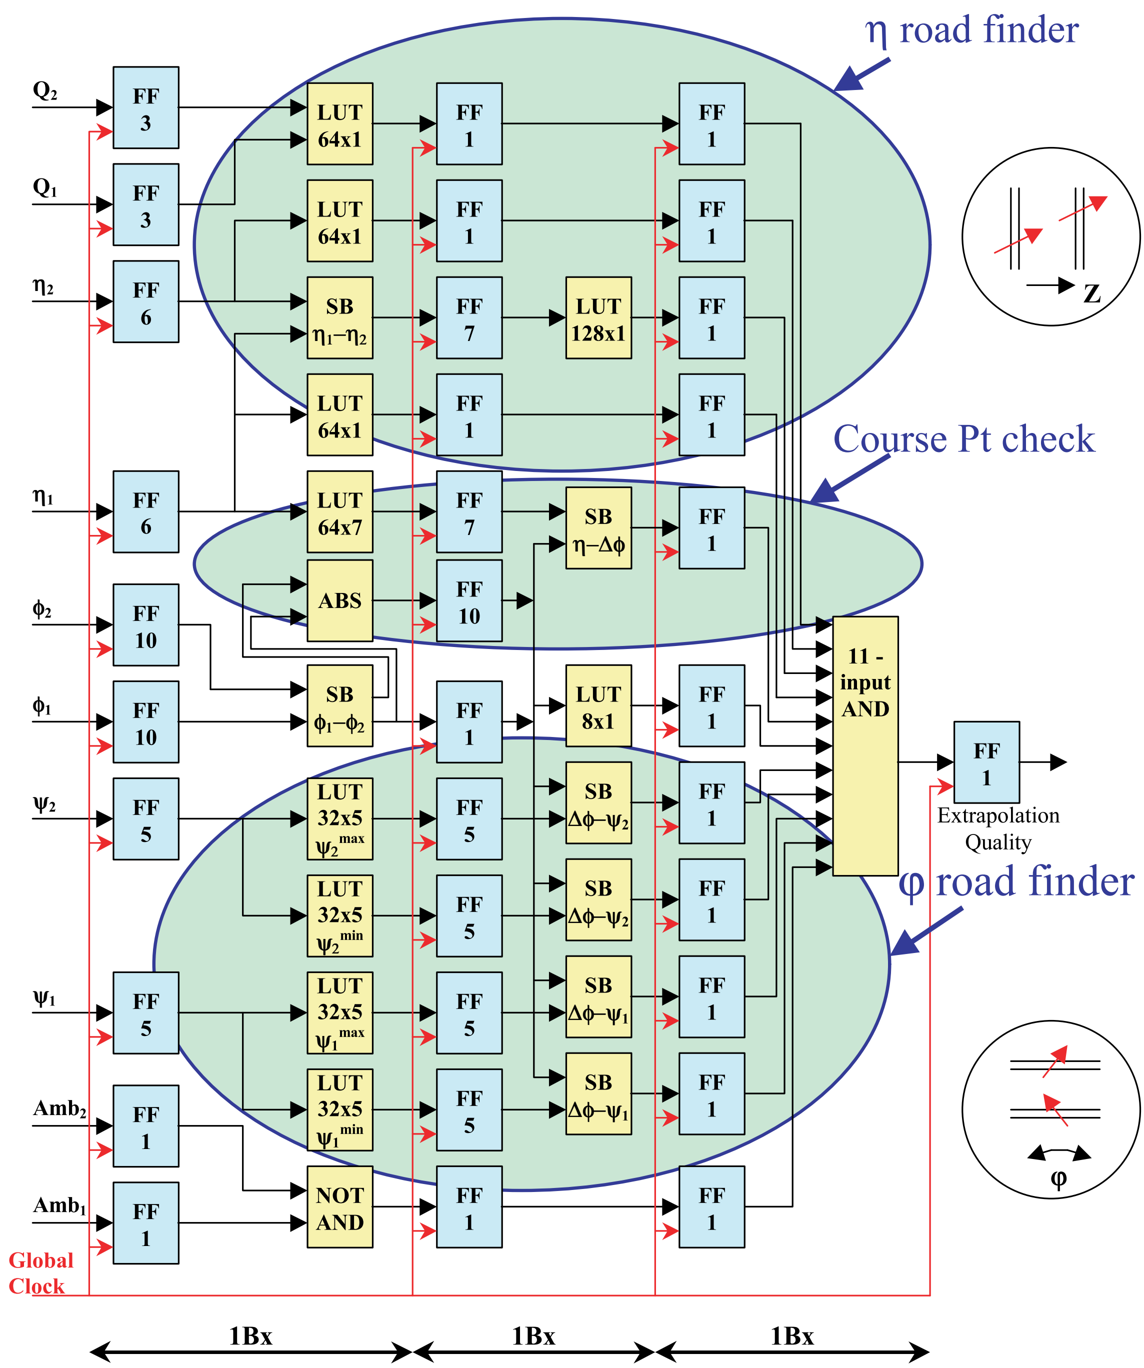


**Extrapolation Unit for ME2-ME3, ME2-ME4, ME3-ME4.**

Figure 12: Logic diagram for an extrapolation unit in the endcap region which does not include a track segment from ME1.

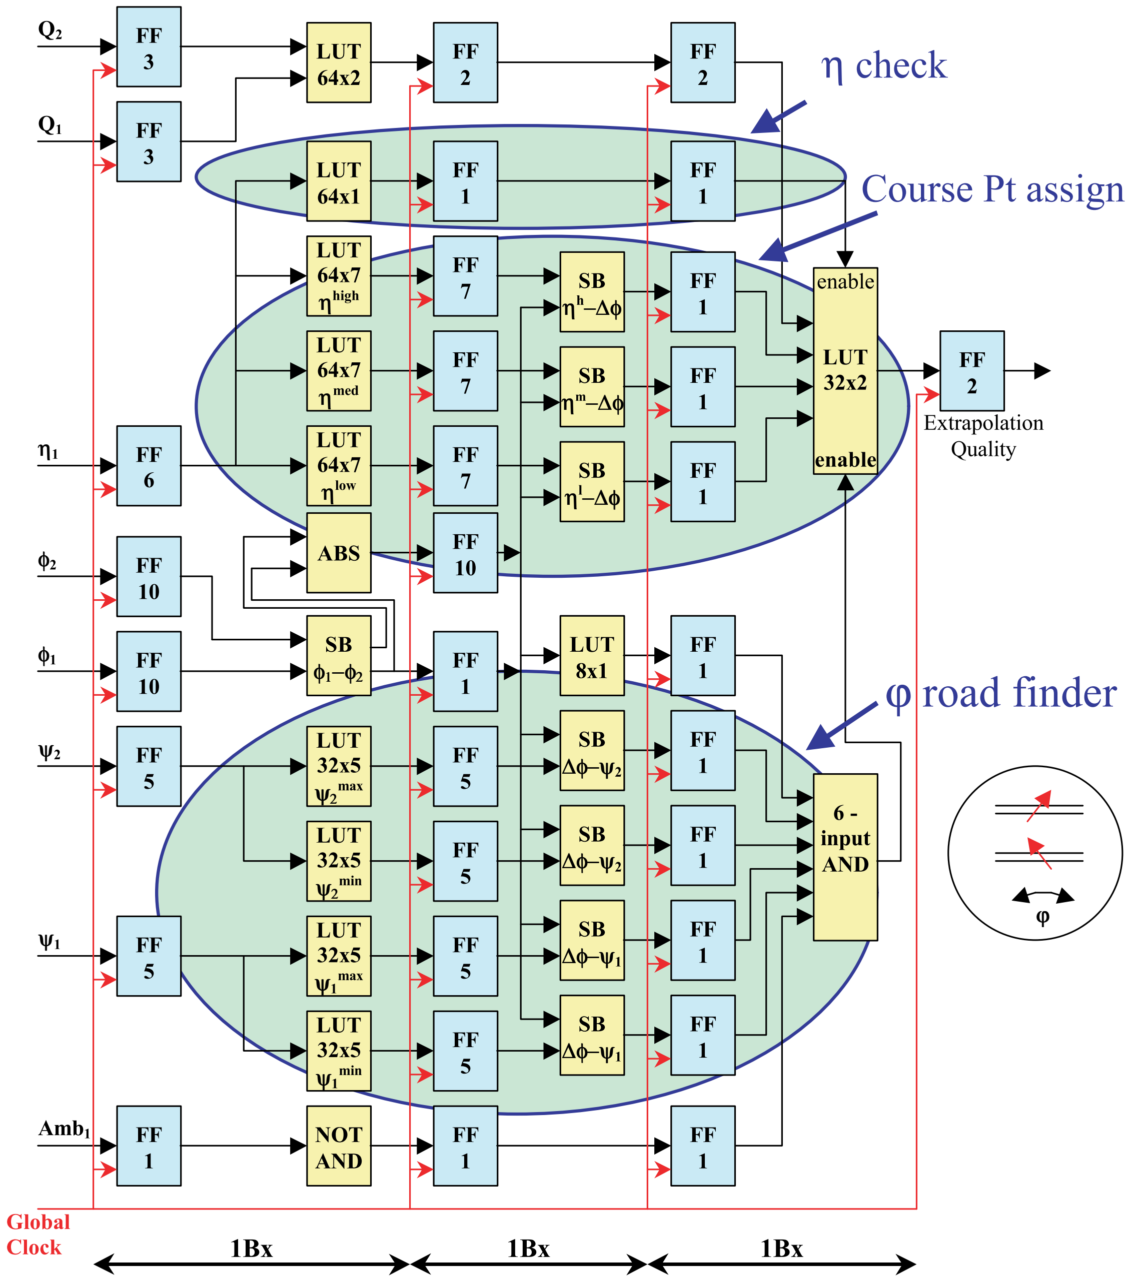


**Extrapolation Unit for MB1-ME2.**

Figure 13: Logic diagram for an extrapolation unit in the overlap region which includes a track segment from MB1.

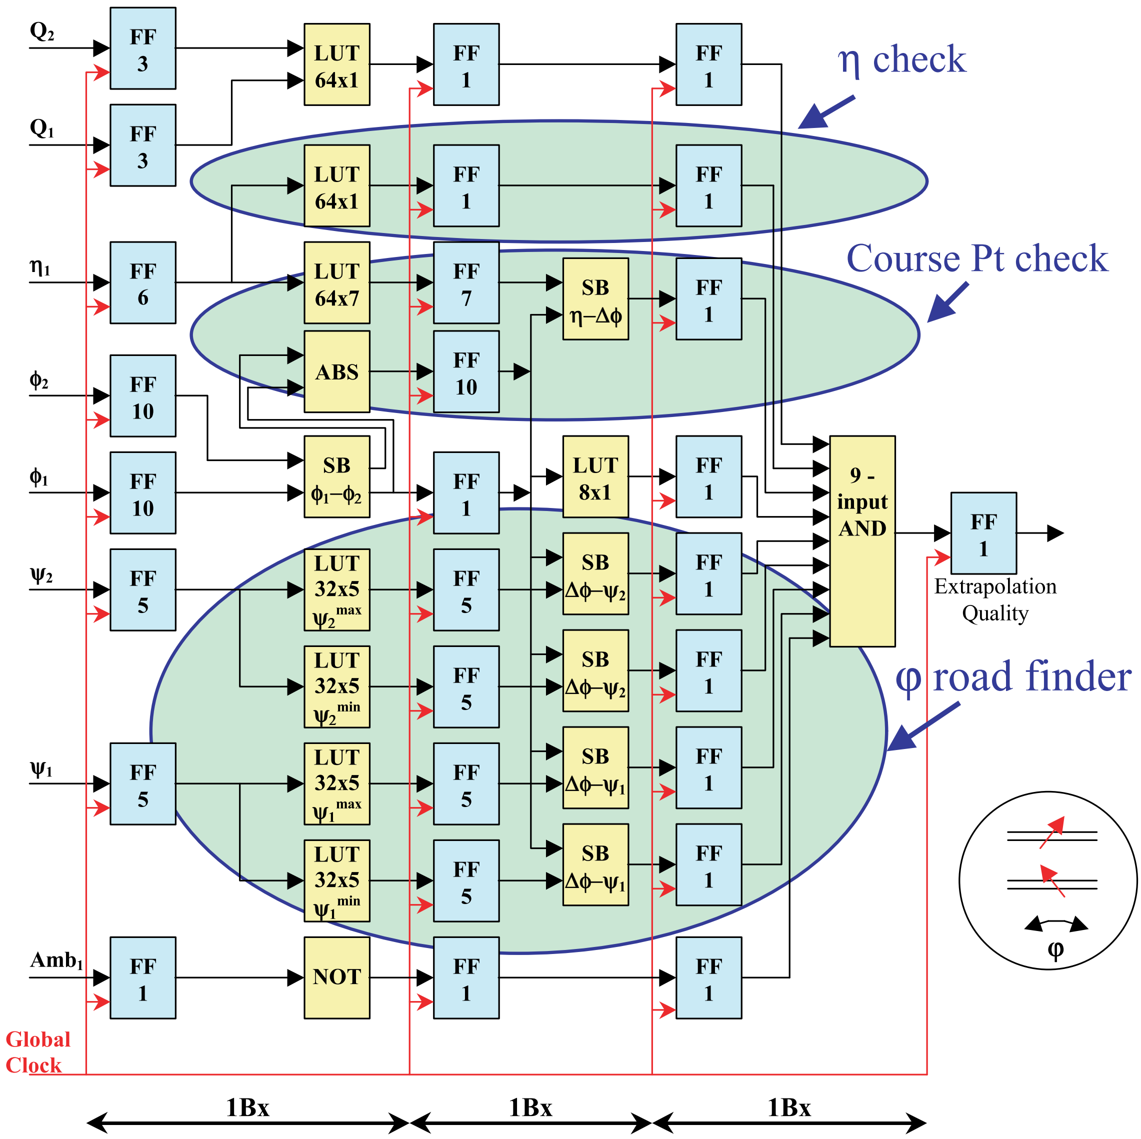


**Extrapolation Unit for MB2-ME2.**

Figure 14: Logic diagram for an extrapolation unit in the overlap region which does not include a track segment from MB1.

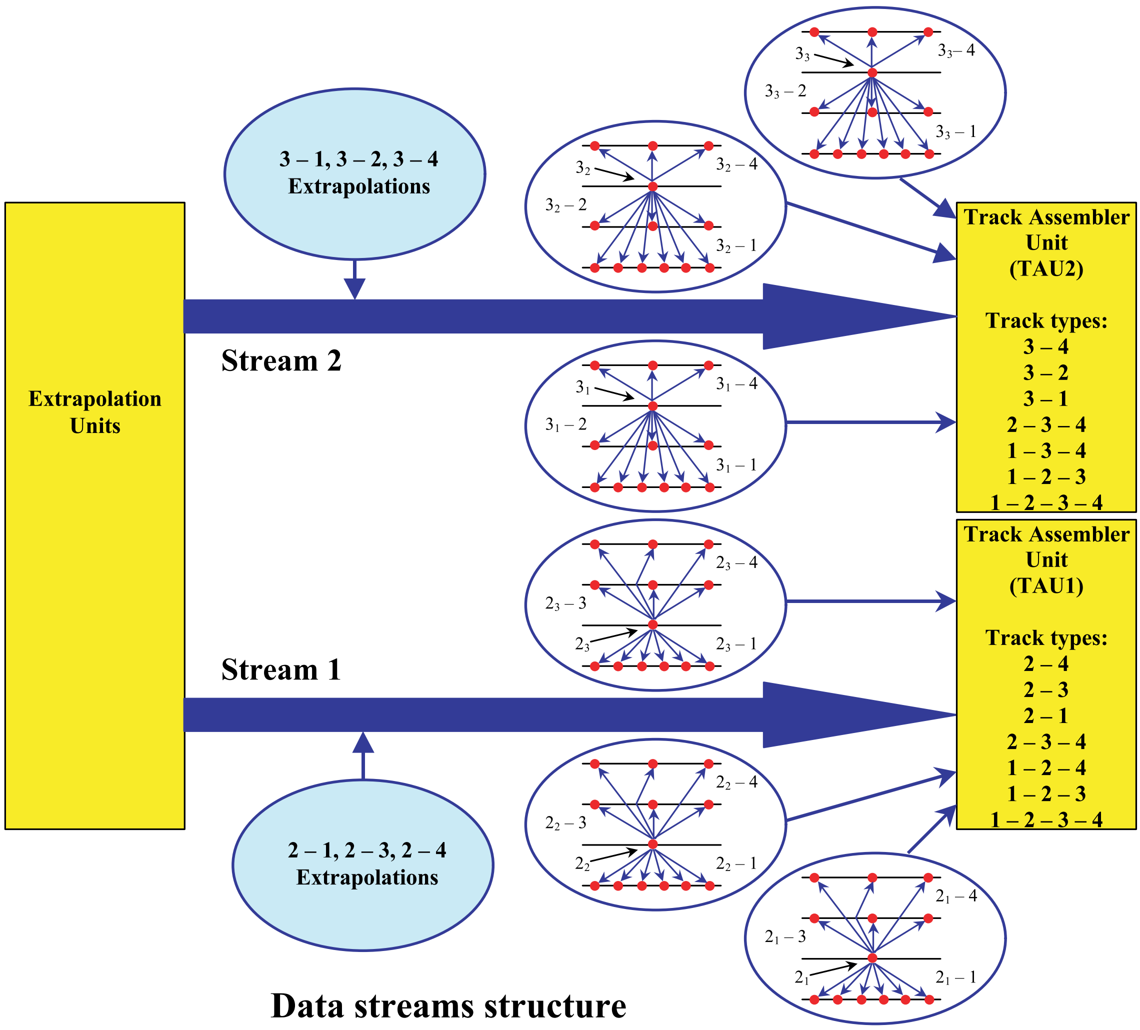


Figure 15: Illustration of the track assembly procedure for the endcap region separated into two data streams.

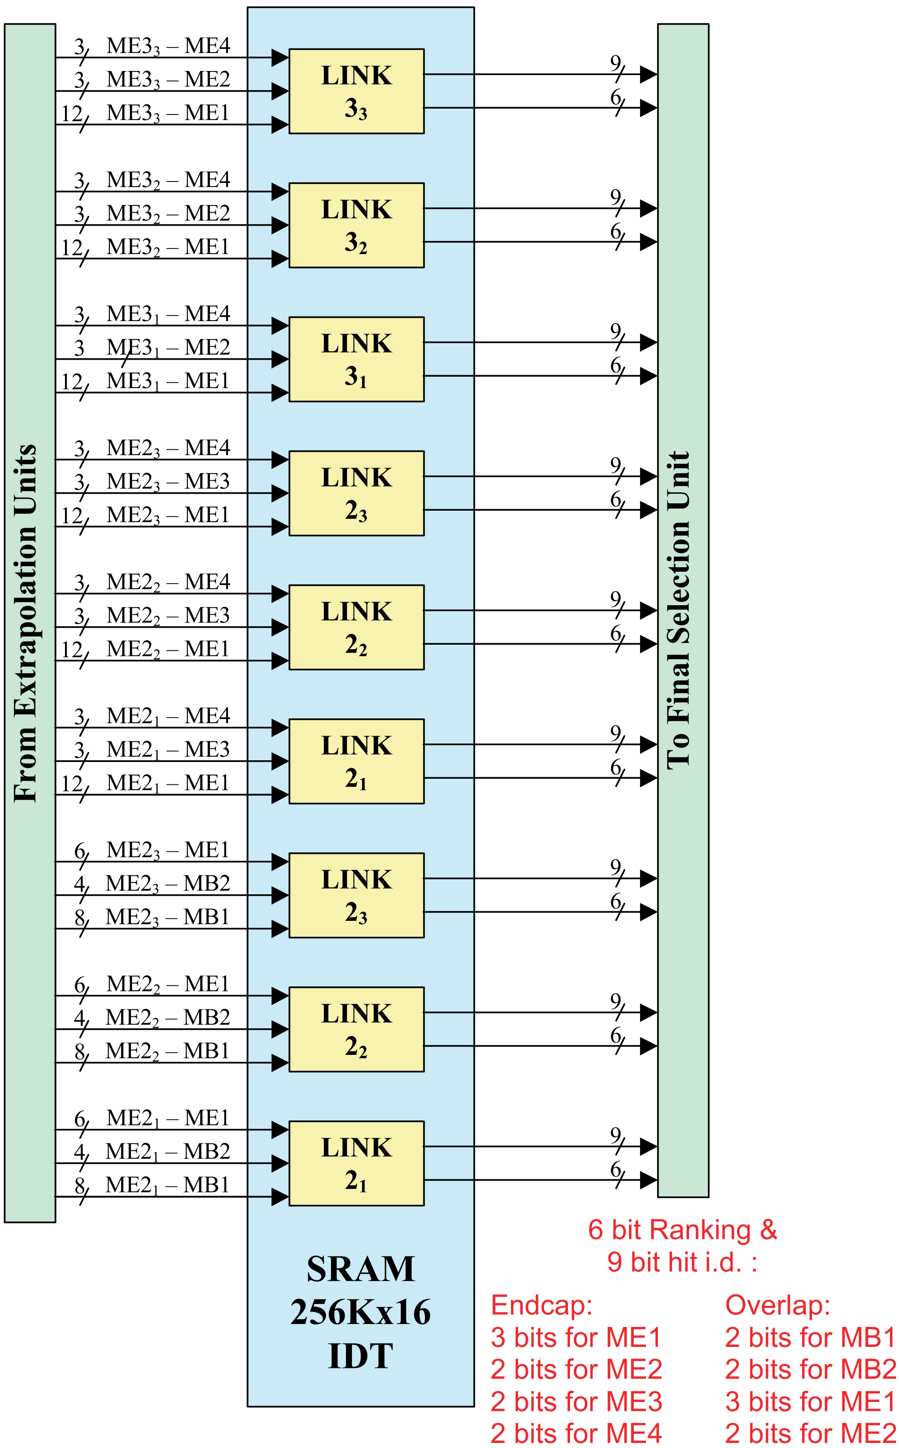


Figure 16: The Track Assembler Unit implemented as 9 static RAM memories for the endcap and overlap region.

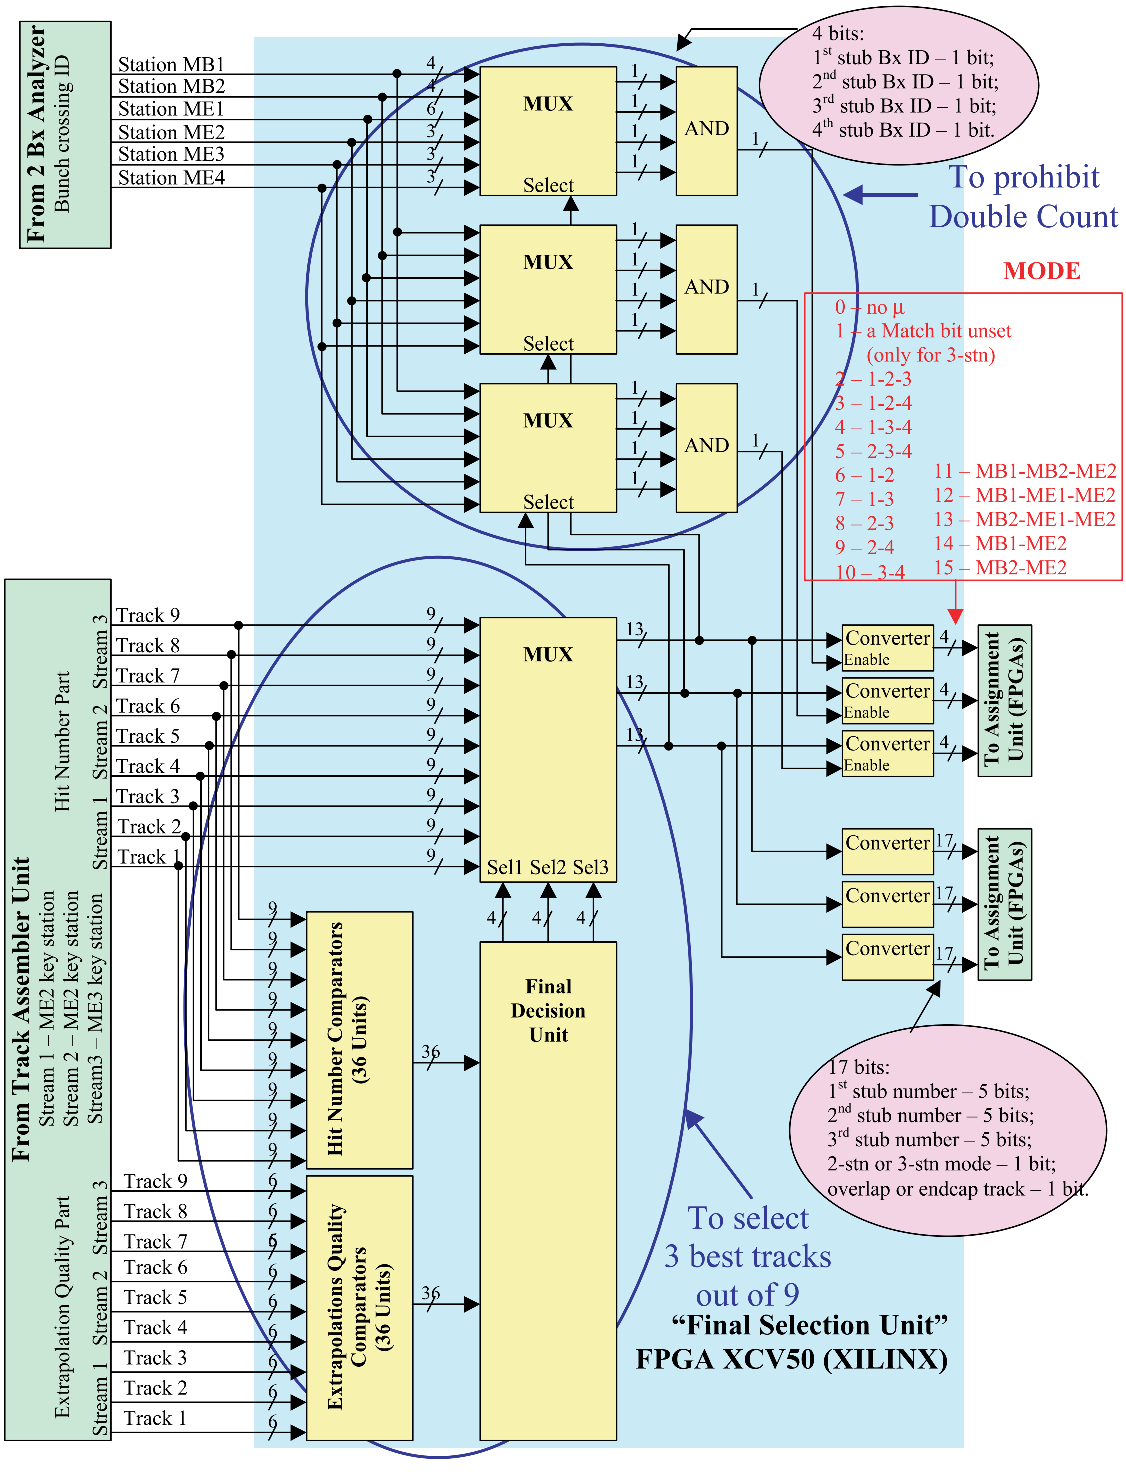


Figure 17: Block diagram of the Final Selection Unit.

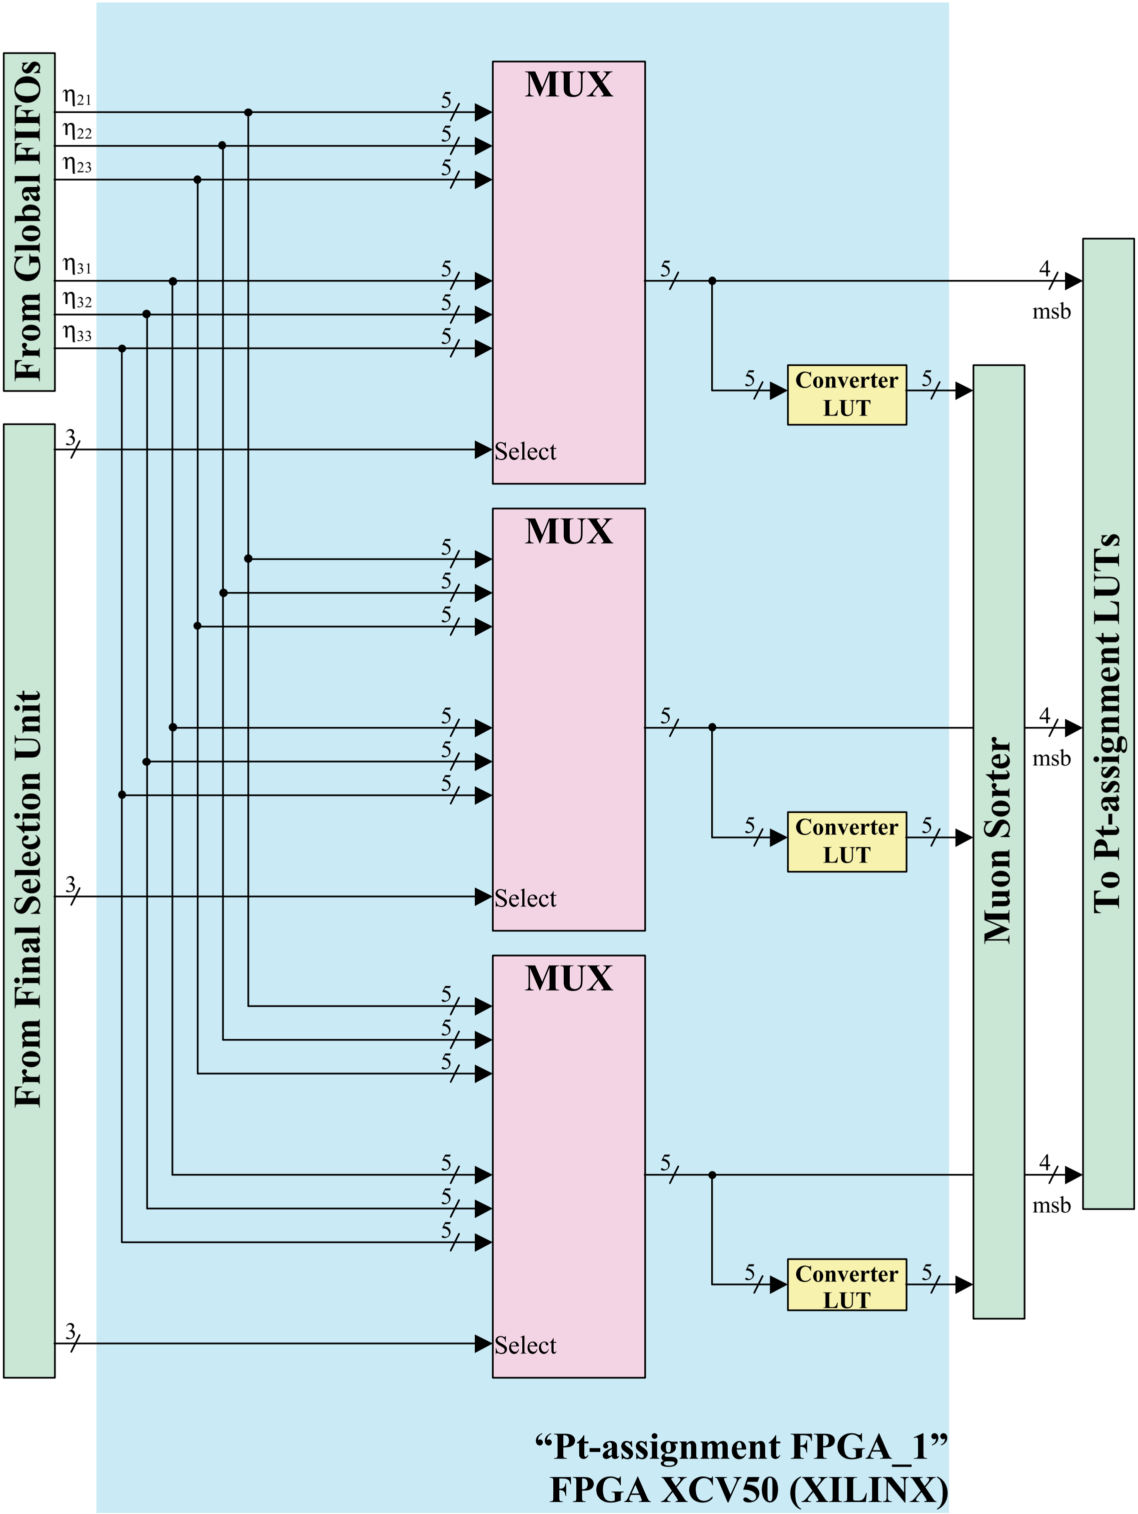


Figure 18: First step in the Assignment Unit to extract the needed data and to calculate some simple quantities.

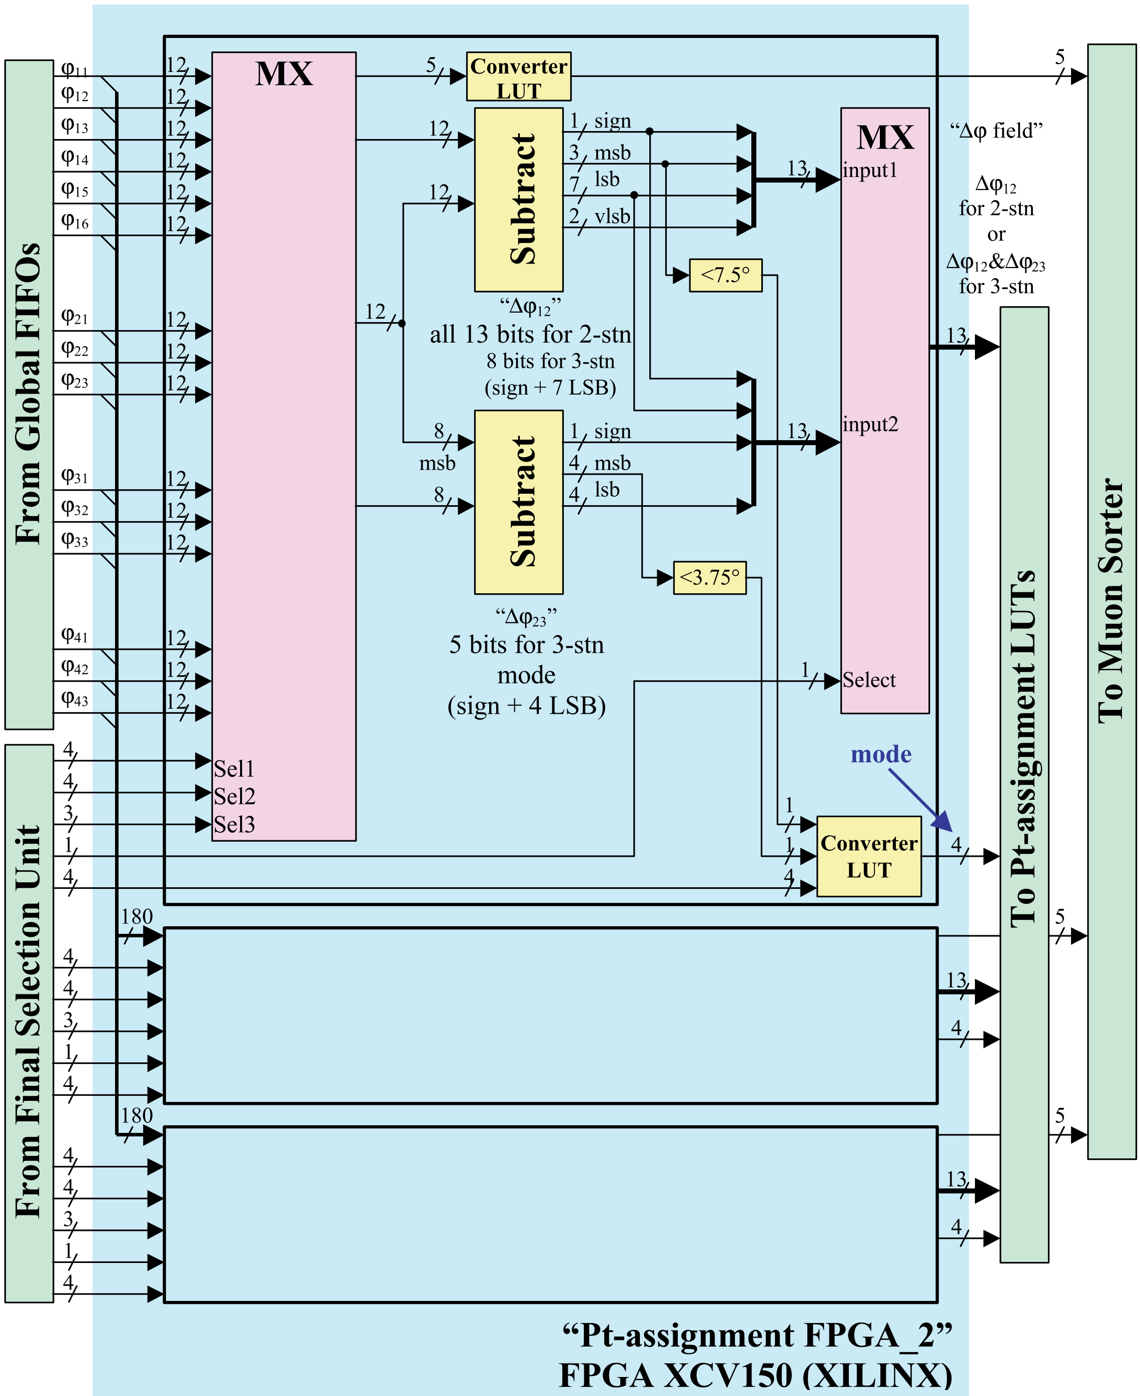


Figure 19: Second step in the Assignment Unit to calculate the difference in $\phi$ between the first and second station, and between the second and third (when it is available).

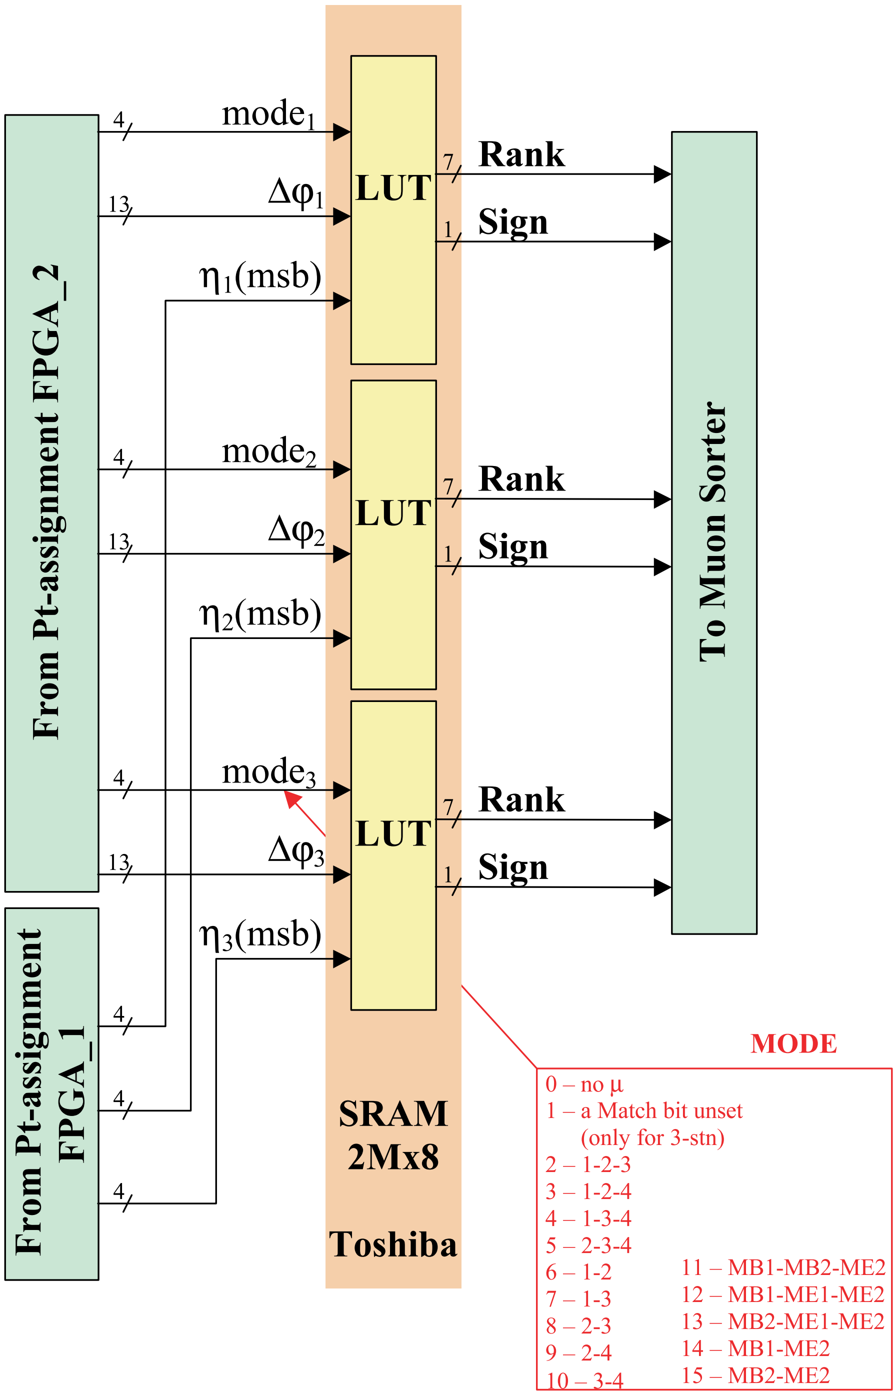


Figure 20: Final assignment of the signed $p_{\mathrm{T}}$ and quality (stored as a 7-bit rank and a sign bit) in three SRAM memories.

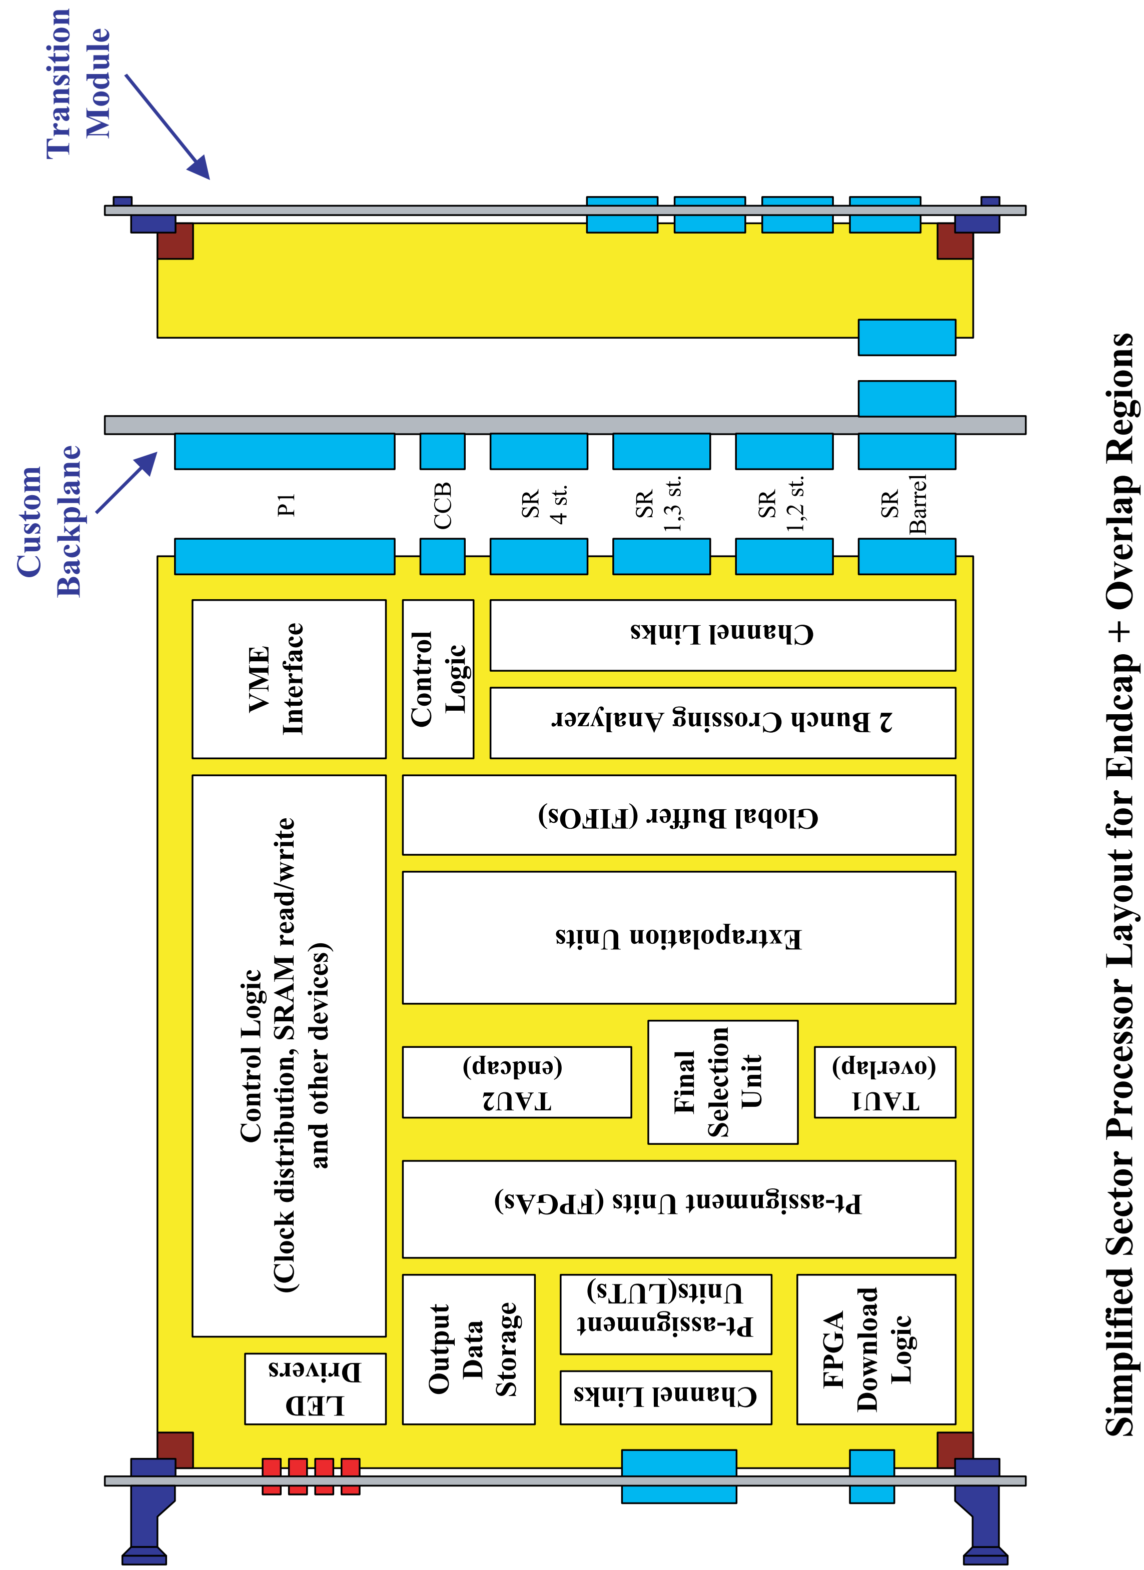


Figure 21: Functional layout of the Sector Processor on a 9U VME board, including processing of the DT/CSC overlap.

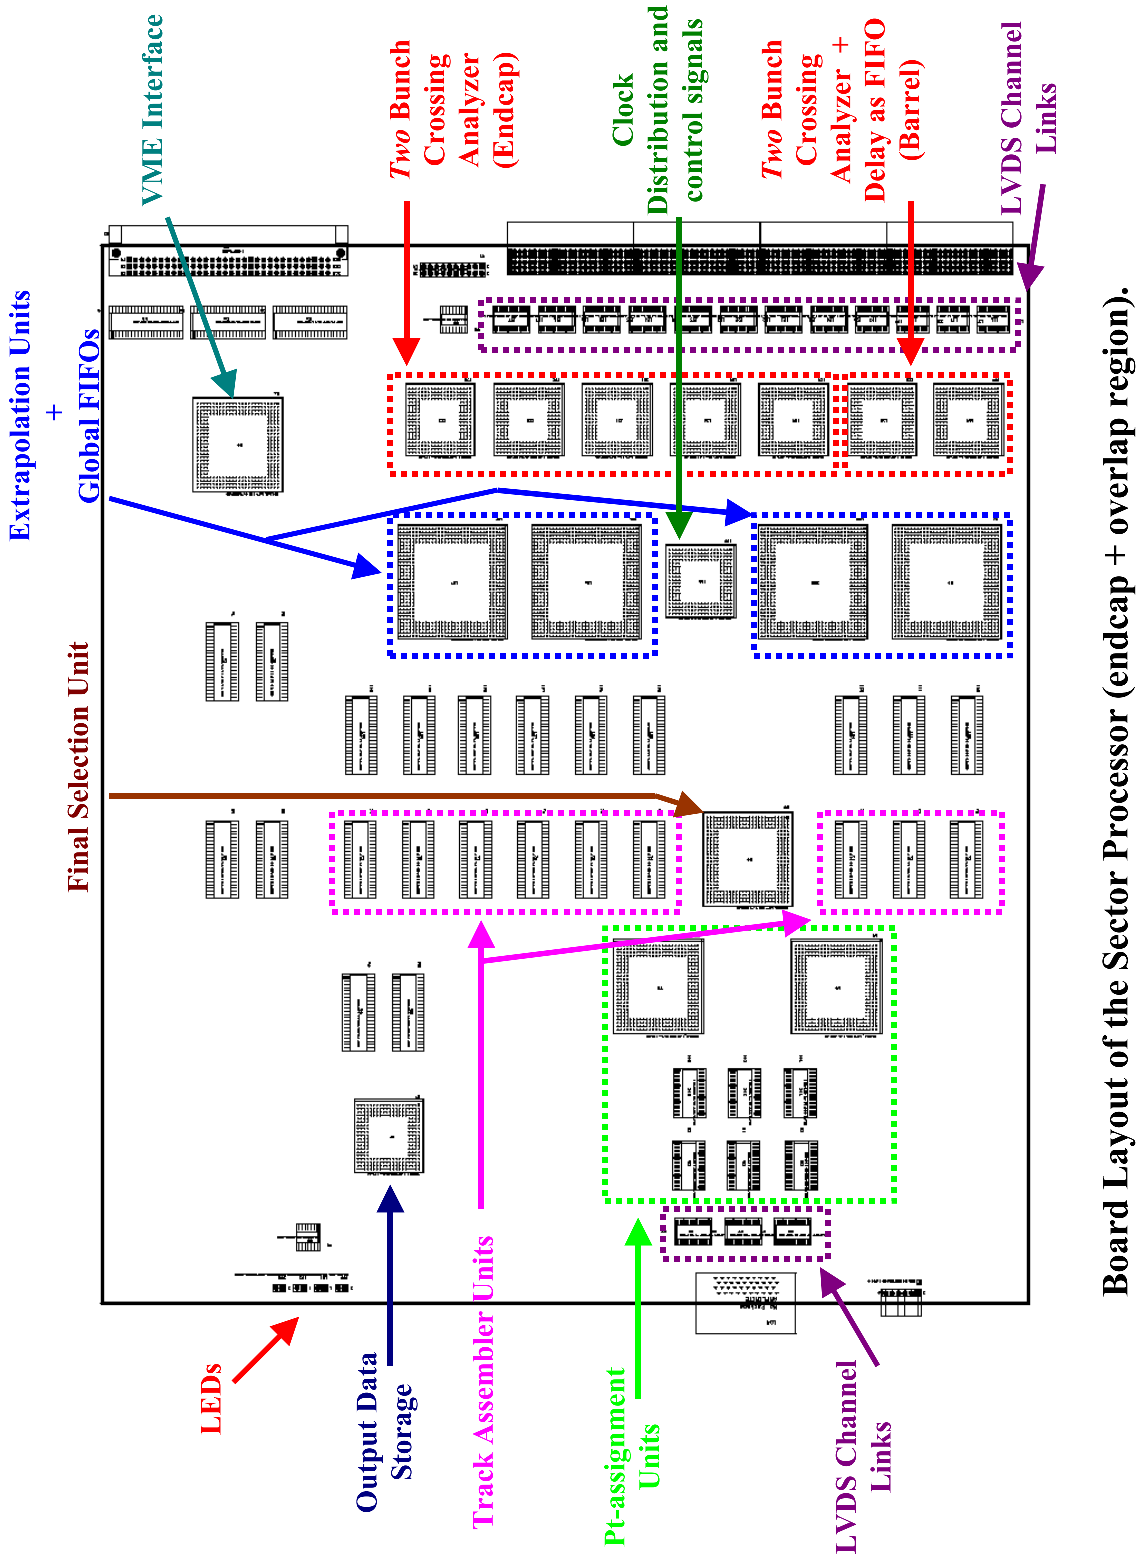


Figure 22: Preliminary layout consideration of the Sector Processor on a 9U VME board, including processing of the DT/CSC overlap.