12 Cathode Strip Chamber Track-Finder

12.1 Requirements

12.1.1 Physics Requirements

The L1 trigger electronics of the CMS muon system must measure the momentum of penetrating particles in order to reduce the several megahertz rate of low-momentum muons produced at full LHC luminosity. The task of the Cathode Strip Chamber (CSC) Track-Finder is to reconstruct tracks in the CSC endcap muon system and to measure the transverse momentum (p_T), pseudo-rapidity (η), and azimuthal angle (ϕ) of each muon. Although this task is the same for the Drift-Tube (DT) and CSC muon systems, the optimization of the design of the Track-Finder is significantly different for each muon system because of the different logical partitioning of the trigger primitives and the non-axial magnetic field in the endcap region. The algorithms of the CSC Track-Finder are inherently 3-dimensional to achieve maximum background rejection, as illustrated in Fig. 12.1. Moreover, the measurement of p_T uses spatial information from up to three stations to achieve a precision similar to that of the DT Track-Finder despite the reduced magnetic bending in the endcap. A p_T resolution of 25% is necessary (see Ref. [12.1]) to have sufficient rate reduction at L1 with a reasonable threshold.

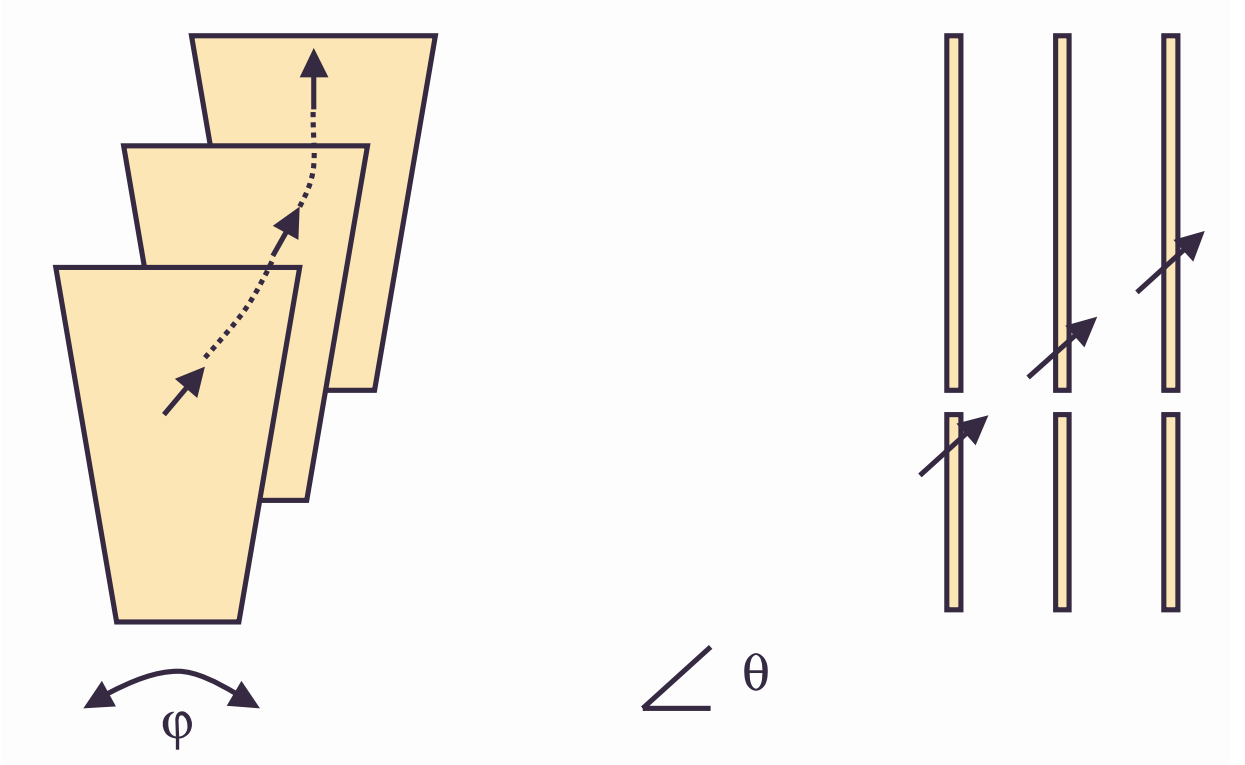

Fig. 12.1: Illustration of the three-dimensional track-finding procedure.

12.1.2 Boundary Between DT and CSC Track-Finders

The region of overlap between the DT and CSC muon systems should be covered as efficiently as possible by the L1 trigger; however, it is important to define a sharp boundary between the DT and CSC Track-Finders to avoid a duplication of triggers for single muons in this region. Issues related to this separation are documented in Refs. [12.2] and [12.3].

The boundary in the coverage between the DT and CSC Track-Finders is currently set at $\eta = 1.04$, as shown in Fig. 12.2. This extends the coverage of the DT Track-Finder to the high-η limit of the MB2/2 chambers, and the coverage of the CSC Track-Finder extends approximately to the outer radius of the ME2/2 chambers. This boundary also corresponds to the separation between the barrel and endcap RPC trigger system, which simplifies the association made in the Global Muon Trigger between RPC muons and DT/CSC muons. The CSC Track-Finder requires at least one track segment in the CSC muon system, since otherwise the track should be found by the DT Track-Finder.

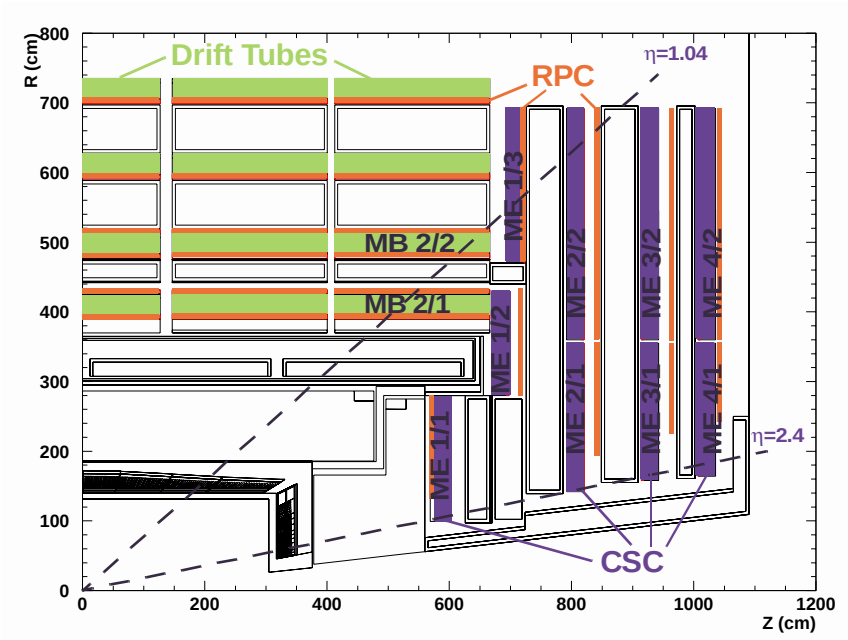

Fig. 12.2: Illustration of the CMS muon system showing the boundary division between the DT and CSC Track-Finders at $\eta = 1.04$.

To cover the region of overlap efficiently, information from MB2/1 chambers is shared with the CSC Track-Finder, and information from the ME1/3 chambers is shared with the DT Track-Finder.

12.2 System Overview

The CSC Track-Finder is defined to be the collection of electronic boards which are on the receiving end of the optical links sent by the CSC local trigger and that transmit L1 muons to the Global Muon Trigger (GMT). The CSC muon system is logically partitioned into 12 azimuthal

sectors (6 per endcap) for purposes of regional track-finding. Thus, 12 Sector Processors (SP) identify up to the three best muons in each 60° azimuthal sector. Each processor is a 9U VME card housed in a crate in the counting room of CMS. Three Sector Receiver (SR) cards per sector also reside in the crate to collect the optical signals sent from the Muon Port Cards of one sector, although technology may permit the SR functionality to be merged onto the SP board. A maximum of six track segments are delivered to a Sector Processor from the first muon station (ME1) of a sector. These track segments arrive from three Muon Port Cards, each delivering up to 2 track segments in a 20° subsector, as illustrated in Fig. 11.1. For the other muon stations (ME2-ME4), one Muon Port Card per station delivers 3 track segments. In addition, up to four track segments from the DT muon system (2 from each 30° subsector) are propagated to a transition board in the back of the crate and delivered to each Sector Processor as well. The output from all 12 Sector Processors is sent to a Muon Sorter (MS) which selects the 4 best muons out of 36 for transmission to the GMT. A block diagram of the CSC Track-Finder architecture is shown in Fig. 12.3.

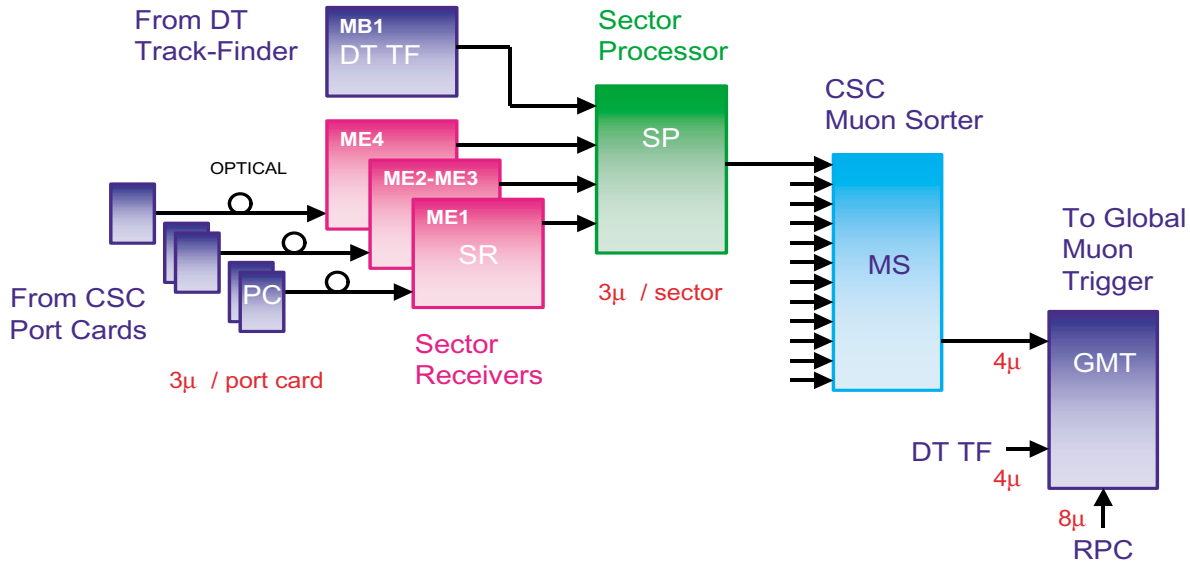

Fig. 12.3: Architecture of the CSC Track-Finder. The Sector Receiver functionality may be incorporated onto the same board as the Sector Processor.

12.3 System Interfaces

12.3.1 Crate and Backplane

The CSC Track-Finder is contained in several 9U VME crates. In the present prototype design, 2 Sector Processors, 6 Sector Receivers (3 per SP), and 1 Clock and Control Board occupy one crate; so 6 crates are needed for the entire CSC Track-Finder. The Muon Sorter is contained in an additional crate. VME addressing up to A24/D16 is carried over a standard VME 3U backplane. A custom 6U point-to-point backplane is used to carry all track segment data from the Sector

Receivers and DT Track-Finder to the Sector Processor. Channel-Link LVDS transmitters from National Semiconductor drive the data onto the custom backplane. The layout of one sector of one crate is shown in Fig. 12.4.

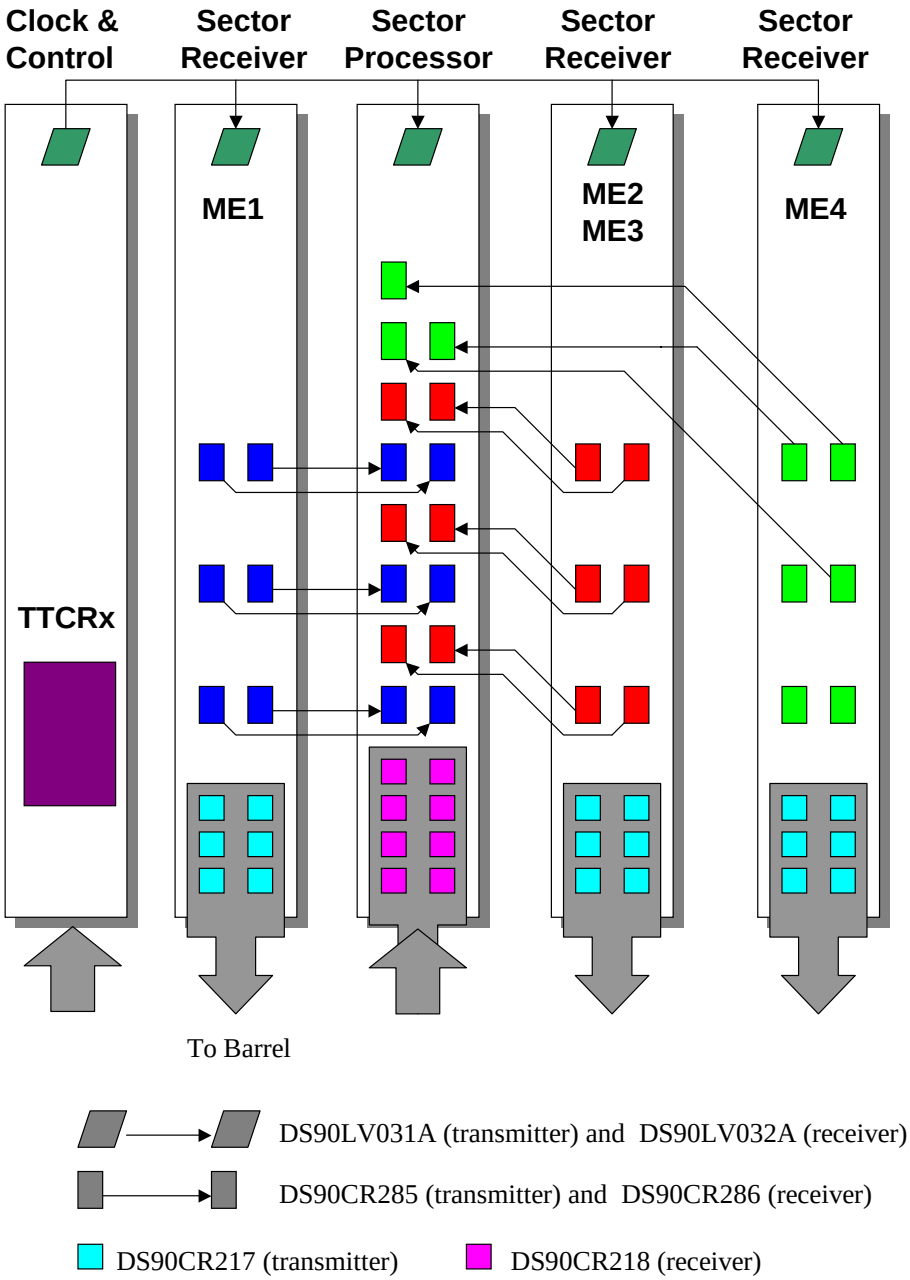

Fig. 12.4: Illustration of the card placement and backplane connections for one sector in the Track-Finder crate, according to current prototypes.

However, new optical link technology and recent high-density FPGAs may allow SR and SP functionality to be merged onto the same board, thus allowing the entire CSC Track-Finder

to be housed in one 9U VME crate along with the CSC Muon Sorter. The custom backplane in this case would deliver signals from the 12 SPs to the Muon Sorter.

12.3.2 Transition Modules

Transition Module to the DT Track-Finder

CSC track segments from the ME 1/3 chambers in the DT/CSC overlap region are sent from the ME1 Sector Receiver to the DT Track-Finder via a transition board on the back of the crate. One connection is needed to transmit two CSC track segments to one sector of the DT Track-Finder, and two such connections are provided from the ME1 Sector Receiver. The information sent is the quality and the ϕ coordinate of each track segment as well as the bunch crossing number (BXN), as listed in Table 12.1. Look-up tables on the Sector Receiver convert the track segment quantities into the format expected by the DT Track-Finder. The transmission technology currently proposed is Channel-Link LVDS. A connector on the backplane carries the signals to the transition board from the Sector Receiver.

Table 12.1: Information delivered from one SR to one DT Track-Finder sector.

Variable	Function	bits / muon	bits / 2 muons
ϕ	azimuth coordinate	12	24
quality	quality	3	6
BXN	LSBs of bunch i.d.	–	2

Transition Module from the DT Track-Finder

The DT track segments from the MB2/1 chambers in the DT/CSC overlap region are delivered to the Sector Processor via a transition board on the back of the crate. Connections from two sectors of the DT Track-Finder are required. The information received is the quality, ϕ coordinate, and ϕ_b bend angle of each track segment as well as the BXN, as listed in Table 12.2. Look-up tables on the DT Track-Finder convert the track segment quantities into the format expected by the Sector Processor. The chosen transmission technology is Channel-Link LVDS. A connector on the backplane carries the signals from the transition board to the Sector Processor.

12.3.3 Clock and Control Module

The TTC interface

The Clock and Control Board (CCB) provides the crate level interface to the TTC system. The TTC interface is based on the TTCrx chip [12.4]. The general sequence of L1A, Reset and BC0 commands is described in Chapter 16. Particularly, for the Reset sequence, the TTC sends a broadcast command (either system, or user) to the TTCrx indicating that the next L1A has to be treated as a Reset. After this data is sent, a single L1A is generated. CCB decoding logic recognizes this command and treats the next incoming L1A as a RESET signal. After some predetermined interval the next broadcast command is transmitted over the TTC indicating that the next L1A

Table 12.2: Information delivered from one DT Track-Finder sector to the SP

Variable	Function	bits / muon
ϕ	azimuth coordinate	12
ϕ_b	ϕ bend angle	5
quality	quality	3
BXN	LSBs of bunch i.d.	2
Synch./Calib.	Special Mode	1
Flag bit	Denote if 2nd muon	1

should be treated as Bunch Crossing 0. In a similar fashion, CCB recognizes this command and treats the next L1A as a BC0. After that CCB internal logic enables generation of L1A to backplane upon every L1A from TTC system. The TTCrx chip can be programmed using an I2C interface and interface controller PCF8584 over VME.

Clock adjustments and settings

There are fine and coarse delays for clock and command signals incorporated on the TTCrx chip [12.4]. We use the TTCrx Clock40Des1 deskewed signal as the main clock signal from the TTCrx board. Three other possible clock sources are ECL and NIM clocks from external connectors on the front panel, and clock from quartz oscillator.

The selected clock signal acts as a main master clock for the CCB internal logic. The phase of the clock signal provided to synchronization logic, can be adjusted with 2 ns step accuracy with respect to the main master clock. The phase of L1A, BC0, RESET, and two reserve signals (RSV1 and RSV2) distributed to all slots in a crate can be adjusted with 2 ns accuracy in respect to the main CCB master clock.

The selected clock is distributed to the SP and SR modules in the crate via its custom backplane. The phase of each clock signal distributed over this backplane can be adjusted with 2 ns step accuracy with respect to the main master clock individually to each slot in the crate. There is also the possibility to send just a single 25 ns clock pulse to all modules in crate upon special VME command.

12.3.4 Crate Power and Cooling

Crate power supply delivers ±12V, +5V, and +3.3V. All other voltages, such as 2.5V for high-density FPGAs, are obtained using DC-to-DC converters on the boards.

12.4 Sector Receiver

Each Sector Receiver (SR) receives via optical links the Local Charged Track (LCT) information for 3 muons from each of two Muon Port Cards (MPCs) located at the periphery of the CMS detector (except for the first station, where 2 muons from each of three Muon Port Cards are received). This information is then synchronized and reformatted within the SR (via look-up

tables) into angular variables for the muons: the azimuthal angle (ϕ), the local slope angle in ϕ (ϕ_b), and the rapidity (η). These data, along with bits summarizing muon quality and other diagnostic information are then communicated to the Sector Processor and the (barrel muon) DT Track finder in the (differing) format expected by each. Complete input information is also stored for readout by the DAQ system for accepted events. In addition, a VME-readable counter and log of detected hardware errors is included. The basic elements of the SR are shown in Fig. 12.5.

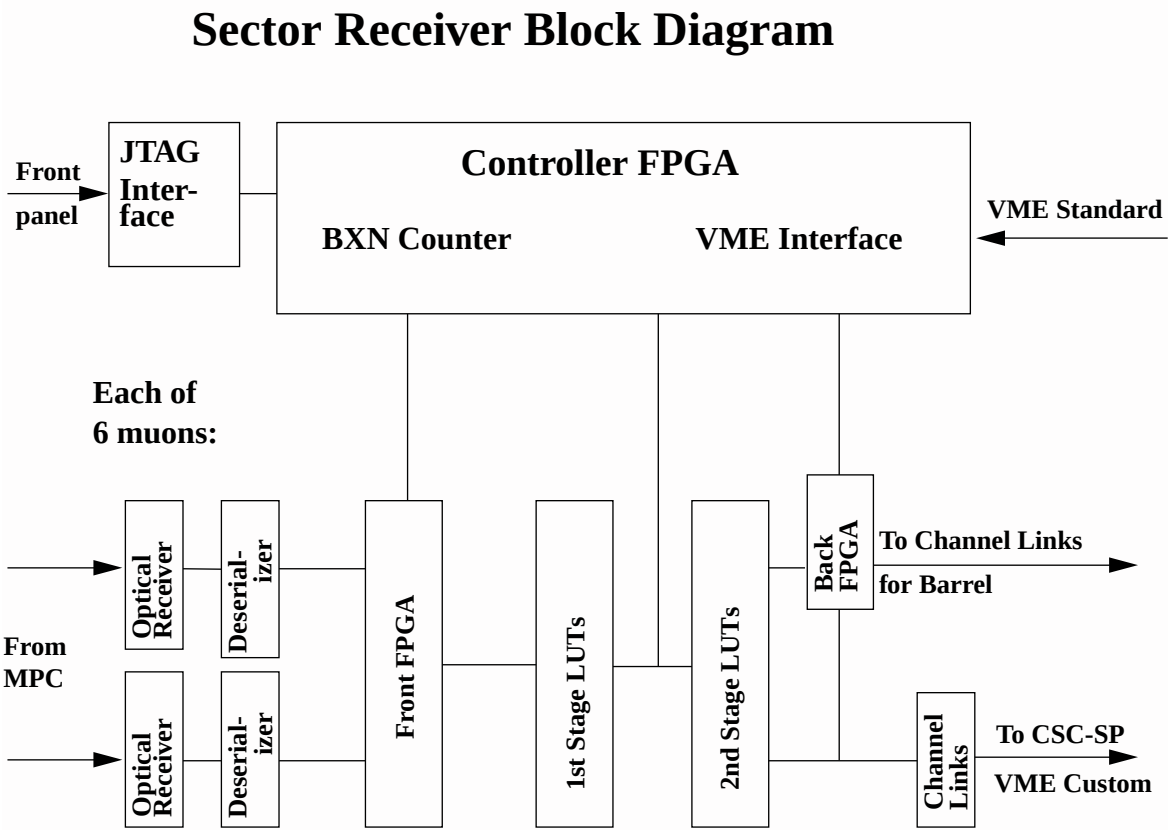

Fig. 12.5: Block diagram of the Sector Receiver logic

12.4.1 Optical link inputs

Each MPC transmits two or three muon LCTs to a SR. Each SR can receive data from two or three MPCs. For each 60^o sector in the current prototype architecture, there are:

1. A SR receiving data from the three MPCs from ME1
2. A SR servicing the MPC from each of ME2 and ME3
3. A half-used SR servicing the MPC from ME4, in the re-scoped scenario with the fourth CSC station

A single SR design has the flexibility to handle these cases. In total, there are 120 bits sent from each MPC to a SR.

The bits from the MPCs are sent via optical links. The present design uses the HP serializer/de-serializer chipset (HDMP 1022/1024) with Methode optical transceivers. The HDMPs operate in simplex mode with frames 24 bits long, with one frame transmitted per bunch crossing. (The HDMPs automatically divide the time between bunch crossings into 24 bits, each approximately 1 ns long.) Of these 24 bits, 3 are for defining the frame and one is a “flag bit” for checking transmitter/receiver synchronization, leaving 20 bits available for data. Thus, each MPC drives 6 links, organized as 3 pairs.

12.4.2 Backplane inputs.

The SR gets its clock and control signals from the CCB off the backplane, as described in Section 12.3.3. The SR maintains an internal bunch crossing counter which is incremented each clock cycle. With the arrival of a RESET signal, the counter halts and re-initializes to a preset value. It then starts counting upon receipt of the BC0 signal. The 5 least significant bits of this locally determined bunch crossing number are compared to the Anode BXN arriving from the MPC, and an error condition results if they are different. One of the two reserved signals will be used to specify a test mode.

12.4.3 Sector Receiver Outputs to Trigger Path

There are two separate streams of outputs from the SR, one to the SP via the backplane, and one to the DT Track-finder via transition modules and cables. These output streams are similar but are customized for each recipient. Table 12.3 lists the bits transmitted to the SP. Table 12.1 lists the bits transmitted to the DT Track-finder. The muon track variables ϕ, ϕ_b, and η are described further below. The accelerator muon bit is simply copied from the input from the MPC. The quality bits are computed from MPC inputs as described below. For each set of 3 muons coming from a MPC, the “CSC tag” encodes in 2 bits which pair of muons, if any, had a common CSC ID; this information helps resolve ambiguities in combining anode and cathode views. Finally, the Error bit is set if the data is invalid, for example if an optical link error has been detected.

12.4.4 Sector Receiver Outputs to DAQ Path

The SR will store all incoming data in a buffer for a time long enough (several μs) so that the DAQ system can read it for accepted events. The data will be transmitted from the SR to a Front-End Driver using an optical link.

Another buffer will contain a VME-readable error counter and error log, with bunch crossing number and error type for any errors found. Error conditions detectable by the SR include: the above-mentioned mismatch in BXN; a set error bit in data from the MPC, indicating error further upstream; and error bit provided by the HDMP chipset, indicating loss of synchronization in the optical link.

12.4.5 Hardware Implementation

Optical receiving, de-serialization, synchronization

As discussed above, the data from the MPCs arrive on optical cables, and is de-serialized, forming a 120-bit-wide data set arriving from each MPC each beam crossing. The data

Table 12.3: Information delivered from one SR to the SP

Variable	Function	bits / muon	bits / 6 muons
ϕ	azimuth coordinate	12	72
ϕ_b	ϕ bend angle	5	30
η	pseudo-rapidity	6	36
Accelerator μ	η bend angle	1	6
quality	–	3	18
CSC ghost	2 hits in same CSC	–	4
Error	Data not valid	–	1

are latched into the Front FPGA with clocks derived from the CCB-derived 40-MHz clock. Data arriving on separate links can be re-synchronized by programmable delays within the FPGA.

Front FPGA

For each muon, the Front FPGA receives the data from the MPCs. It is connected to all address lines of the muon's first stage of memory lookup tables, to the DAQ output path, and to the SRs VME interface. Thus, there is some room for flexibility and evolution in the design presented here, if design requirements evolve. The Front FPGA keeps a copy of the incoming data in a buffer several μs deep for readout via DAQ. For diagnostic purposes, the Front FPGA can insert dummy data into the Sector Receiver, simulating data arriving from the MPC, either in single-event-mode, or burst-mode at 40-MHz for 256 beam crossings. Finally, the Front FPGA contains logic for downloading the LUT contents via the VME interface.

Memory Look-up Tables

The LUT functions are shown schematically in Fig. 12.6. In the current prototype, 6 identical 256K by 16 bit memories are used for each of the six muons. In the first stage, the cathode LCT information is translated into local or approximate values of ϕ and ϕ_b, and anode LCT information is translated into the approximate η. In the second stage, each coordinate is corrected using information from other coordinates. The second stage also corrects for any alignment problems as well as slanted anode wires. The details of these two stages are as follows. In addition, a sixth LUT is used to compute quality bits, as described in the next subsection.

The 8-bit "1/2 strip ID" corresponds to a coarse ϕ position at layer 3 within a particular "station" (ME1–ME4) of 6 layers. For a track traversing this station at an angle, the true ϕ depends on the depth within the 6 layers. Furthermore, at each station, chambers alternating in ϕ are overlapped, with "front" chambers closer to the interaction region than "back" chambers. The Sector Processor can function most effectively if the ϕ passed to it is at a conventional distance from the interaction region. In the first LUT, such a ϕ is computed, using in addition: the 8-bit "CLCT pattern" which encodes which strip pattern was recorded in the 6 layers (and which also carries information for decoding the "1/2 strip ID"); the "L/R Bend" bit, and a "Front/Rear chamber bit" which is derived in the Front FPGA from the 4 CSC ID bits. The resulting "local ϕ" is 10 bits, in half-strip units, spanning only that particular chamber. From the same LUT also comes

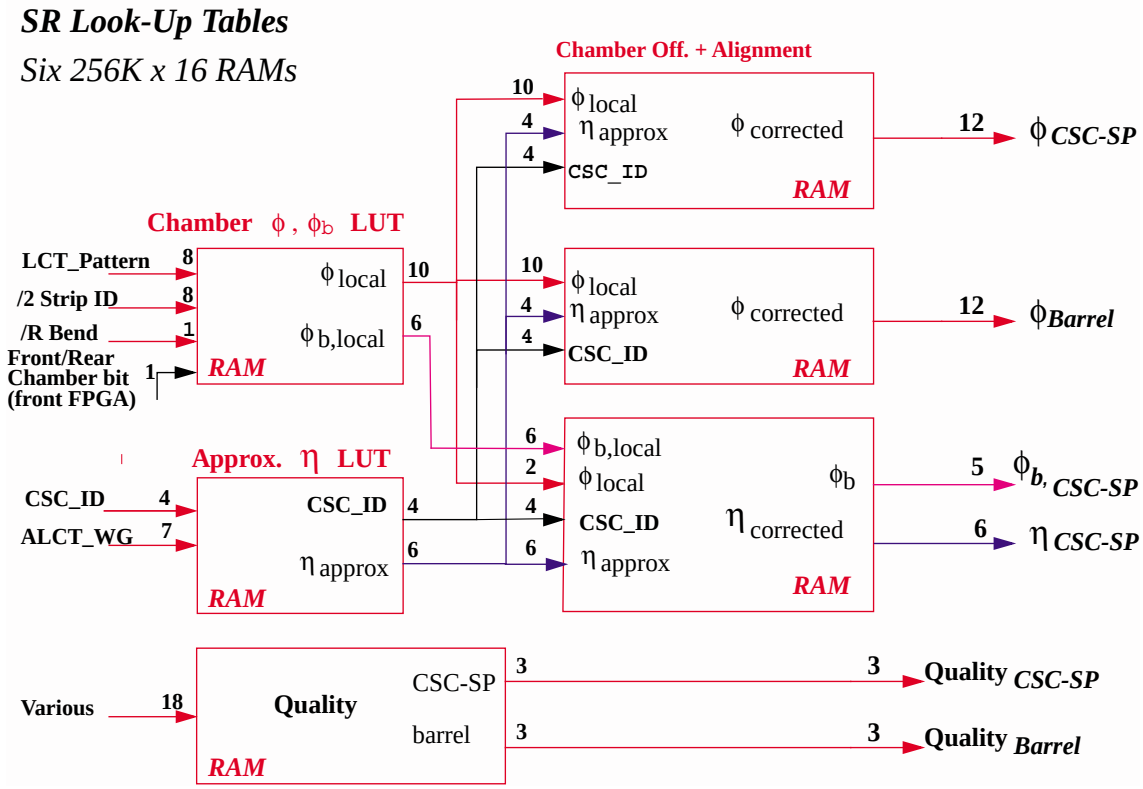

Fig. 12.6: Sector Receiver Memory Lookup Connections

a raw ϕ_b, which with 6 bits encodes the change in ϕ between layer 1 and layer 6, in fractional strip units, inferred from the CLCT pattern. (Although there are 256 possible CLCT patterns, many of them correspond to identical values of ϕ_b.)

In the other first-stage LUT, the approximate 6-bit η is computed from the 7-bit ALCT wire group and the 4-bit CSC ID (which is necessary since wire grouping depends on the chamber). Unlike the first-stage local ϕ, this η is already global, although still requiring correction in the second stage. In this LUT, we also use some of the extra output bits to pass on an in-time copy of the 4 CSC ID bits to the second stage LUTs.

In the second-stage LUTs, four independent quantities are calculated:

i) For the SP, a 12-bit global ϕ is computed from the 10-bit local ϕ using the CSC ID, with alignment corrections also possible using the most significant 4 bits of η from the first stage.
ii) For the DT Track-finder, a similar 12-bit global ϕ is computed independently in order to provide for coordinate conventions different from the SP.
iii) For the SP, 5-bit corrected ϕ_b in radians is computed from the raw value in fractional strip units. This requires η since the strip width depends on the position along the strip. The CSC ID is also input to this LUT in case there are chamber-dependent corrections.
iv) For the SP, the 6-bit corrected η is computed from the first-stage η. This is corrected using the 2 most significant bits of local ϕ from the first stage, primarily to account for slanted

anode wires it ME1/1. The (essentially negligible) ϕ-dependence of η along other anode wires can also be corrected for, at no cost.

Items i) and ii) each have a dedicated LUT, while items iii) and iv) are combined into one LUT. The memories used in the working prototype are described in Section 12.10.1.

Computation of quality bits

For each muon, 3 quality bits are passed to both the DT Track-finder and the SP. They are functions of the following inputs to the SR: the Valid Pattern flag, the cathode pattern number, the anode pattern quality, the ALCT/CLCT BXN match, all 4 TMB status bits, the incoming error bit, as well as SRs internally generated error flags. Flexibility is needed, since the preferred function may change with operating conditions and as understanding of the trigger performance increases. The simple functions now envisioned could probably be implemented in the Front FPGA. However, in order to maintain maximum flexibility and eliminate potential latency problems with this calculation, we are dedicating the sixth LUT to it for each muon.

Back FPGA

After the second LUT stage, the data go directly to both the Back FPGA and the Channel-Link chips for transmission to the SP on the backplane (or to SP functions on the same board). Unlike the data to the SP, the data to the DT Track-Finder must go through logic in the Back FPGA in order to select 2 of 3 muons.

Since the Back FPGA is connected to all output signals as well as the VME interface, there is flexibility for various diagnostic tests, all run at 40 MHz. In a typical test, the Back FPGA records the data from 256 events as it passes by on the way to the Sector Processor. For more specialized tests of the SR-SP interface, 256 events can be loaded into the Back FPGA buffer and then clocked out to the SP. These capabilities facilitated debugging of both the SR and the SP. As with the Front FPGA, the Back FPGA is used to load and read the memory LUT contents.

Computation of CSC ghost bits for SP

It is possible that two track segments delivered by a single MPC may have come from the same CSC chamber, in which case there is an ambiguity in the association of the anode and cathode LCTs. The SR can compare the CSC IDs of each track segment coming from a MPC and set a “CSC ghost” flag for the SP, which in turn will try all η, ϕ combinations for this pair of track segments to resolve the ghosts in the track-finding. For example, for stations ME2-ME4, the three muons (designated A, B, and C) in the top half of one SR all come from one MPC. Each has its own CSC ID, and at most two of them can have the same ID. To assist the SP in resolving the ambiguities, 2 “CSC ghost” bits (see Table 12.3) are sent to the SP with the following binary code: 00 means all three IDs are different, 01 means A=B, 10 means A=C, and 11 means B=C. Calculation of the code is in one of the FPGAs for that set of three muons. Similarly, a separate 2-bit code is computed for the three muons in the bottom half of the SR. The case for ME1 is slightly different since 3 MPCs each deliver two track segments to one SR. In that case, only 3 bits are needed, where each bit denotes whether the two track segments from one MPC are from the same chamber.

Reduction from three to two muons for the DT

As agreed on with the DT group, in the endcap-barrel overlap region, the DT Track-Finder will use muons only from (part of) chamber ME1/3. The MPC servicing this chamber provides up to three muons to the SR, including muons from ME1/1 and ME1/2. The DT Track-Finder can accept at most two muons. Therefore, special logic on the SR is required to select (at most) two muons relevant to the DT Track-Finder from the muons coming from the MPC, transfer those muons, and mask out the rest. The logic for this reduction is in a Back FPGA.

VME and JTAG Interfaces

The VME interface is in a FPGA near the backplane. This FPGA also contains miscellaneous logic such as the bunch crossing counter. It drives a JTAG controller (design copied from the SP) for servicing the Front and Back FPGAs. There is also a front-panel JTAG connector for bench tests; this will not be used in the working experiment. Currently we use it to load the EEPROMs at the top of the board; the FPGAs themselves are then loaded from the EEPROMs.

Operational Modes

The Sector Receiver has VME addressable registers to define various modes for reading and writing to memory LUTs, for normal operation, and for test modes. The number of clock cycles per start command and various clock delays are all VME-programmable. Depending on the mode, access to the memory LUT address and data lines is reconfigured using tri-state and bi-directional buffers. The desired interactions among the firmware of the three functional types of FPGAs have been successfully demonstrated in the prototypes.

12.5 Sector Processor

12.5.1 Overview

The Sector Processor reconstructs tracks from the track segments delivered by the Sector Receivers and the DT Track-Finder. The number of CSC track segments collected by one Sector Processor is 15 per bunch crossing, assuming that ME4 participates. Six track segments are delivered from ME1; three each are delivered from ME2–ME4. Additionally, 4 DT track segments are delivered from the MB 2/1 chambers in the outer wheel of the barrel muon system.

A description of the algorithms for the Sector Processor can be found in Refs. [12.5] and [12.6]. 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, the track segments from the CSC and DT trigger systems must be synchronized and possibly held for more than one bunch crossing to accommodate bunch-crossing misidentification from the LCT and BTI processors. Next, 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. If an ambiguity is created when two muons enter the same CSC chamber, all possible η, ϕ combinations are tried. The process is not complete, however, since the Track-Finder must report the number of *distinct* muons to the L1 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 Sector Processor is illustrated in Fig. 12.7. Each of the important blocks is described in detail below.

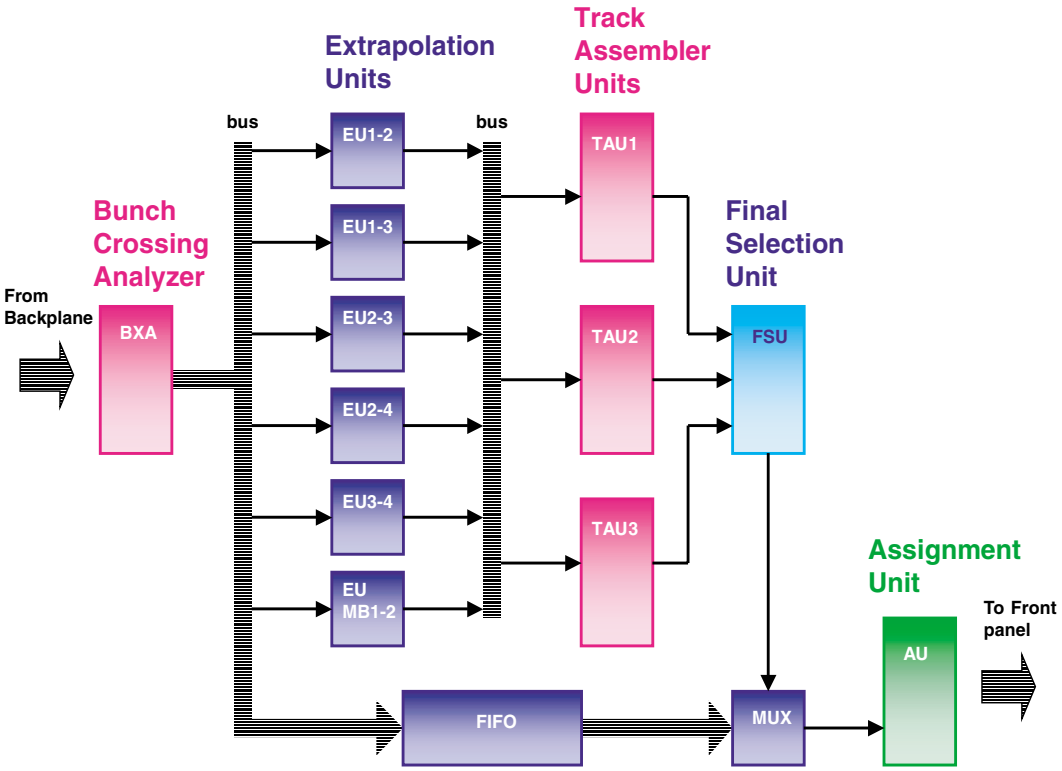

Fig. 12.7: Block diagram of the Sector Processor logic

12.5.2 Bunch Crossing Analyzer

The input data to the Sector Processor from the DT and CSC trigger systems is synchronized to the local clock before being sent to the Extrapolation Units. A provision was made in the design include some ability to analyze track segments received in out-of-time bunch crossings for several reasons:

- The DT Track-Finder sends two track segments from one chamber over consecutive bunch crossings
- The bunch crossing assignment of the DT and CSC local triggers is not 100% accurate, although it is nearly so for the CSC system

– It will be easier to commission the system when cable delays are not exactly known

To incorporate a multi-bunch mode, we take advantage of the sparseness of the data. If the data is not sparse, CSC track segments would be lost already at the Muon Port Card, which selects only the three best track segments from 9 chambers. 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.

The window over which track segments are collected is at least two bunch crossings wide. 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 N crossings based on the track segment quality and on the deviation from the current crossing. The same can be done for the best track segments from MB2/1. 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, the Muon Port Card sends the best three track segments in ranked order.

This scheme is shown in Fig. 12.8 for 3 track segments as input. 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. A flag is set to record whether a track segment comes from the current bunch crossing or a different one. This flag will be used in the Final Selection Unit of the Sector Processor to determine if a trigger should be inhibited so that the Sector Processor does not generate extra triggers over several bx.

12.5.3 Extrapolation Unit

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. All possible extrapolation pairs should be tested in parallel to minimize the trigger latency. However, we have excluded direct extrapolations from ME1 to ME4 in order to reduce the number of combinations and to reduce some random coincidences (since those chambers are expected to have the highest rates). The exclusion 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 FPGAs. A logic diagram for the extrapolation of one track segment from station A to another in station B is shown in Fig. 12.9. The flip-flops for data pipelining are not shown. The extrapolation unit is composed of several sub-units which analyze the η coordinates of the two track segments from different stations, the ϕ coordinates, and the quality of the resulting extrapolation. These sub-units are described below.

Eta Road-Finder:

The tests involving the η information from the two track segments are the following:

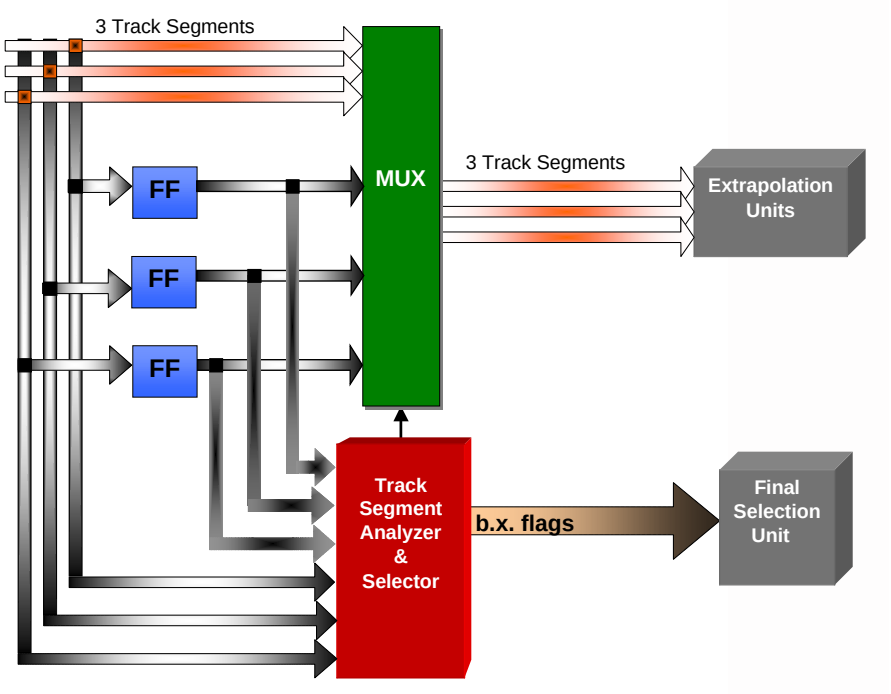

Fig. 12.8: Block diagram of the Bunch Crossing Analyzer

1. Determine if each track segment is in the allowed trigger region in η
2. Compare the η values of the two track segments to determine if both lie along a straight line projection to the collision vertex within a certain tolerance
3. Check that the bend angle in η for at least one track segment is consistent with a track originating from the collision vertex. Presently, only one bit (the Accelerator Muon bit) is used to flag if a track segment is parallel to the beam axis rather than projective.

The “AND” of all 3 conditions results in one bit which is sent to all other extrapolation units involving the same pair of stations. In the event that two track segments come from the same CSC chamber, there is an ambiguity in the association of the η and ϕ hits which gives rise to ghost hits. The Sector Processor can test all possible combinations by swapping the η coordinates of two track segments which come from the same CSC chamber. This is accomplished by sharing the result of the η tests with other extrapolation units. The overall output of an η unit, then, is the “OR” of its own test with the result from another η unit with a different track segment in the same source chamber (but the same track segment in the target station). This CSC ghosts handling is only foreseen for ME1 currently.

Those extrapolation units that test track segments from the DT muon system have modified conditions for the η unit because no η information is sent from the DT trigger system. In general, the conditions listed here apply only to the CSC track segment for tracks in the overlap region. The value of η from the CSC track segment is used for further tests in the ϕ road-finder.

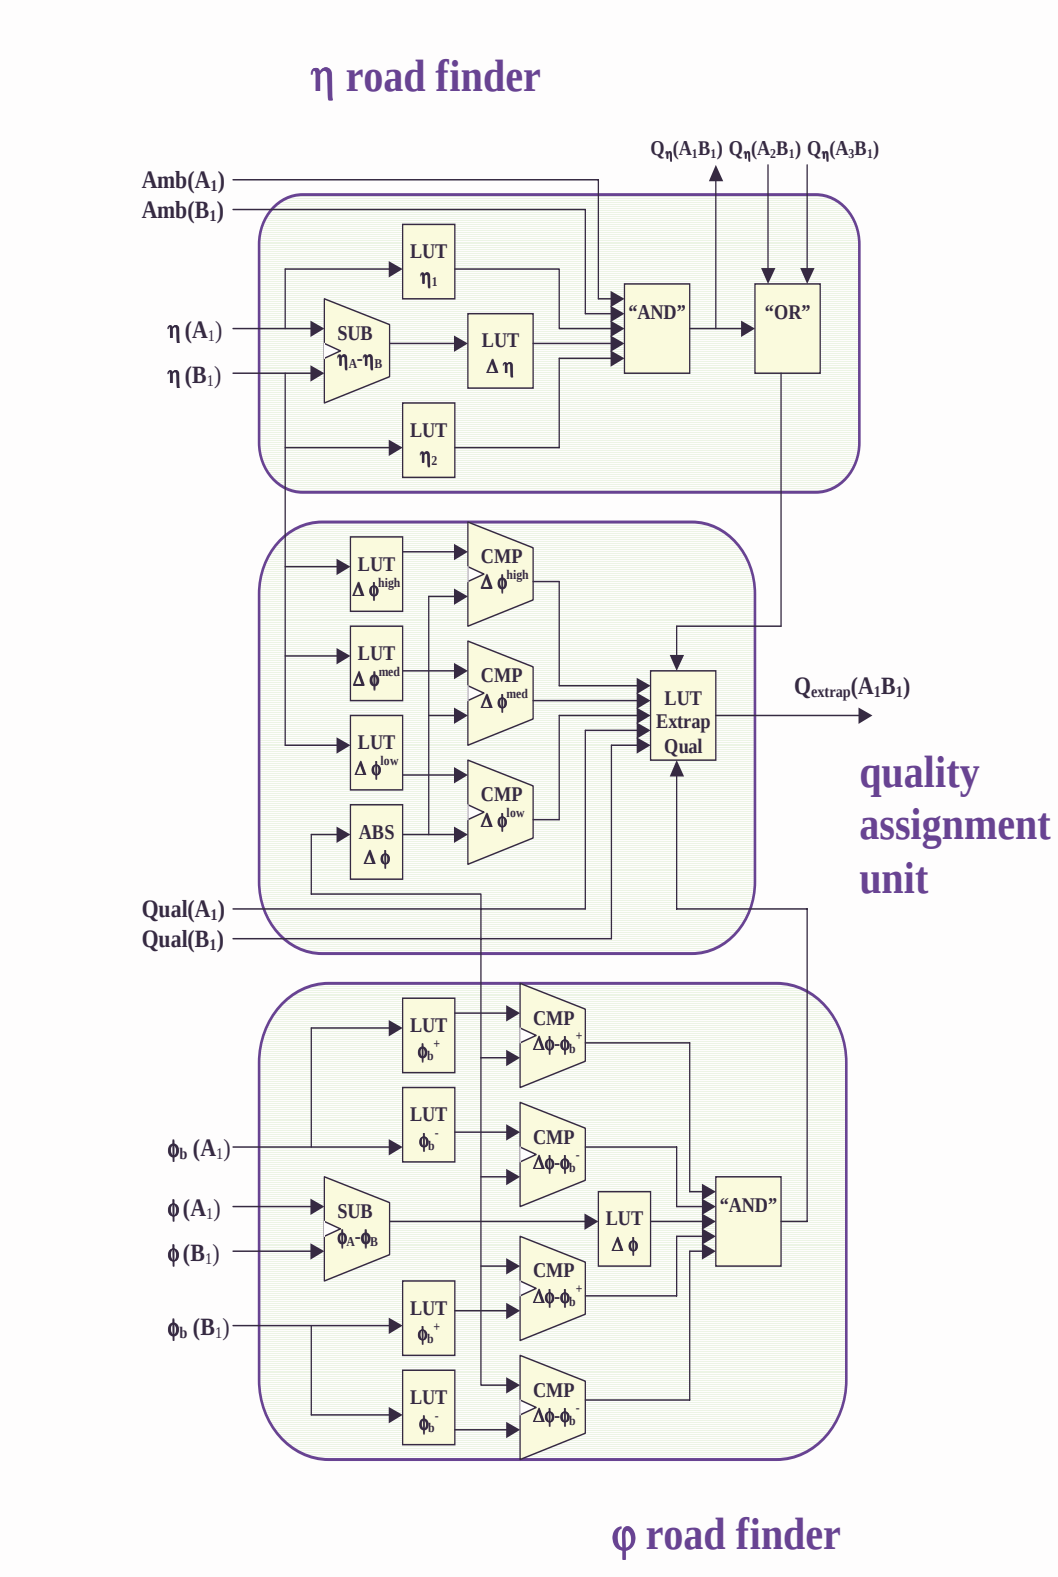

Fig. 12.9: Block diagram of the extrapolation unit logic, which compares a track segment in one station (A_1) with that in another (B_1).

Phi Road-Finder:

The tests involving the ϕ information from the two track segments are the following:

1. Compute the difference in ϕ between the two track segments
2. Check that the difference in ϕ is consistent with the bend angles in ϕ_b measured at each station
3. Compare the difference in ϕ to the maximum allowed at that η. Several thresholds may be employed to provide a coarse p_T measurement.

Quality Assignment Unit:

The final quality assignment for the extrapolation is based on the bits from the η and ϕ road-finder units as well as the track segment quality bits. It is generated by a small look-up table. 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 (ME1 or MB1). Otherwise, the quality just represents whether the extrapolation was successful or not. The expected p_T resolution for a ϕ resolution of 10 bits is about 30% when ME1 is involved. The quality word is used later when muon candidates are sorted.

12.5.4 Track Assembly Unit

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 can be separated into three data streams: one for patterns keying off ME3, one for patterns keying off ME2 in the endcap region, and one for patterns keying off ME2 in the DT/CSC overlap region. This is illustrated in Fig. 12.10. Only ME2 is used as a key station in the overlap region, since ME3 has no coverage there and ME1 has too many track segments. Some muons will be found by more than one stream, 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. 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, 9 total for all three 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. 12.11. The latency, therefore, is just one beam crossing. The output

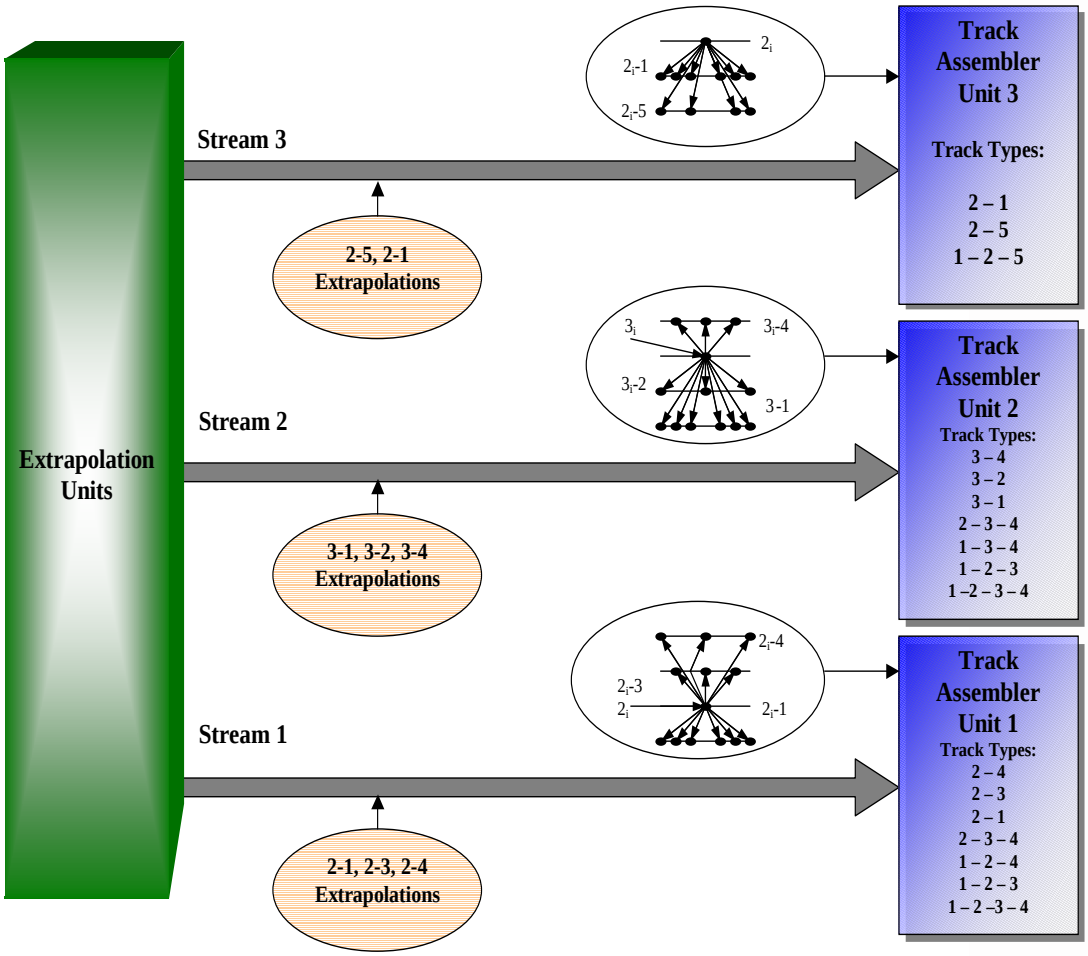

Fig. 12.10: Illustration of the track assembly procedure separated into three data streams.

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. It is possible, however, that a reasonable latency also can be achieved using FPGA track-assembly logic, so this option is kept open as well.

12.5.5 Final Selection Unit

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 CSC 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

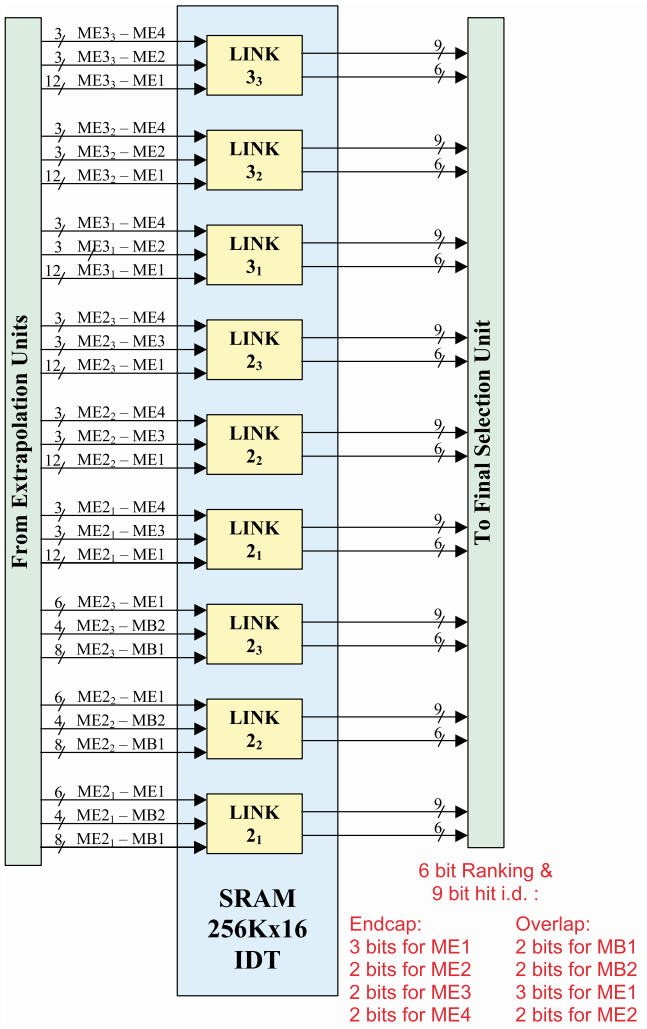

Fig. 12.11: The Track Assembler Unit implemented as 9 static RAM memories for the endcap and overlap region

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 Selection Unit is a sorter with cancellation logic. It sorts and cancels 9 tracks down to 3 since there are two endcap data streams and one overlap data stream.

A block diagram of the Final Selection Unit is shown in Fig. 12.12. 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 Section 12.5.2.

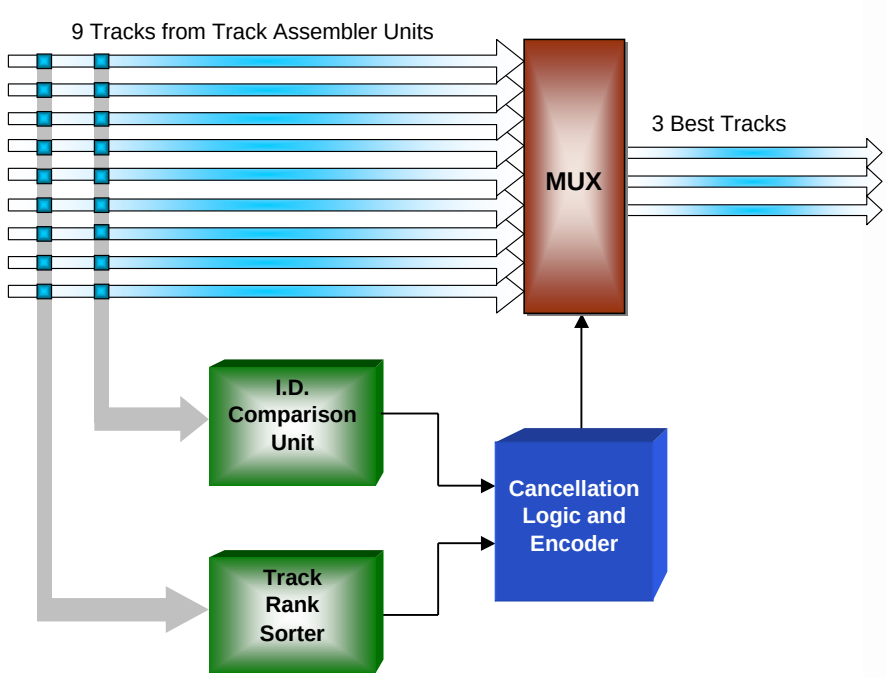

Fig. 12.12: Block diagram of the Final Selection Unit

12.5.6 Assignment Unit

The Sector Processor measures the momentum of the identified muons in the final stage of processing. This 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 format of the data is specified in Table 12.4. In particular, p_T and the track quality are combined into an overall rank before transmission to the Muon Sorter. 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.

Table 12.4: Information delivered from one SP to the Muon Sorter

Variable	unit / precision	range	bits / muon	bits / 3 muons
ϕ	2.5°	0–60°	5	15
η	0.075 η unit	0.9–2.4	5	15
Rank	p_T (nonlinear) and Quality	2–140 GeV/*c*	5+2	21
Sign	Muon sign	–	1	3
BXN	–	–	–	4
Error	–	–	–	1

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 [12.1] 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. (The best momentum resolution possible from an offline standalone muon measurement in the endcap is 15%, from Ref. [12.7].) 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 megabyte-sized memory for assignment of the track rank (p_T and quality) and sign. Tracks composed of only two track segments are allowed also in certain cases. This scheme is illustrated in Table 12.13 for the parameter assignment of one muon. Three such units are necessary for the three best muons selected by the Final Selection Unit. Since the first stage of the parameter assignment is done in an FPGA, additional logic may be added to cancel certain track classes when they occur near sector boundaries. This may help to reduce fake di-muon triggers.

The 2-bit quality assigned to a muon reflects the uncertainty in the p_T assignment. Specifically, the highest quality is assigned to tracks that have segments in 3 or 4 CSC stations, including ME1, since the best resolution is possible. Medium quality is assigned to tracks that have

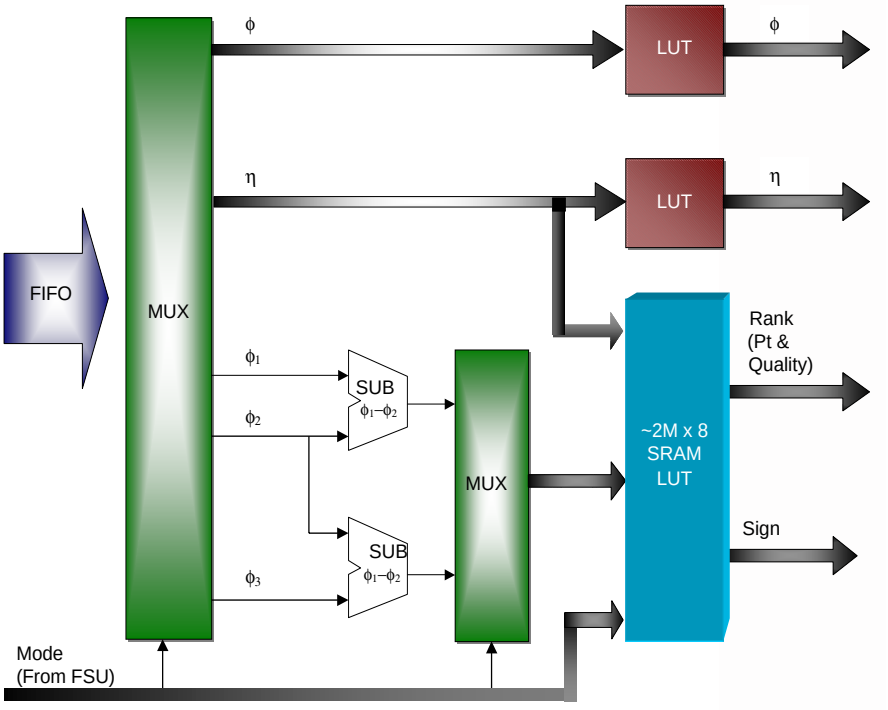

Fig. 12.13: Block diagram of the Assignment Unit for one muon.

segments in only two stations, of which ME1 must be one. Finally, lowest quality is assigned to all tracks that do not include ME1, since only a very poor p_T resolution is possible.

12.5.7 Hardware Implementation

The Sector Processor logic should be fully programmable, so FPGAs and SRAM should be used. Microprocessors and DSPs are too slow for L1. The hardest challenge for the hardware implementation is the design of the Extrapolation Units, which have a high I/O count and a large amount of logic. It is expected that the Extrapolation Unit logic will be implemented in high-density FPGAs such as the Virtex family from Xilinx. The Bunch Crossing Analyzer and global FIFO are expected to be implemented in more moderately-sized FPGAs. The Track Assembler Units, and the Assignment Unit, will be implemented in SRAM memory in conjunction with FPGA logic.

Approximately 500 signals will be received by one Sector Processor every bunch crossing. Thus, both connector space and board routing will be challenge. The Sector Processor also contains several high-density FPGAs with a very large pin-count (presumably in a ball-grid array), which further complicates the board routing.

12.6 CSC Muon Sorter

A total of 12 SP are needed for both endcap regions and the DT/CSC overlap region. Each SP outputs three muons, but only the four best muons from the CSC chambers are to be

reported to the Global Muon Trigger (GMT), in ranked order. Thus the purpose of the CSC Muon Sorter is to select four best muons out of up to 36 muons coming from 12 SP, and transmit them to the GMT every 25 ns. The format of the data sent to the GMT is specified in Table 12.5.

Table 12.5: Information delivered from the CSC Muon Sorter to the GMT

Variable	unit / precision	range	bits / muon	bits / 4 muons
ϕ	2.5°	0–360°	8	32
η	0.075 η unit	0.9–2.4	6	24
p_T	non-linear	2–140 GeV/*c*	5	20
Quality	–	–	3	12
Sign	Muon sign	–	1	4
BXN	–	–	–	4
Error	–	–	–	1

12.6.1 Algorithm

The sorting is based on a 7-bit rank, which is provided by the SP. Higher ranks (i.e. larger 7-bit rank patterns) correspond to “better” muons for the purposes of sorting, so the MS selects the four muons with the largest rank and outputs them in descending order. The best muon should always be present on the first link to the GMT, the second best muon – on the second link to the GMT and so on. The rest of the bits belonging to each incoming muon are stored in pipeline logic until the sorting result is obtained.

Given the different demands upon the Muon Sorter, unlike the RPC project [4], we have chosen to implement the sorter algorithm in PLD logic and not an ASIC. Such a solution can provide a lot of flexibility and can be done faster than the current RPC ASIC. Also we can benefit from the rapid growth in PLD/FPGA technology. Our first sorter implementation is targeted to the 20KE Altera PLDs, the fastest available Altera PLD family. We also concentrate on a single chip solution for the sorting logic, which would provide the minimal latency and optimal board design.

All sorting schemes are based on multiple comparisons and data multiplexing. Different design approaches and schemes require different number of comparison steps and number of comparisons at each step. Our main goal is to reduce the latency of sorting. We assume that latency is the time interval between the latching of input patterns into sorter chip and moment when the addresses of selected patterns are available for latching at the external logic outside sorter chip.

The block diagram of the sorter PLD which was designed and simulated is shown in Table 12.14. Two sorting steps “4 out of 18” are realized in parallel at the beginning of sorting tree, and then one sorting step “4 out of 8”. In case of (n) input patterns the total number of comparisons between all patterns is $N=n(n-1)/2$. If n=36, then N=630, and if n=18, then N=153. The first step of our scheme requires 18×17=306 comparisons, the second one – only 28.

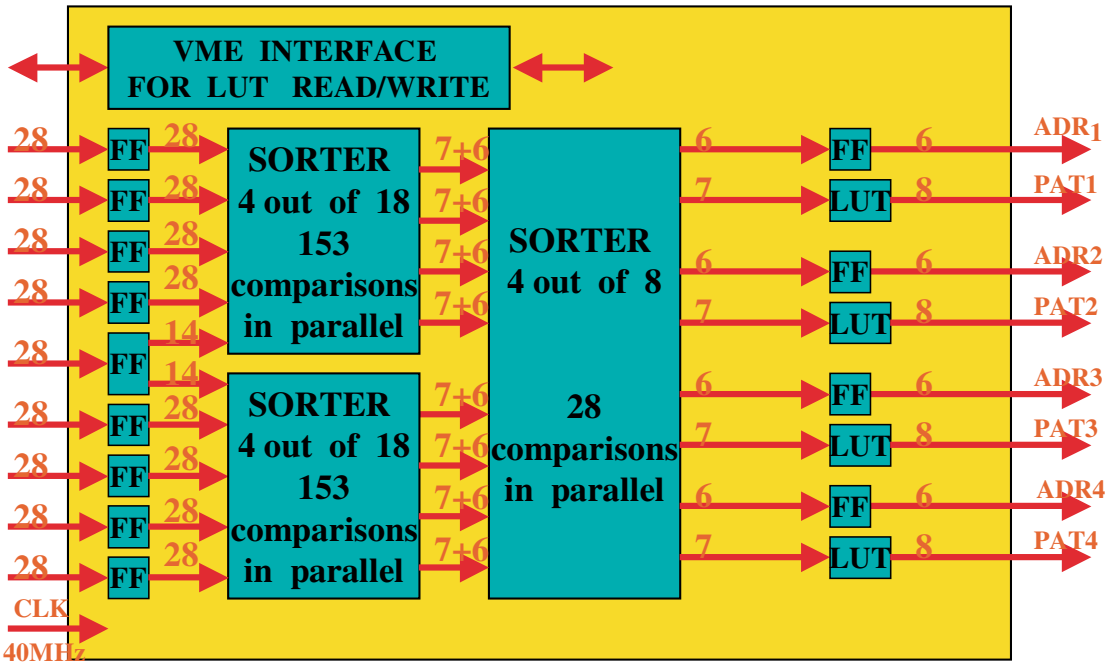

Fig. 12.14: Block diagram of the sorter PLD. It receives 7 bits of rank for each of 36 muons found by the Sector Processors and selects the 4 highest rank. "FF" denotes a flip-flop, "LUT" denotes a look-up table RAM.

The sorting PLD contains input, output, and intermediate (not shown on Fig. 12.14) flip-flops (FF) for proper data pipelining in order to provide a synchronous operation at 40 MHz. Our initial single-chip design is based on Altera EPF20K200EFC484-1 PLD. Sorting latency is four clock cycles, or 100 ns. The sorting PLD outputs the 6-bit addresses of the first, second, third and fourth best muons. These addresses enable multiplexing of the pipelined muons to the sorter board outputs. One clock cycle later, sorting logic outputs four 8-bit patterns (5-bit p_T + 3-bit Quality) which correspond to the selected muons. So the total number of sorter logic outputs is [6-bit address (ADR) + 8-bit pattern (PAT)] × 4 = 56.

12.6.2 Hardware Implementation

As discussed in previous sections, the MS will accept trigger data from 12 separate Sector Processors. To match the current Sector Processor prototype, it should contain 12 input connectors and receive 60×12=720 input signals using 12×4=48 16-bit input LVDS receivers. It should also contain four connectors and eight parallel LVDS transmitters to the GMT. Due to large number of inputs and outputs, a single chip solution for the whole sorter board is not feasible at the moment. We propose to use several PLDs for the sorting and pipelining of 36 muons.We intend to implement the sorting logic, the interface to the GMT, the VME control logic, and the interface to the Clock and Control Board (CCB) on a single 9U × 400 mm board. This board would need to carry four stacked mezzanine receiver boards, each of them consisting of the connectors, interface,

and pipeline logic, for communication with three SP of the current prototype design. This leads to a 5 board Muon Sorter which we intend to locate in a separate 9U crate in the counting room.

However, if the Sector Receiver and the Sector Processor are combined onto one board using improved technology, it is possible to fit the entire CSC Track-Finder into one VME crate, including the Muon Sorter. In this case, a custom backplane will deliver the signals from the 12 Sector Processors to the Muon Sorter.

12.7 Synchronization and Latency

12.7.1 Synchronization Procedure

The Anode LCT is used to synchronize the trigger system. The Anode LCT can identify the correct bunch crossing with greater than 99% efficiency. Thus the BXN generated by the ALCT will be histogrammed and compared to the bunch crossing structure of the LHC beam. By using the repeating nature of the bunch structure it is estimated that a determination can be made in 25 minutes of running at $10^{32}cm^{-2}s^{-1}$. Once that has been determined, each subsequent board in the chain is counting the BXN for itself. By comparing with the BXNs from the prior boards it will be possible to determine the offsets for each board in the system.

12.7.2 Latency Determination

The estimated latency of the CSC Track-Finder is 26.5 bx, from the time data is available at the end of the optical fiber in the counting room (51.5 bx after the collision) until it is delivered to the Global Muon Trigger crate. Thus, the CSC trigger data is available at the Global Muon Trigger 78 bx after the collision. The accounting of this latency is shown in Table 12.6.

Table 12.6: Latency of the CSC Track-Finder

Description	bx this step	Total bx
Delivery of optical signals to CSC Track-Finder	–	51.5
SR optical receiving and synchronization	2	53.5
SR Processing and transmission to SP	3.5	57
SP processing	11	68
SP to Muon Sorter transmission over 5m cable	2.5	70.5
Muon Sorter processing	4	74.5
Muon Sorter to GMT transmission over 11m cable	3.5	78

12.8 System Monitoring

The Sector Receiver and Sector Processor will have VME-readable registers to log errors and other diagnostic information so that the system can be monitored through the VME bus. For example, if any board detects that a synchronization error has occurred, all subsequent trigger data will be flagged as “data not valid” and the error and bunch crossing will be logged in the VME registers.

The trigger data received and generated by the CSC Track-Finder for each event will be stored in FIFOs on the Sector Receiver and Sector Processor for subsequent readout by the DAQ if an L1A is issued. A dedicated link to a Front-End Driver of the DAQ is anticipated, since VME is too slow for the event readout. This trigger data will be used to monitor the performance and efficiency of the CSC Track-Finder from the offline data stream.

12.9 Simulation Results

The performance of the CSC Track-Finder has been simulated using the GEANT-based CMSIM and ORCA software packages. Most results are obtained using the ORCA4 framework, where GEANT hits were obtained using version 118 of CMSIM; however, some earlier studies (see Ref. [12.1]) of the p_T resolution reported here were obtained using version 114 of CMSIM and a previous LCT simulation. Efficiency studies based on ORCA were carried out with a sample of positive and negatively-charged single muons generated flat in the ϕ coordinate ($0 < \phi < 2\pi$ rad), flat in pseudo-rapidity ($|\eta| < 2.4$), and flat in p_T ($5 < p_T < 100$ GeV/c). The trigger rate studies have been performed on the minimum bias samples described in Chapter 8.

An object-oriented description of the CSC Track-Finder has been written to exactly mimic the functionality of the current prototypes. In fact, this code has been used to validate the hardware, where perfect agreement has been achieved for 200,000 single muon events.

The p_T measurement in the Sector Processor is based on the sagitta of the track induced by the magnetic bending. The sagitta is determined from the difference in the azimuthal angle ϕ of the track segments measured in different CSC stations. It is possible to determine the p_T of a track from the difference in ϕ values measured in any pair of CSC stations at a given pseudorapidity. However, a more precise p_T measurement can be obtained if more of the track's sagitta is used; *i.e.* using track segments measured in three CSC stations. The techniques for the two-station and three-station p_T measurements used by the CSC Track-Finder are described in Ref. [12.1].

Figure 12.15 shows the resolution of p_T for the 2-station measurement as a function of η for several p_T values as determined in the earlier CMSIM 114 studies of Ref. [12.1]. The p_T was reconstructed from $\Delta\phi$ measured between MB2/1 and ME1/3 for the overlap region, and between ME1 and ME2 in the endcap region. The resolution is defined to be the width of the the residual distribution $(1/p_T^{meas} - 1/p_T^{true}) \times p_T^{true}$, where the measured p_T is reported at 50% efficiency. The resolution improves with decreasing p_T, and tends to stay at about 30% at low p_T. At high p_T the resolution is dominated by the width of the cathode strip, whereas at low p_T the resolution is dominated by multiple scattering. The resolution gets worse (~70% at low p_T) if MB1 is excluded in the overlap region or ME1 is excluded in the endcap region.

Figure 12.16 shows the resolution of the p_T for the three-station measurement (ME1–ME2–ME3) as a function of η compared to the two-station (ME1–ME2) measurement at

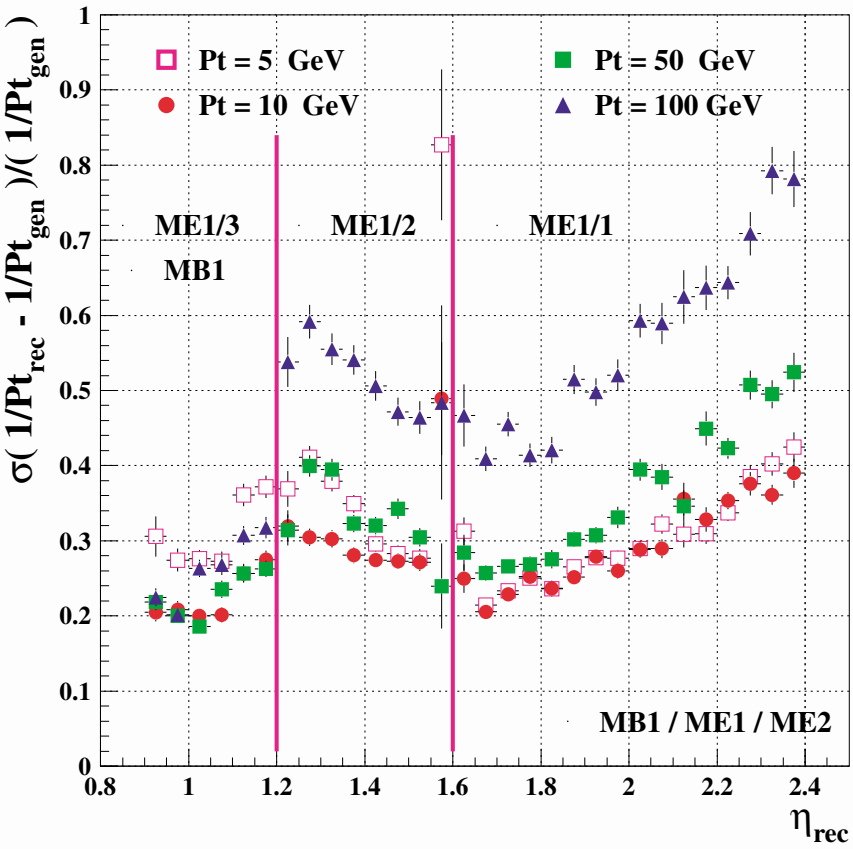

Fig. 12.15: The resolution of p_T for the 2-station measurement as a function of η for several p_T values. The p_T was reconstructed from $\Delta\phi$ measured between MB2/1 and ME1/3 for the overlap region, and between ME1 and ME2 in the endcap region.

p_T=5 GeV/*c*. There are no results for the three-station p_T measurement for $|\eta| < 1.2$ because this method of measurement was only implemented in the endcap region during this study. The figure shows that the three-station p_T measurement provides a significant improvement in the p_T resolution for low p_T muons in the endcap region as compared to the two-station p_T measurement.

These results are confirmed with the ORCA simulation. In particular, Fig. 12.17 shows the residual distribution of $(1/p_T^{meas} - 1/p_T^{true}) \times p_T^{true}$ for reconstructed muons with segments in at least two CSC stations, including ME1, with $5 < p_T^{true} < 50$ GeV/*c* and $1.2 < |\eta| < 2.0$. The distribution is centered at zero with a width of 29%.

The CSC trigger efficiency, as estimated from ORCA, for single muons as a function of η is shown in Fig. 12.18 for two sets of criteria applied to tracks found by the CSC Track-Finder. The solid line shows the efficiency for tracks that have segments in at least two muon stations, of which one must be ME1 for the endcap region ($|\eta| > 1.2$) in order to maintain satisfactory p_T

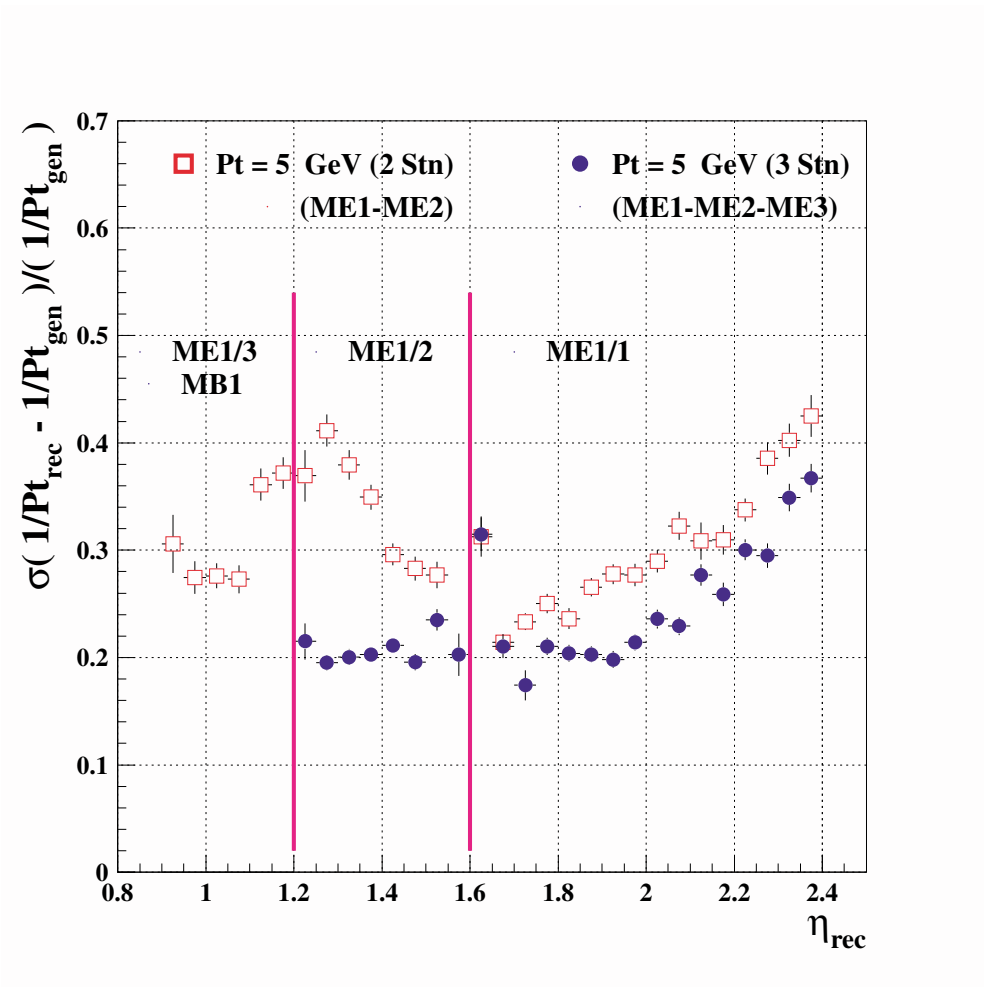

Fig. 12.16: The resolution of p_T for the 3-station measurement (ME1–ME2–ME3) as a function of η compared to the 2-station (ME1–ME2) measurement at $p_T = 5$ GeV/*c*.

resolution. Any two stations are allowed for the DT/CSC overlap region ($1.05 < |\eta| < 1.2$) to keep the efficiency high. In contrast to this loose set of conditions, the dashed line shows the efficiency for tracks when segments in at least three muon stations are required, including ME1 in the endcap region and MB1 in the DT/CSC overlap, so that the best p_T resolution is achieved. The efficiency is reduced in this case, particularly in the DT/CSC overlap, but maximum background rejection will be achieved. The configuration used in the trigger is a trade-off between efficiency and increased rate from low quality tracks. At high luminosity, we expect that three stations will be required if the CSC system is used standalone, without coincidence with the RPC system. It should be noted that in this efficiency plot, and those that follow, that all four CSC stations (ME1–ME4) are assumed to be present.

The CSC trigger efficiency for single muons as a function of ϕ is shown in Fig. 12.19. The solid line corresponds to the loose set of track requirements just described, and the dashed line corresponds to the tighter set. The overall CSC trigger efficiency is 92% for muons generated in

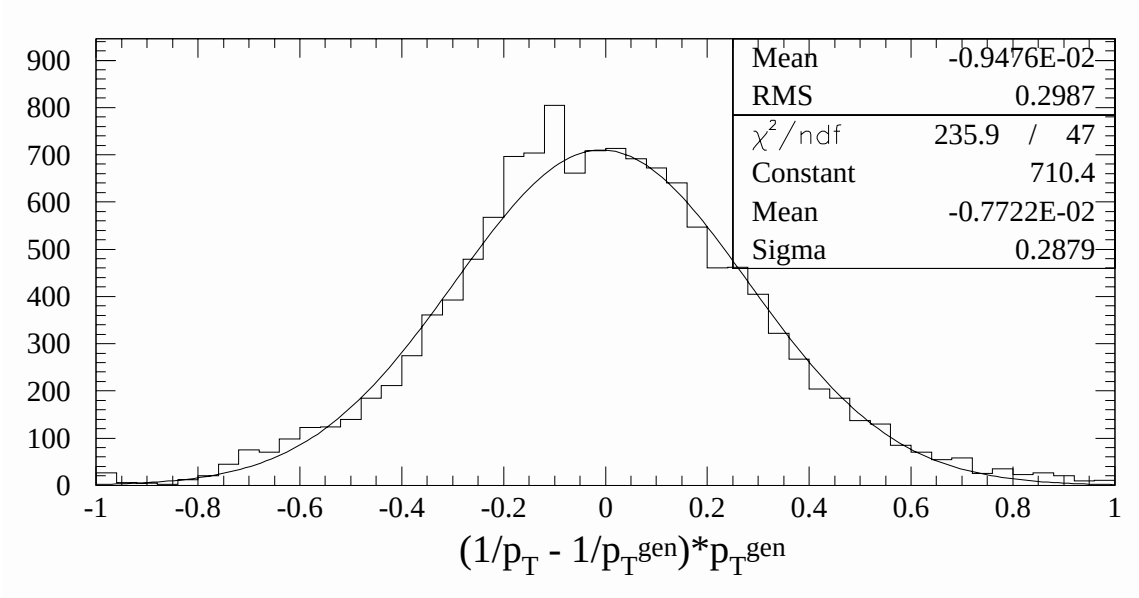

Fig. 12.17: Residual distribution of the inverse transverse momentum measured by the CSC Track-Finder for single muons generated flat in $5 < p_T < 50$ GeV/*c* and $1.2 < |\eta| < 2.0$. A track segment in ME1 is required.

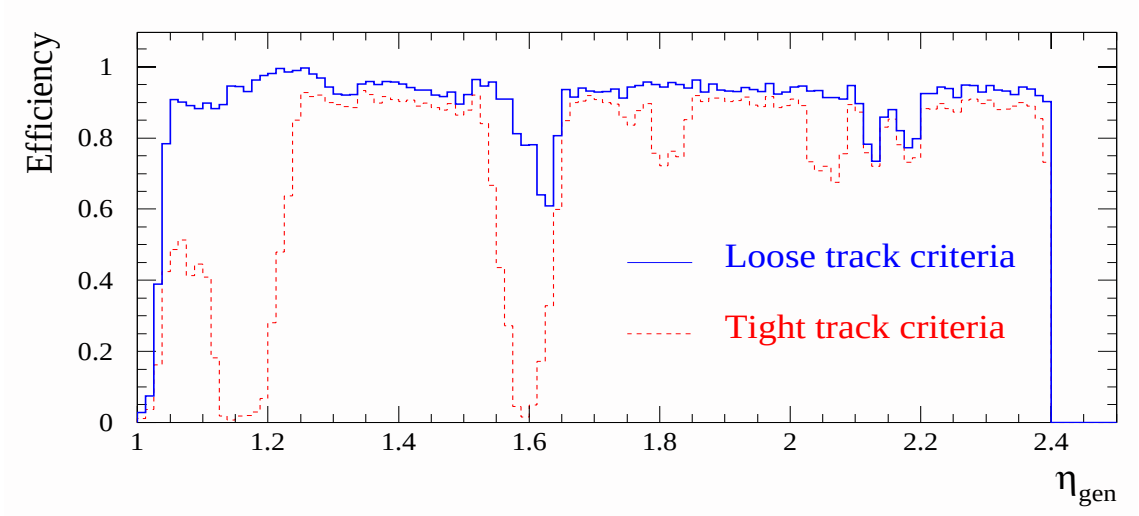

Fig. 12.18: The CSC Track-Finder efficiency as a function of η for single muons. The solid line corresponds to loose requirements on the track quality, the dashed line corresponds to tight requirements.

the range $1.05 < |\eta| < 1.2$ using the loose set of criteria, and is flat in ϕ. For the tight track criteria, the efficiency drops to 70% because of the geometric holes in the η coverage seen in Fig. 12.18.

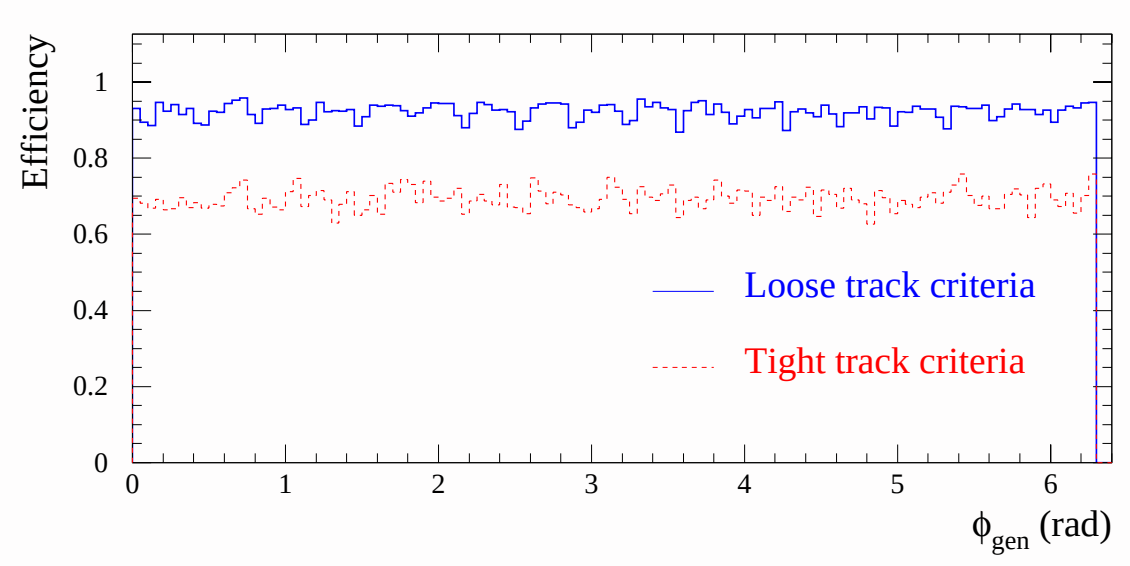

Fig. 12.19: The CSC Track-Finder efficiency as a function of ϕ for single muons. The solid line corresponds to loose requirements on the track quality, the dashed line corresponds to tight requirements.

Figure 12.20 shows the trigger efficiency turn-on curves as a function of the true p_T for several trigger thresholds (defined at 90% efficiency) for the loose set of track criteria in the endcap region. The sample of single muons used for the study are generated flat in pseudorapidity, $1.2 < |\eta| < 2.4$, so the p_T resolution is an average over this interval.

The CSC trigger rate has been studied using the samples of minimum bias events described in Chapter 8. In particular, an LHC luminosity of $10^{34}cm^{-2}s^{-1}$ is assumed, which implies 17.3 minimum bias events are piled-up on average every beam crossing. Figure 12.21 shows the single muon trigger rate from the CSC trigger as a function of the p_T threshold applied (defined at 90% efficiency) for the two sets of requirements on the track quality shown in the efficiency plots. The solid line corresponds to the loose set of criteria, which yields good trigger efficiency but high rate (>10 kHz for any threshold). This rate can be reduced to acceptable levels by the Global Muon Trigger, however, when the CSC tracks are combined with RPC tracks, as discussed in Chapter 14. On the other hand, the CSC trigger alone can reduce the rate to acceptable levels when the requirement on the tracks reported by the CSC Track-Finder is tightened, as shown by the dashed line in Fig. 12.21. A single muon rate of about 5 kHz is achieved when the single muon threshold is set to 25 GeV/*c* because of the improved p_T resolution of high quality tracks, at the expense of some efficiency loss.

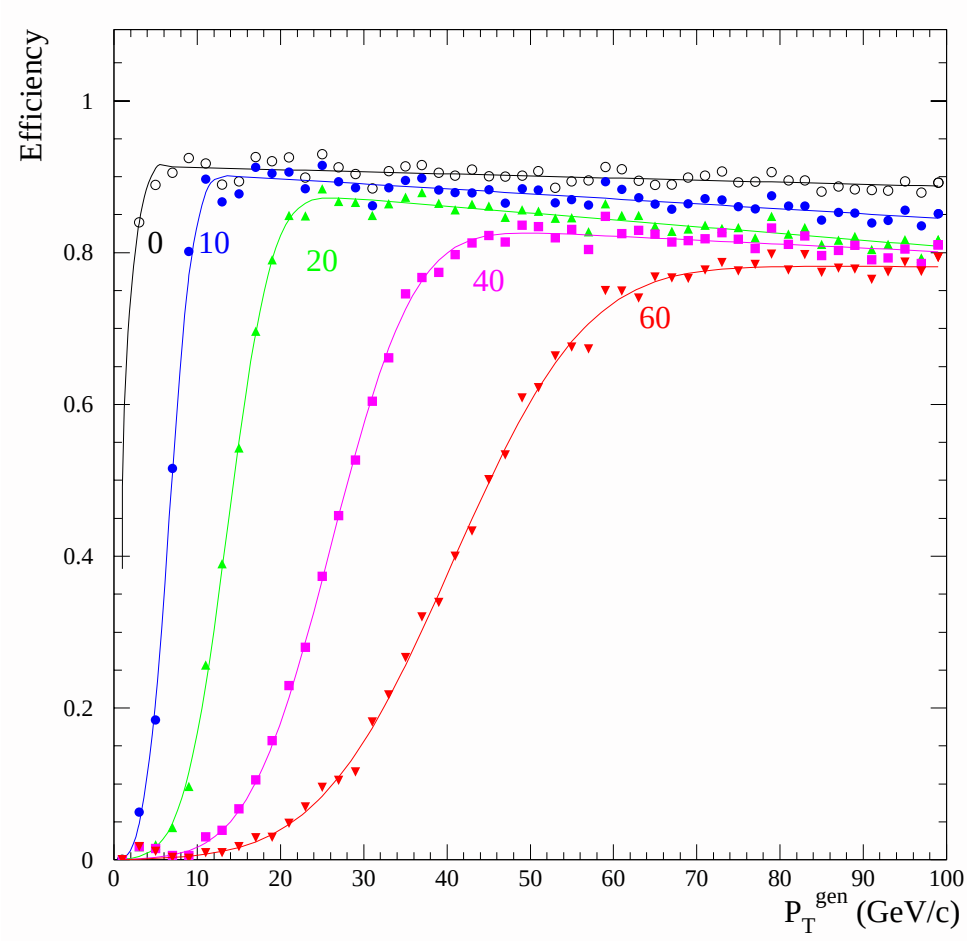

Fig. 12.20: CSC trigger efficiency turn-on curves as function of the true p_T for several trigger thresholds. The p_T thresholds are defined at 90% efficiency.

12.10 Prototypes and Tests

12.10.1 Sector Receiver

In order to test the Sector Processor with a full complement of inputs, three prototype Sector Receivers for the CSC Track-Finder have been built and tested. A photograph of one of them is shown in Fig. 12.22. The 12 optical receiver blocks are visible on the left followed (left to right) by de-serializers, the 6 front FPGAs, two columns of 16 memory LUTS with buffers in between, 6 Back FPGAs, and more buffers and Channel-Link drivers. The VME interface FPGA is in the upper right corner. Front-panel LED's display the status of a number of internal and VME bits. At the top of the board are Xilinx EEPROMs which load all FPGAs upon power-up. The board has 10 layers and about 9400 vias.

The (thirteen) Front, Back, and VME Interface FPGAs were implemented in Xilinx Virtex FPGAs, part number XCV50-6BG256C. The memory LUTs were implemented in 36

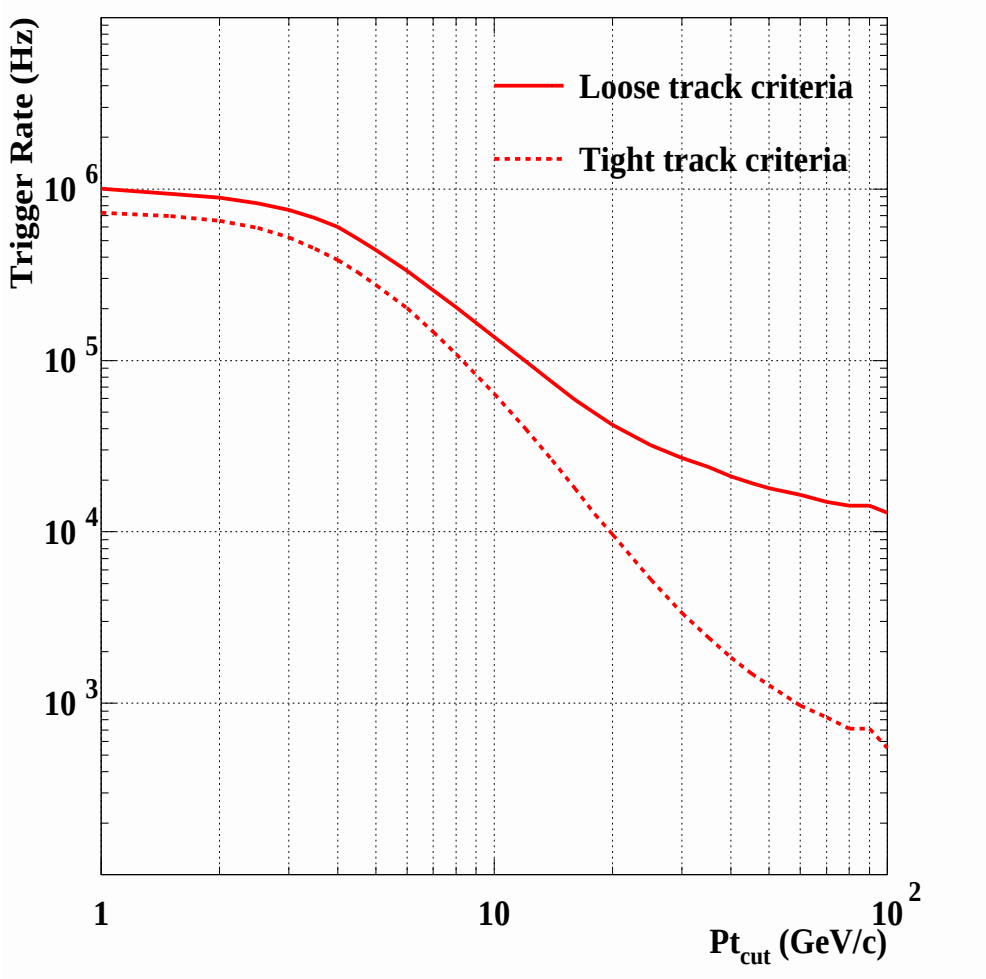

Fig. 12.21: Single muon trigger rate from the CSC system as a function of the p_T threshold (defined at 90% efficiency) for a luminosity of 10^{34} $cm^{-2}s^{-1}$. The solid line corresponds to loose requirements on the track quality, the dashed line corresponds to tight requirements.

identical 256Kx16 synchronous static RAMs, GSI part number GS74116TP. The use of SRAM allows the output of the second-stage RAMs to appear as soon as the inputs to the first-stage RAM propagate freely through the chips. Two of the SR prototypes were built with 8-ns RAM, the fastest available, to guarantee that the entire propagation through the two chips could take place within one cycle at 40 MHz. The third prototype was built with 10-ns RAM, which has worked equally well in all tests thus far. Further tests will determine if the (cheaper, more readily available) 10-ns RAM has a sufficient safety margin for reliability in production boards.

For tests involving the Sector Processor, ORCA-simulated CMS data and corresponding LUT contents are used. Thus, many of the LUTs contain the same contents, and the LUT addresses checked reflect the population of tracks in CMS. For further testing of a stand-alone Sector Receiver, random numbers are used for all LUT contents and event data, allowing more stringent tests of the addressing. Events are loaded into the Front FPGA in groups of 252, passed through

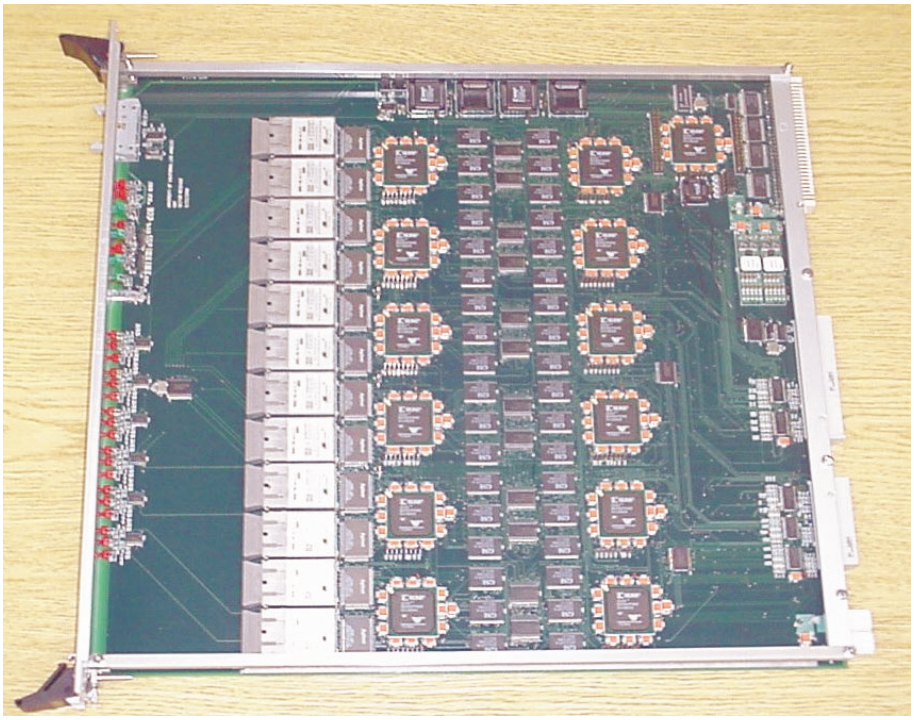

Fig. 12.22: Picture of the Sector Receiver prototype board

the SR at 40 MHz, and read out of the storage buffer in the Back FPGA. With random data, a typical test run of 100000 cycles of 252 events per cycle showed perfect agreement between the expected and observed data, for all three SR prototypes. In some longer trials, we have recently observed some rare discrepancies which are under investigation.

12.10.2 Sector Processor

A prototype Sector Processor for the CSC Track-Finder has been completed and tested, and first results are reported in Ref. [12.8]. Differences with respect to the design described in this Chapter have to do with the specification of the DT interface. In particular, the prototype design accommodates track segments from both MB2/1 *and* MB2/2 chambers (rather than MB2/1 only), and up to two track segments per chamber can be delivered on the same bunch crossing (rather than serialized across two bunch crossings).

The prototype receives its input from a custom point-to-point backplane operating at 280 MHz. The signals are transmitted and received using Channel-Link LVDS from National Semiconductor. The hardware implementation of the Sector Processor trigger logic is listed below.

Bunch Crossing Analyzer: The logic of the Bunch Crossing Analyzer is partitioned across 7 moderately-sized FPGAs from the Xilinx Virtex series (XCV50-6BG256C).

Extrapolation Units: The extrapolation logic, as well as the global FIFO which stores the information for the Assignment Unit, occupies 4 large Xilinx Virtex FPGAs (XCV400-6BG560C).

Track Assembler Units: As discussed in the text, the Track Assembler Units are realized with nine 256Kx16 SRAM memory chips from Integrated Device Technology.

Final Selection Unit: The Final Selection Unit is implemented in one Xilinx Virtex FPGA (XCV150-6BG352C).

Assignment Unit: The Assignment Unit is implemented in three Xilinx Virtex FPGAs (XCV50-6BG256C) and three 2Mx8 SRAM memory chips from Toshiba.

In addition to the fast trigger logic, a Xilinx Virtex FPGA makes up the VME interface for the board, and a parallel-to-serial interface chip (SCANPSC100F) from National made up the JTAG interface. A front panel connector provides the SP output in parallel LVDS.

In total, 17 Xilinx FPGAs with a ball-grid array footprint, 12 memory chips, and 25 32-bit buffers were used. The entire board was routed using 12 layers and approximately 10,000 vias. A picture of the prototype is shown in Fig. 12.23.

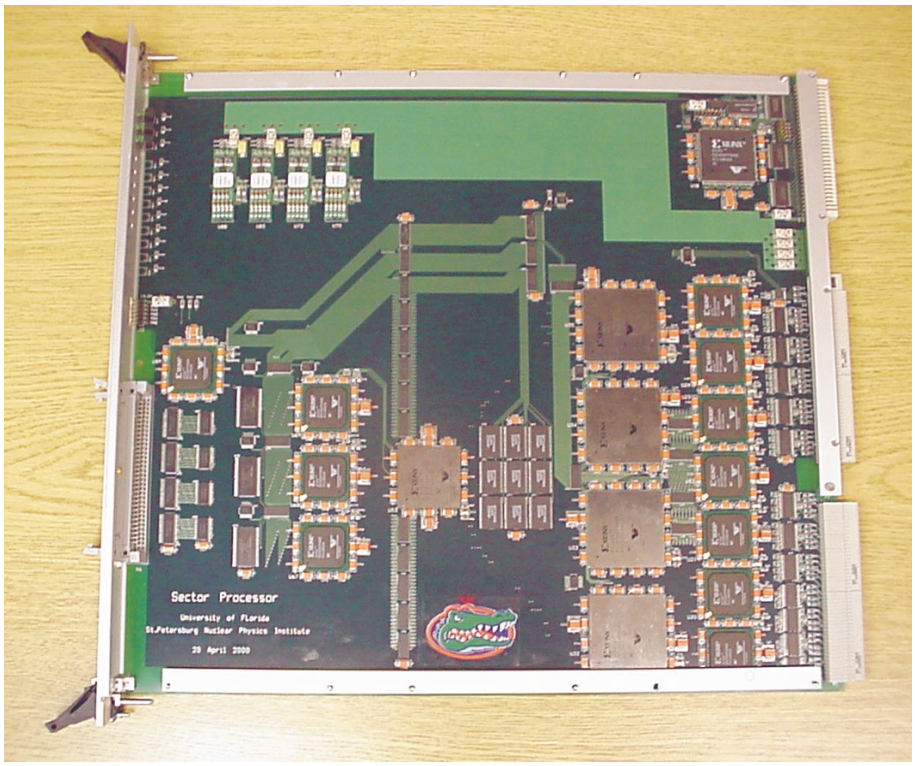

Fig. 12.23: Picture of the Sector Processor prototype board

12.10.3 Track-Finder Crate Test

The CSC Track-Finder prototypes (SR, SP, and CCB), as well as the MPC (described in Chapter 11) which sends CSC trigger primitives, underwent crate tests during the summer and fall of 2000. Figure 12.24 shows the arrangement of the prototypes in a single 9U VME crate. All critical functions of the boards were tested. The CCB prototype was used to distribute clock and

control signals to each board. In particular, the CCB issues a BC0 to initiate the start of a chain test between two or more prototypes.

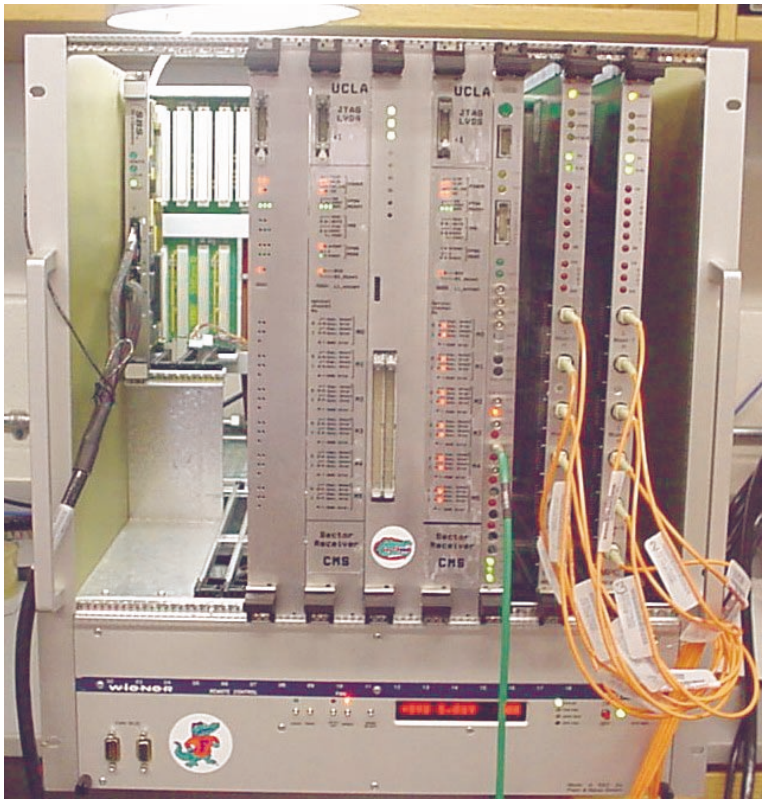

Fig. 12.24: Photograph of the CSC trigger prototypes undergoing system tests in a 9U VME crate. From right to left: two Muon Port Cards, the Clock and Control Board, a Sector Receiver, the Sector Processor, and two other Sector Receivers. The Port Cards send trigger primitives to the Sector Receivers via optical cables, and the Sector Receiver and Sector Processor communicate through a custom Channel-Link backplane. Data was successfully transmitted and processed in this chain test with no errors.

VME Interface: Downloading of FPGA programs, LUT contents, and test data was achieved with a PCI to VME interface (Bit3 Model 917 from SBS Technology) connected to a PC. Each board contained an FPGA for the VME interface that was automatically loaded on power-up. The data for the rest of the FPGA JTAG chain was sent in parallel over the VME bus and de-serialized at up to 25 MHz on each board.

SR Functionality: The data conversion FPGAs and memories of the Sector Receiver have been successfully tested dynamically at 40 MHz using pseudo-random input data as well as simulated muon data from ORCA. Perfect agreement was achieved between the hardware and simulation for 30,000 cycles of 256 bx. The latency of the FPGA and SRAM conversion is 2 bx, excluding de-serialization/serialization on the input/output. In particular, data are successfully sent through two consecutive SRAMs within 1 bx.

SP Functionality: All the track-finding algorithms of the Sector Processor have been tested dynamically at 40 MHz, except for the Bunch Crossing Analyzer which was configured as an input FIFO for the tests. In particular, perfect agreement was achieved between the ORCA simulation and the SP hardware for about 200,000 single muon events. Moreover, agreement also was achieved when the same events were piled-up three at a time to mimic triple-muon events, which is a stronger test of the track assembly and sorting functions on the board. Each sub-processor on the board was tested separately and in concert. The maximum clock frequency of the extrapolation and track assembly logic was measured to be 63 MHz. The latency of the SP prototype is determined to be 15 bx, excluding the Channel-Link de-serialization.

MPC to SR Optical Communication and Data-Flow Tests: After an SR passed its stand-alone tests, a Muon Part Card was connected to it with 100-m optical fibers driven and received by HP GLinks. Unsorted tracks were put into a buffer in the front of the MPC. These were clocked through the MPC, with the sorted highest-priority tracks going on to the SR, through the SR, where they were recorded by the Back FPGA. In October 2000, success was first achieved in a test in which 1.6 million random events were processed without error before the test was stopped. Since then, these boards and a second MPC have been tested together in a different crate. No errors were encountered using random data or simulated muon data in the tests. The overall latency from the input of the MPC to the input of the SR is 28 bx, broken down as follows: 5 bx for MPC processing, 1 bx for HP GLink serialization, 20 bx for transmission through 100 m of optical fiber, and 2 bx for HP GLink de-serialization.

SR to SP Backplane Communication: The integrity of the data sent from two SR prototypes (corresponding to stations ME2–ME4) to the SP prototype through the custom Channel-Link backplane operating at 280 MHz has been verified. A FIFO in either the input or output of each SR can be loaded with data 256 bx deep. The overall latency from the input of the SR to the output of the SP is 20 bx, broken down as follows: 2 bx for SR processing, 4 bx for Channel-Link serialization/de-serialization and transmission over the custom backplane, and 15 bx for SP processing.

MPC to SR to SP Chain Test: The entire trigger path from the MPC to the SP has been tested successfully using simulated muons from ORCA as input. Two MPCs representing ME2 and ME3 of one trigger sector sent data to one SR, which in turn re-formatted and transmitted the data to one SP. The SP successfully reconstructed tracks in agreement with the ORCA simulation.

12.11 Maintenance and Operation

We plan to have spare boards and components sufficient for 10 years of LHC operations. Each board will have a VME-readable register that labels the particular trigger configuration that was loaded into the board. The entire set of LUT contents and FPGA programs will be given one unique identifier. There are many FPGAs and LUTs in the design, so it is important to verify that the correct patterns were loaded. We do not foresee generating LUT contents on the boards from a VME-addressable register. Rather, we will have precompiled FPGA programs (and LUT contents) that are certified to work properly and that meet all timing specifications.

12.12 Status and Schedule

Prototypes of all boards for the CSC Track-Finder, except the CSC Muon Sorter, were constructed and tested by October 2000. The schedule for the CSC trigger development is shown in Fig. 11.29. Final design of all boards should be finished by mid-2002. Final production is expected to be completed by the end of 2003, with full system testing beginning in 2004.

References

[12.1] S.M. Wang and D. Acosta, "Simulation Studies on the Transverse Momentum Resolution of the CSC Track-Finder", CMS IN 2000/026.

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

[12.3] D. Acosta *et al.*, "Coverage of the DT/CSC Overlap Region by the Level-1 Track-Finders", CMS Note in preparation.

[12.4] J. Christiansen, A. Marchioro, P. Moreira, TTCrx Reference Manual, RD12 Working Document.

[12.5] D. Acosta *et al.*, "The Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System", CMS Note 1999/060.

[12.6] D. Acosta *et al.*, "The Track-Finder Trigger for the Level-1 Trigger of the CMS Endcap Muon System", Proceedings of the Fifth Workshop on Electronics for LHC Experiments (LEB'99), Snowmass, CERN 99-09, CERN/LHCC/99-33.

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

[12.8] S.M. Wang *et al.*, "A Prototype Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System", Proceedings of the Meeting of the Division of Particles and Fields of the American Physical Society (DPF2000), Columbus, to be published in the International Journal of Modern Physics A.