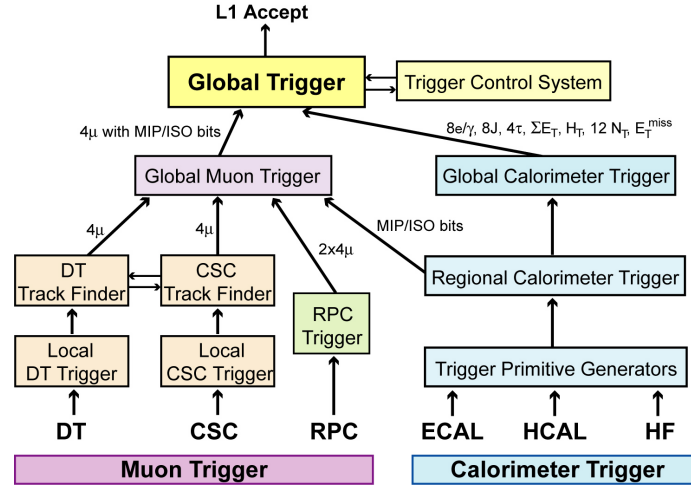


## Chapter 8

# Trigger

The LHC provides proton-proton and heavy-ion collisions at high interaction rates. For protons the beam crossing interval is 25 ns, corresponding to a crossing frequency of 40 MHz. Depending on luminosity, several collisions occur at each crossing of the proton bunches (approximately 20 simultaneous pp collisions at the nominal design luminosity of  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ ). Since it is impossible to store and process the large amount of data associated with the resulting high number of events, a drastic rate reduction has to be achieved. This task is performed by the trigger system, which is the start of the physics event selection process. The rate is reduced in two steps called Level-1 (L1) Trigger [187] and High-Level Trigger (HLT) [188], respectively. The Level-1 Trigger consists of custom-designed, largely programmable electronics, whereas the HLT is a software system implemented in a filter farm of about one thousand commercial processors. The rate reduction capability is designed to be at least a factor of  $10^6$  for the combined L1 Trigger and HLT. The design output rate limit of the L1 Trigger is 100 kHz, which translates in practice to a calculated maximal output rate of 30 kHz, assuming an approximate safety factor of three. The L1 Trigger uses coarsely segmented data from the calorimeters and the muon system, while holding the high-resolution data in pipelined memories in the front-end electronics. The HLT has access to the complete read-out data and can therefore perform complex calculations similar to those made in the the analysis off-line software if required for specially interesting events. Since HLT algorithms will evolve with time and experience they are not described here. More information may be found in [189]. For reasons of flexibility the L1 Trigger hardware is implemented in FPGA technology where possible, but ASICs and programmable memory lookup tables (LUT) are also widely used where speed, density and radiation resistance requirements are important. A software system, the Trigger Supervisor [190], controls the configuration and operation of the trigger components.

The L1 Trigger has local, regional and global components. At the bottom end, the Local Triggers, also called Trigger Primitive Generators (TPG), are based on energy deposits in calorimeter trigger towers and track segments or hit patterns in muon chambers, respectively. Regional Triggers combine their information and use pattern logic to determine ranked and sorted trigger objects such as electron or muon candidates in limited spatial regions. The rank is determined as a function of energy or momentum and quality, which reflects the level of confidence attributed to the L1 parameter measurements, based on detailed knowledge of the detectors and trigger electronics and on the amount of information available. The Global Calorimeter and Global Muon Triggers



**Figure 8.1:** Architecture of the Level-1 Trigger.

determine the highest-rank calorimeter and muon objects across the entire experiment and transfer them to the Global Trigger, the top entity of the Level-1 hierarchy. The latter takes the decision to reject an event or to accept it for further evaluation by the HLT. The decision is based on algorithm calculations and on the readiness of the sub-detectors and the DAQ, which is determined by the Trigger Control System (TCS). The Level-1 Accept (L1A) decision is communicated to the sub-detectors through the Timing, Trigger and Control (TTC) system. The architecture of the L1 Trigger is depicted in figure 8.1. The L1 Trigger has to analyze every bunch crossing. The allowed L1 Trigger latency, between a given bunch crossing and the distribution of the trigger decision to the detector front-end electronics, is  $3.2 \mu\text{s}$ . The processing must therefore be pipelined in order to enable a quasi-deadtime-free operation. The L1 Trigger electronics is housed partly on the detectors, partly in the underground control room located at a distance of approximately 90 m from the experimental cavern.

## 8.1 Calorimeter trigger

The Trigger Primitive Generators (TPG) make up the first or local step of the Calorimeter Trigger pipeline. For triggering purposes the calorimeters are subdivided in trigger towers. The TPGs sum the transverse energies measured in ECAL crystals or HCAL read-out towers to obtain the trigger tower  $E_T$  and attach the correct bunch crossing number. In the region up to  $|\eta| = 1.74$  each trigger tower has an  $(\eta, \phi)$ -coverage of  $0.087 \times 0.087$ . Beyond that boundary the towers are larger. The TPG electronics is integrated with the calorimeter read-out. The TPGs are transmitted through high-speed serial links to the Regional Calorimeter Trigger, which determines regional candidate electrons/photons, transverse energy sums,  $\tau$ -veto bits and information relevant for muons in the form of minimum-ionizing particle (MIP) and isolation (ISO) bits. The Global Calorimeter Trigger determines the highest-rank calorimeter trigger objects across the entire detector.

## Calorimeter trigger primitive generators


The ECAL on-detector front-end electronics boards, each serving 25 crystals, receive the ADC signals from the very front-end electronics located at the rear of the detector modules. They contain most of the TPG pipeline in six radiation-hard  $0.25\ \mu\text{m}$  CMOS ASIC chips named FENIX. An off-detector Trigger Concentrator Card (TCC) collects the primitives from 68 front-end boards in the barrel and 48 boards in the endcaps through optical links. The TCCs finalize the TPG generation and encoding, store the trigger primitives during the L1 latency time and transmit them to the RCT by dedicated daughter boards, the Synchronization and Link Boards (SLB), upon reception of a L1A signal. The SLBs synchronize the trigger data through circuits that histogram the LHC bunch crossing structure. Each trigger tower is aligned with the bunch crossing zero signal. A Data Concentrator Card (DCC) performs the opto-electronic conversion and deserialization of the serial input data streams and sends the read-out data collected from the front-end boards to the DAQ. Clock and Control System (CCS) boards distribute the clock, the L1A and control signals to the TCC, the DCC and the on-detector electronics. The ECAL TPG hardware is contained in twelve 9U VME crates for the barrel and six for the endcaps.

The front-end modules of the hadron calorimeter contain Charge Integrator and Encoder (QIE) ADC chips to digitize the signals from the photo detectors. Optical links transmit the data to the HCAL Trigger and Readout (HTR) boards. Each HTR board processes 48 channels. It linearizes, filters and converts the input data to generate the HCAL trigger primitives. The energy values of front and back towers are added and the bunch crossing number is assigned by a peak filtering algorithm. As for the ECAL, the primitives are sent to the RCT by SLBs, and read-out data are collected by a DCC. The HCAL trigger electronics is contained in 26 9U VME crates. Each crate houses 18 HTR boards and one DCC.


## Regional Calorimeter Trigger

The Regional Calorimeter Trigger [191] determines electron/photon candidates and transverse energy sums per calorimeter region. Information relevant for muons about isolation and compatibility with minimally ionizing particles is also calculated. A region consists of  $4 \times 4$  trigger towers except in HF where a region is one trigger tower. Electromagnetic and hadronic transverse energies are summed in each tower.

The  $e/\gamma$  trigger algorithm (figure 8.2) starts by determining the tower with the largest energy deposit and is applied across the entire ECAL region. The energy of the tower with the next-highest deposit in one of the four broad side neighbours is then added. Isolated and non-isolated  $e/\gamma$  within  $|\eta| \leq 2.5$  are determined by the trigger. A non-isolated  $e/\gamma$  requires passing of two shower profile vetoes. The first one is based on a fine-grain crystal energy profile reflecting the lateral extension of a shower. The fine-grain bit is set by the TPG if the shower is contained in a matrix of  $2 \times 5$  crystals. The matrix is dimensioned such that it also allows for the detection of bremsstrahlung due to the magnetic field. The second one is based on the ratio of the deposited energies in the hadronic and in the electromagnetic sections. A typical maximal value of 5% is allowed for that ratio. An isolated electron/photon candidate has to pass the previous vetoes for all eight neighbouring towers. In addition, at least one quiet corner made of four groups of five electromagnetic towers surrounding the hit tower is required. Four isolated and four non-isolated  $e/\gamma$  per region are forwarded to the GCT.



**Figure 8.2:** Electron/photon algorithm.



**Figure 8.3:** Electron Isolation Card.

The RCT also sums the transverse energy in a given region of the central calorimeter (HF is not included) and determines  $\tau$ -veto bits for the identification of jets from one- and three-prong  $\tau$ -decays, which are narrower than ordinary quark/gluon jets. A  $\tau$ -veto bit is set unless the pattern of active towers corresponds to at most  $2 \times 2$  contiguous trigger towers within a  $4 \times 4$  tower region. Jets can be classified as  $\tau$ -jet only at  $|\eta| < 3.0$  (not in HF).

The RCT hardware consists of 18 regional 9U VME crates and one 6U clock distribution crate located in the underground control room. Each crate covers a region of  $\Delta\eta \times \Delta\phi = 5.0 \times 0.7$ . Receiver cards are plugged into the rear of the regional crates. Seven cards per crate receive the ECAL and HCAL primitives. The HF primitives are directly received on a Jet/Summary card. The serial input data are converted to 120 MHz parallel data, deskewed, linearized and summed before transmission on a 160 MHz custom monolithic backplane to seven Electron Isolation Cards (EIC) and one Jet/Summary Card (JSC) mounted at the front side of the crate. Different ASICs perform the algorithm calculations. An EIC is shown in figure 8.3.

### Global Calorimeter Trigger

The Global Calorimeter Trigger determines jets, the total transverse energy, the missing transverse energy, jet counts, and  $H_T$  (the scalar transverse energy sum of all jets above a programmable threshold). It also provides the highest-rank isolated and non-isolated  $e/\gamma$  candidates.


Jets are found by a four-stage clustering technique based on jet finders operating in  $2 \times 12$  cells in  $\phi$  and  $\eta$ , spanning  $40^\circ$  and half the detector, respectively, in these directions. The cell at  $\eta=0$  is duplicated. In the first stage mini-clusters are created by summing energies within  $2 \times 3$  cells if a central cell has more energy than neighbouring cells. In the second stage the three largest mini-clusters in each of the two  $\phi$ -strips are transferred in opposite  $\phi$ -directions. These are compared against the existing mini-clusters on the receiving  $\phi$ -strip. Mini-clusters adjacent or diagonally adjacent to a larger mini-cluster are removed. In the third and fourth stages the received mini-clusters that survive have their three adjacent cells in the receiving  $\phi$ -strip combined to make a  $3 \times 3$  cell. A jet is classified as a  $\tau$  jet if none of the corresponding RCT regions had a  $\tau$ -veto bit set. After sorting, up to four jets and four  $\tau$  jets from the central HCAL and four jets from HF are forwarded to the GT. The magnitude and direction of the missing energy and the total transverse energy are

computed from the regional transverse energy sums in the two coordinates transverse to the beam within  $|\eta| < 5$ . Twelve jet counts for different programmable  $E_T$ -thresholds and optionally also different  $(\eta, \phi)$ -regions are computed. Muon MIP/ISO bits are received from the RCT along with the  $e/\gamma$  data and are forwarded to the GMT through a dedicated muon processing system. Apart from triggering, the GCT also acts as the read-out system for the RCT. The GCT has in addition the capability to monitor rates for certain trigger algorithms and from those deduce information about the LHC luminosity as seen by the CMS trigger system.

All GCT electronics is located in the underground control room. The large amount of data from the RCT crates are transmitted electronically to 63 Source Cards, which reorder the data onto 252 optical fibres. The core of the GCT processing is performed by Leaf Cards, which can be configured as electron or jet cards. Several Leaf Cards can be connected with each other in order to perform complex tasks such as the jet finding. There are two Leaf Cards for electrons and six for jets. Each electron leaf card receives the  $e/\gamma$  data from one half of the RCT crates on 27 fibres and sorts them. Each jet card receives 30 regional sum fibres from three RCT crates via the source cards. They perform the jet clustering and transmit the jet candidates to two Wheel Cards for sorting and data compression. They also calculate partial energy sums and jet counts and forward them to the Wheel Cards. A Concentrator Card finally collects the data from all Electron Leaf and Wheel Cards and performs the final sorting for electrons/photons, completes the jet finding in the boundaries between groups of three Leaf Cards, sorts all jets, calculates the global energy and jet count quantities and sends the final results to the GT and the DAQ. In addition to the tasks involving  $e/\gamma$ 's, jets and energy sums, the GCT also handles MIP/ISO bits for muons. They are processed by three muon processing cards, which receive 6 muon fibres each from Source Cards. The processor design is built on an evolution of the leaf concept and uses a modular, low-latency architecture based on the  $\mu$ TCA industry standard [193]. An active custom backplane based on the principle of a crosspoint switch allows a programmable routing of the 504 MIP/ISO bits, which are then transmitted to the GMT on 24 links.

## 8.2 Muon trigger

All three muon systems – the DT, the CSC and the RPC – take part in the trigger. The barrel DT chambers provide local trigger information in the form of track segments in the  $\phi$ -projection and hit patterns in the  $\eta$ -projection. The endcap CSCs deliver 3-dimensional track segments. All chamber types also identify the bunch crossing from which an event originated. The Regional Muon Trigger consists of the DT and CSC Track Finders, which join segments to complete tracks and assign physical parameters to them. In addition, the RPC trigger chambers, which have excellent timing resolution, deliver their own track candidates based on regional hit patterns. The Global Muon Trigger then combines the information from the three sub-detectors, achieving an improved momentum resolution and efficiency compared to the stand-alone systems. The initial rapidity coverage of the muon trigger is  $|\eta| \leq 2.1$  at the startup of LHC. The design coverage is  $|\eta| \leq 2.4$ .



**Figure 8.4:** Drift Tube Local Trigger.

### Drift Tube local trigger

The electronics of the DT local trigger consists of four basic components (figure 8.4): Bunch and Track Identifiers (BTI), Track Correlators (TRACO), Trigger Servers (TS) and Sector Collectors (SC). While the SCs are placed on the sides of the experimental cavern, all other trigger and read-out electronics is housed in minicrates on the front side of each chamber. All devices are implemented in custom-built integrated circuits. The BTIs are interfaced to the front-end electronics of the chambers. Using the signals from the wires they generate a trigger at a fixed time after the passage of the muon. Each BTI searches for coincident, aligned hits in the four equidistant planes of staggered drift tubes in each chamber superlayer. The association of hits is based on a mean-timer technique [194], which uses the fact that there is a fixed relation between the drift times of any three adjacent planes. From the associated hits, track segments defined by position and angular direction are determined. The spatial resolution of one BTI is better than 1.4 mm, the angular resolution better than 60 mrad. The BTI algorithm is implemented in a 64-pin ASIC with CMOS  $0.5 \mu\text{m}$  Standard Cell technology. There are a few hundred BTIs per chamber.

The DT chambers have two  $\Phi$ -type superlayers, measuring  $\phi$  coordinates. The TRACO attempts to correlate the track segments measured in each of them. If a correlation can be found, the TRACO defines a new segment, enhancing the angular resolution and producing a quality hierarchy. Four BTIs in the inner  $\Phi$ -type superlayer and 12 BTIs in the outer  $\Phi$ -type superlayer are connected to one TRACO. The number of TRACOs is 25 for the largest muon chamber type. The TRACO is implemented in a 240-pin ASIC with CMOS  $0.35 \mu\text{m}$  Gate Array technology. The trigger data of at most two track segments per bunch crossing reconstructed by each TRACO are transmitted to the TS, whose purpose is to perform a track selection in a multitrack environment.

The TS has two components, one for the transverse projection ( $\text{TS}\phi$ ) and the other for the longitudinal projection ( $\text{TS}\theta$ ). The first one processes the output from the TRACO, whilst the second uses directly the output of the BTIs of the  $\theta$  view delivered by the  $\Theta$ -type superlayers present in the three innermost muon stations. The  $\text{TS}\phi$  consists itself of two components, the

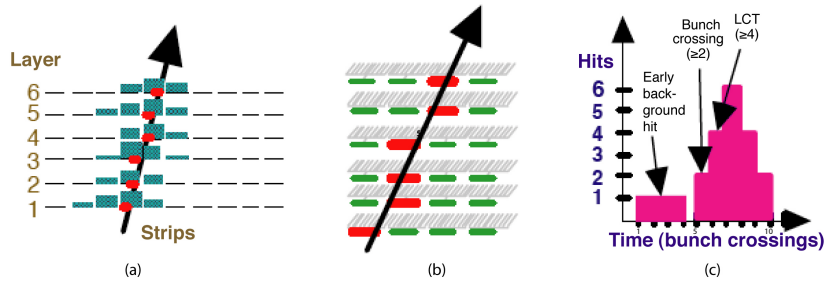
Track Sorter Slave (TSS) and the Track Sorter Master (TSM). The TSS preselects the tracks with the best quality and the smallest bending angle based on a reduced preview data set coming from the TRACOs in order to save processing time. A select line in the TRACO with the best track is then activated and the TRACO is allowed to send the full data to the TSM. The corresponding preview data are also sent to the TSM for a second stage processing. The TSM analyzes up to seven preview words from the TSSs. The output consists of the two tracks with the highest transverse momentum. There is one TSM per muon station. In the longitudinal view the  $TS\theta$  groups the information from the 64 BTIs per chamber. From each BTI two bits are received, a trigger bit and a quality bit. A logic OR of groups of eight bits is applied. The output data consist of 8 bits indicating the position of the muon and 8 quality bits.

The requirement of robustness implies redundancy, which introduces, however, a certain amount of noise or duplicate tracks giving rise to false triggers. Therefore the BTIs, the TRACOs and the different parts of the TS contain complex noise and ghost reduction mechanisms. The trigger and also the read-out data from each of the sixty  $30^\circ$ -sectors of CMS are sent to Sector Collector (SC) units, where the trigger information — the position, transverse momentum and track quality — is coded and transmitted to the DT regional trigger, called the Drift Tube Trigger Track Finder (DTTF), through high-speed optical links.

### **Cathode Strip Chamber local trigger**

The endcap regions are challenging for the trigger since many particles are present and muons at a given  $p_T$  have a higher momentum than in the barrel, which gives rise to more bremsstrahlung photons. In addition, photon conversions in a high-radiation (neutron-induced) environment occur frequently. Therefore the CSCs consist of six layers equipped with anode wires and cathode strips, which can be correlated. Muon track segments, also called Local Charged Tracks (LCT), consisting of positions, angles and bunch crossing information are first determined separately in the nearly orthogonal anode and cathode views. They are then correlated in time and in the number of layers hit. The cathode electronics is optimized to measure the  $\phi$ -coordinate, the anode electronics to identify the bunch crossing with high efficiency.

An electric charge collected by the anode wires induces a charge of opposite sign in the cathode strips nearby. The trigger electronics determines the centre of gravity of the charge with a resolution of half a strip width, between 1.5 and 8 mm, depending on the radius. By demanding that at least four layers are hit, the position of a muon can be determined with a resolution of 0.15 strip widths. Due to the finite drift time the anode signals in the six chamber layers are spread out over an interval of more than two bunch crossings. As for the cathodes, at least four coincident hits are required, since in contrast to neutron-induced background a real muon leaves coincident signals in at least four layers with a probability that exceeds 99%. Actually a coincidence of two signals (pre-trigger) is used to identify the crossing, in order to allow for long drift time hits to arrive. A validation of the track occurs if in the following two bunch crossings at least four coincident signals are found. In order to reduce the number of trigger channels 10 to 15 anode wires are ORed. Figure 8.5 shows the principles of the cathode and anode trigger electronics and the bunch crossing assignment.



**Figure 8.5:** Cathode Strip Chamber Local Trigger: (a) Cathode LCT formation from strips, (b) Anode LCT formation from wire group hits, (c) Bunch crossing assignment.

The track segments from the cathode and anode electronics are finally combined into three-dimensional LCTs. They are characterized by the high-precision  $\phi$ -coordinate in the bending plane, the bending angle  $\phi_b$ , a rough  $\eta$ -value and the bunch crossing number. The best two LCTs of each chamber are transmitted to the regional CSC trigger, called the CSC Track Finder (CSCTF), which joins segments to complete tracks.

The hardware of the CSC local trigger consists of seven types of electronics boards. Cathode and anode front-end boards (CFEB and AFEB) amplify and digitize the signals. Anode LCT-finding boards (ALCT) latch the anode hits at 40 MHz, find hit patterns in the six chamber layers that are consistent with having originated at the vertex, and determine the bunch crossing. They send the anode information to the Cathode LCT-finding plus Trigger Motherboard (CLTC/TMB) cards. The CLCT circuits look for strip hit patterns consistent with high-momentum tracks. The TMB circuits perform a time coincidence of cathode and anode LCT information. If a coincidence is found, they send the information to the Muon Port Cards (MPC). The TMB selects up to two LCTs based on quality cuts. In order to cancel out ghosts a coincidence with RPC hits is established if two or more LCTs are found. A MPC receives the LCTs from the CLTC/TMBs of one endcap muon station sector, selects the best two or three LCTs depending on the station number and sends them over optical links to the CSC Track Finder. The anode and cathode LCTs and the raw hits are recorded by DAQ motherboards (DAQMB) and transmitted to the CSC detector-dependent units (DDU) belonging to the DAQ system upon reception of a L1A signal. The LHC timing reference, the L1A decision, the bunch crossing number and bunch counter reset signals are distributed by the Clock and Control Boards (CCB). The front-end boards and the ALCTs are mounted directly on the chambers. The rest of the local trigger electronics is housed in 48 peripheral crates on the endcap disks. The optical fibres to the control room depart from there. Except for the comparator-network ASIC implemented in the CLCT module, the CSC trigger electronics is built in FPGA technology.

### Resistive Plate Chamber trigger

The RPCs are dedicated trigger detectors. Several layers of double-gap RPCs are mounted on the DT and CSC tracking chambers, six in the central region (two layers on the inside and outside of the two innermost muon stations, one on the inside of the two outermost stations) and four in the forward parts (one layer on the inside of each station). Their main advantage is their excellent



timing resolution of about 1 ns, which ensures an unambiguous bunch crossing identification. For triggering purposes the measurement of the momentum of a particle is also important. In the magnetic field, muons are bent in the plane transverse to the LHC beams. It is sufficient to measure the azimuthal coordinate  $\phi$  at several points along the track to determine the bending and thus the  $p_T$ . Therefore the RPC strips run parallel to the beam pipe in the barrel, and radially in the endcaps. There are about 165 000 strips in total, which are connected to front-end boards (FEB) handling 16 channels each.


The RPC trigger is based on the spatial and temporal coincidence of hits in several layers. It is segmented in 33 trigger towers in  $\eta$ , which are each subdivided in 144 segments in  $\phi$ . As opposed to the DT/CSC, there is no local processing on a chamber apart from synchronization and cluster reduction. The Pattern Comparator Trigger (PACT) logic [195] compares strip signals of all four muon stations to predefined patterns in order to assign  $p_T$  and electric charge, after having established at least three coincident hits in time in four planes. Spatially the PACT algorithm requires a minimum number of hit planes, which varies depending on the trigger tower and on the  $p_T$  of the muon. Either 4/6 (four out of six), 4/5, 3/4 or 3/3 hit layers are minimally required. A quality parameter reflects the numbers of hit layers. For six planes there are typically 14 000 possible patterns. The outer section of the hadron calorimeter (HO) consists of scintillators placed after the magnet coil up to  $|\eta| < 1.24$ . Their signals can also be taken into account by the RPC trigger in order to reduce rates and suppress background [196]. The algorithm requires HO confirmation for low-quality RPC triggers. The optical links from the four HO HTR boards are received by the RPC trigger boards, and the signals are treated and incorporated in the PACT logic like an additional RPC plane, with the required number of planes hit increased by one.

The RPC signals are transmitted from the FEBs, which contain ASICs manufactured in 0.8  $\mu\text{m}$  BiCMOS technology, to the Link Boards (LB), where they are synchronized, multiplexed, serialized and then sent via 1732 optical links to 108 Trigger Boards in 12 trigger crates in the control room. The 1640 LBs are housed in 136 Link Board Boxes. The Trigger Boards contain the complex PACT logic which fits into a large FPGA. There are 396 PACT chips in the system. Since duplicate tracks may be found due to the algorithm concept and the geometry, a *ghost busting* logic is also necessary. The RPC muon candidates are sorted separately in the barrel and forward regions. The best four barrel and the best four forward muons are sent to the Global Muon Trigger. The RPC data record is generated on the Data Concentrator Card, which receives data from the individual trigger boards.

### **Drift Tube and Cathode Strip Chamber track finders**

The regional muon trigger based on the precision tracking chambers consists of the Drift Tube Track Finder (DTTF) in the barrel [197] and the CSC Track Finder (CSCTF) in the endcaps [198]. They identify muon candidates, determine their transverse momenta, locations and quality. The candidates are sorted by rank, which is a function of  $p_T$  and quality. The DTTF and the CSCTF each deliver up to four muons to the Global Muon Trigger.

The track finding principle relies on extrapolation from a source track segment in one muon station to a possible target segment in another station according to a pre-calculated trajectory originating at the vertex. If a compatible target segment with respect to location and bending angle is



**Figure 8.6:** Track Finder extrapolation scheme.

found, it is linked to the source segment. A maximum number of compatible track segments in up to four muon stations is joined to form a complete track, to which parameters are then assigned. The extrapolation principle is shown in figure 8.6. While the CSCTF incorporates 3-dimensional spatial information from the CSC chambers in the track finding procedure, the DTTF operates 2-dimensionally in the  $\phi$ -projection. A coarse assignment of  $\eta$  is nevertheless possible by determining which chambers were crossed by the track. In most cases an even more refined  $\eta$ -value can be assigned using the information from the  $\theta$ -superlayers. Both for the DTTF and the CSCTF, the track finder logic fits into high-density FPGAs. For the regional trigger the DT chambers are organized in sectors and wedges. There are twelve horizontal wedges parallel to the beams. Each wedge has six  $30^\circ$ -sectors in  $\phi$ . The central wheel has  $2 \times 12$  half-width sectors, whereas the four outer wheels are subdivided in 12 full-width sectors each. In the two endcaps the track finding is partitioned in  $2 \times 6$   $60^\circ$ -sectors. In the overlap region between the DT and CSC chambers, around  $|\eta| \approx 1$ , information from both devices is used.

In the DTTF the track finding in  $\phi$  is performed by 72 sector processors, also called Phi Track Finders (PHTF). Per chamber they receive at most two track segments from the DT local trigger through optical links. The segment information is composed of the relative position of the segments inside a sector, its bending angle and a quality code. If there are two segments present in a chamber, the second one is sent not at the bunch crossing from which it originated but at the subsequent one, provided that in that crossing no other segment occurred. A tag bit to indicate this *second track segment* status is therefore necessary. The sector processors attempt to join track segments to form complete tracks. The parameters of all compatible segments are pre-calculated. Extrapolation windows, which are adjustable, are stored in look-up tables. Muon tracks can cross sector boundaries, therefore data are exchanged between sector processors and a cancellation scheme to avoid duplicated tracks has to be incorporated.

The track finding in  $\eta$ , with the goal to refine  $\eta$ -values, is performed by 12  $\eta$  assignment units, also called Eta Track Finders (ETTF). A pattern matching rather than an extrapolation method is used, since for muon stations 1, 2 and 3 the  $\eta$ -information coming from the DT lo-

cal trigger is contained in a bit pattern representing adjacent chamber areas. The tracks in  $\eta$  are matched with those of the  $\phi$ -projection, if possible. For each wedge, the combined output of the PHTFs and the ETTFs, which consists of the transverse momentum including the electric charge, the  $\phi$ - and  $\eta$ -values and quality for at most 12 muon candidates corresponding to a maximum of two track candidates per  $30^\circ$ -sector, is delivered to a first sorting stage, the Wedge Sorter (WS). There are twelve of these sorters, which have to sort at most 144 candidates in total by  $p_T$  and quality. Suppression of remaining duplicate candidates found by adjacent sector processors and track quality filtering is also performed by these units. The two highest-rank muons found in each WS, at most 24 in total, are then transmitted to the final Barrel Sorter (BS). The latter selects the best four candidates in the entire central region, which are then delivered to the Global Muon Trigger for matching with the RPC and CSC candidates.

The DTTF data are recorded by the data acquisition system. A special read-out unit, the DAQ Concentrator Card (DCC) has been developed. It gathers the data from each wedge, through six Data Link Interface Boards (DLI). Each DLI serves two wedges. The DTTF electronics is contained in three racks in the control room. Two racks contain six track finder crates, which each house the electronics for two wedges as well as a crate controller. There is also one Timing Module (TIM) in each of these crates to distribute the clock and other timing signals. The third rack houses the central crate containing the BS, the DCC, a TIM module and boards for interfacing with the LHC machine clock and the Trigger Control System.

As for the DTTF, the core components of the CSCTF are the sector processors. They receive, through optical links, the LCT data from the Muon Port Cards in the peripheral crates. Each sector processor receives up to six LCTs from ME1 and three LCTs each from stations ME2, ME3 and ME4. Up to four track segments are also transmitted from DT station MB2. First the data are latched and synchronized, and the original LCT information is converted to reflect global  $(\eta, \phi)$ -coordinates. Then nearly all possible pairwise combinations of track segments are tested for consistency with a single track in the processors' extrapolation units. In contrast to the DTTF, no data exchange between neighbour processors is performed. Complete tracks are assembled from the extrapolation results and redundant tracks canceled as in the DTTF. The best three muons per processor are selected and assigned kinematic and quality parameters. The  $p_T$  assignment, through SRAM look-up tables, is based on the  $\phi$ -information from up to three muon stations. The data are collected in a detector-dependent unit (DDU) for the read-out. The twelve sector processors are housed in a single crate in the counting room. This crate also contains a Clock and Control Board (CCB) similar to the ones in the local CSC trigger electronics, which distributes the clock, bunch crossing reset, bunch crossing zero and other timing signals. Over the custom-developed GTL+ backplane a maximum of 36 candidate tracks is transmitted to the forward Muon Sorter board, which determines the best four muons in the two endcaps and sends them to the GMT.

### **Global Muon Trigger**

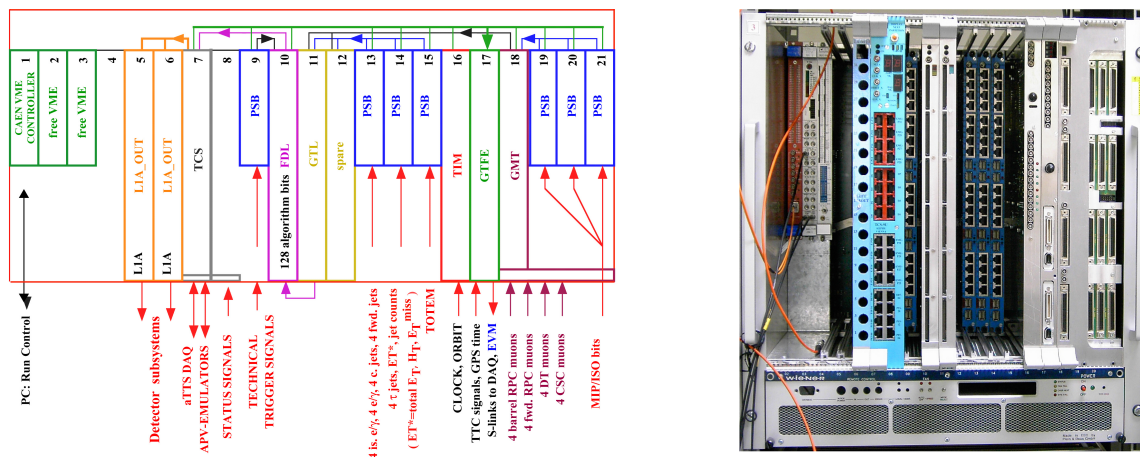
The purpose of the Global Muon Trigger [199] is to improve trigger efficiency, reduce trigger rates and suppress background by making use of the complementarity and redundancy of the three muon systems. It receives for every bunch crossing up to four muon candidates each from the DTs and barrel RPCs, and up to four each from the CSCs and endcap RPCs. The candidate information

consists of  $p_T$  and charge,  $\eta$ ,  $\phi$  and a quality code. From the GCT it also receives isolation and minimally ionizing particle bits for each calorimeter region sized  $\Delta\eta \times \Delta\phi = 0.35 \times 0.35$ . A muon is considered isolated if its energy deposit in the calorimeter region from which it emerged is below a defined threshold. DT and CSC candidates are first matched with barrel and forward RPC candidates based on their spatial coordinates. If a match is possible, the kinematic parameters are merged. Several merging options are possible and can be selected individually for all of these parameters, taking into account the strengths of the individual muon systems. Unmatched candidates are optionally suppressed based on  $\eta$  and quality. Cancel-out units reduce duplication of muons in the overlap region between the barrel and the endcaps, where the same muon may be reported by both the DT and CSC triggers. Muons are back-extrapolated through the calorimeter regions to the vertex, in order to retrieve the corresponding MIP and ISO bits, which are then added to the GMT output and can be taken into account by the Global Trigger. Finally, the muons are sorted by transverse momentum and quality, first separately in the barrel and forward regions, and then together to deliver four final candidates to the GT. A read-out processor collects the input muon data and the output record. The GMT electronics is housed in the same crate as the GT (figure 8.7). The 16 muon cables are directly connected to the GMT logic board, which has a special four VME slot wide front panel. The logic itself, which is contained in FPGA chips, only occupies one slot. The MIP/ISO bits from the GCT are received and synchronized by three Pipeline Synchronizing Buffer (PSB) input modules, which are also used in the GT. The PSB boards receive the bits via 1.4 Gbit/s serial links and are mounted at the back of the crate, behind the wide logic front panel. The MIP/ISO bits are transmitted from the PSBs to the logic board by GTL+ point-to-point links on the GT backplane.

### 8.3 Global Trigger

The Global Trigger [200] takes the decision to accept or reject an event at L1 based on trigger objects delivered by the GCT and GMT. These objects consist in candidate-particle, such as  $e/\gamma$  (isolated and non-isolated), muons, central and forward hadronic jets,  $\tau$  jets, as well as global quantities: total and missing transverse energies, the scalar sum ( $H_T$ ) of the transverse energies of jets above a programmable threshold, and twelve threshold-dependent jet multiplicities. Objects representing particles and jets are ranked and sorted. Up to four objects are available. They are characterized by their  $p_T$  or  $E_T$ ,  $(\eta, \phi)$ -coordinates, and quality. For muons, charge, MIP and ISO bits are also available.

The GT has five basic stages: input, logic, decision, distribution and read-out. The corresponding electronics boards use FPGA technology [201]. All of them, as well as the boards of the GMT, are housed in one central 9U high crate, which is shown in figure 8.7. Three Pipeline Synchronizing Buffer (PSB) input boards receive the calorimeter trigger objects from the GCT and align them in time. The muons are received from the GMT through the backplane. An additional PSB board can receive direct trigger signals from sub-detectors or the TOTEM experiment for special purposes such as calibration. These signals are called *technical triggers*. The core of the GT is the Global Trigger Logic (GTL) stage, in which algorithm calculations are performed. The most basic algorithms consist of applying  $p_T$  or  $E_T$  thresholds to single objects, or of requiring the jet multiplicities to exceed defined values. Since location and quality information is available, more



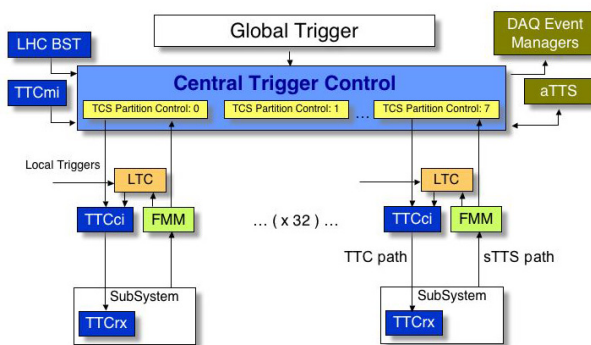
**Figure 8.7:** Global Trigger central crate.

complex algorithms based on topological conditions can also be programmed into the logic. A graphical interface [202] is used to set up the trigger algorithm menu. The results of the algorithm calculations are sent to the Final Decision Logic (FDL) in the form of one bit per algorithm. The number of algorithms that can be executed in parallel is 128. Up to 64 technical trigger bits may in addition be received directly from the dedicated PSB board. For normal physics data taking a single trigger mask is applied, and the L1A decision is taken accordingly. For commissioning, calibration and tests of individual subsystems up to eight final ORs can be applied and correspondingly eight L1A signals can be issued. The distribution of the L1A decision to the subsystems is performed by two L1A\_OUT output boards, provided that it is authorized by the Trigger Control System described in section 8.4. A Timing Module (TIM) is also necessary to receive the LHC machine clock and to distribute it to the boards. Finally, the Global Trigger Frontend (GTFE) board collects the GT data records, appends the GPS event time received from the machine, and sends them to the data acquisition for read-out.

## 8.4 Trigger Control System

The Trigger Control System (TCS) [203] controls the delivery of the L1A signals, depending on the status of the sub-detector read-out systems and the data acquisition. The status is derived from signals provided by the Trigger Throttle System (TTS). The TCS also issues synchronization and reset commands, and controls the delivery of test and calibration triggers. It uses the Timing, Trigger and Control distribution network [204], which is interfaced to the LHC machine.

The TCS architecture is represented in figure 8.8. Different subsystems may be operated independently if required. For this purpose the experiment is divided into 32 partitions, each representing a subsystem or a major component of it. Each partition is assigned to a partition group, also called a TCS partition. Within such a TCS partition all connected partitions operate concurrently. For commissioning and testing up to eight TCS partitions are available, with their own L1A signals distributed in different time slots allocated by a priority scheme or in round robin mode.



**Figure 8.8:** Trigger Control System architecture.

During normal physics data taking there is only one single TCS partition. Subsystems may either be operated centrally as members of a partition or privately through a Local Trigger Controller (LTC). Switching between central and local mode is performed by the TTCci (TTC CMS interface) module, which provides the interface between the respective trigger control module and the destinations for the transmission of the L1A signal and other fast commands for synchronization and control. The TTC Encoder and Transmitter (TTCex) module encodes the signals received from the TTCci and drives optical splitters with a laser transmitter. The LHC clock is received from the TTC machine interface (TTCmi). At the destinations the TTC signals are received by TTC receivers (TTCrx) containing low-jitter quartz PLLs. The Beam Synchronous Timing (BST) system of the LHC sends the GPS time.

The central TCS module, which resides in the Global Trigger crate, is connected to the LHC machine through the TIM module, to the FDL through the GT backplane, and to 32 TTCci modules through the LA1\_OUT boards. The TTS, to which it is also connected, has a synchronous (sTTS) and an asynchronous (aTTS) branch. The sTTS collects status information from the front-end electronics of 24 sub-detector partitions and up to eight tracker and preshower front-end buffer emulators. The status signals, coded in four bits, denote the conditions *disconnected*, *overflow warning*, *synchronization loss*, *busy*, *ready* and *error*. The signals are generated by the Fast Merging Modules (FMM) through logical operations on up to 32 groups of four sTTS binary signals and are received by four conversion boards located in a 6U crate next to the GT central crate. The aTTS runs under control of the DAQ software and monitors the behaviour of the read-out and trigger electronics. It receives and sends status information concerning the 8 DAQ partitions, which match the TCS partitions. It is coded in a similar way as the sTTS. Depending on the meaning of the status signals different protocols are executed. For example, in case of warning on the use of resources due to excessive trigger rates, prescale factors may be applied in the FDL to algorithms causing them. A loss of synchronization would initiate a reset procedure. General trigger rules for minimal spacing of L1As are also implemented in the TCS. The total deadtime estimated at the maximum L1 output rate of 100 kHz is estimated to be below 1%. Deadtime and monitoring counters are provided in the TCS. The central board sends to the DAQ Event Manager (EVM) located in the surface control room the total L1A count, the bunch crossing number in the range from 1 to 3564, the orbit number, the event number for each TCS/DAQ partition, all FDL algorithm bits and other information.