The Compact Muon Solenoid Experiment

CMS Note

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

20 March 2008

The CMS Drift Tube Trigger Track Finder

J. Erö, Ch. Deldicque, M. Galánthay, H. Bergauer, M. Jeitler, K. Kastner, B. Neuherz, I. Mikulec, M. Padrta,
H. Rohringer, H. Sakulin, A. Taurok, C.-E. Wulz

*Institute of High Energy Physics of the Austrian Academy of Sciences,
Nikolsdorfergasse 18, A-1050 Vienna, Austria*

A. Montanari, G.M. Dallavalle, L. Guiducci, G. Pellegrini

*Istituto Nazionale di Fisica Nucleare (INFN), Dipartimento di Fisica,
Viale Berti Pichat 6/2, I-40127 Bologna, Italy*

J. Fernández de Trocóniz, I. Jiménez

*Departamento de Física Teórica, C-XI, Universidad Autónoma de Madrid,
Cantoblanco, E-28049 Madrid, Spain*

Abstract

Muons are among the decay products of many new particles that may be discovered at the CERN Large Hadron Collider. At the first trigger level the identification of muons and the determination of their transverse momenta and location are performed by the Drift Tube Trigger Track Finder in the central region of the Compact Muon Solenoid experiment, using track segments detected in the Drift Tube muon chambers. Track finding is performed both in pseudorapidity and azimuth. Track candidates are ranked and sorted, and the best four are delivered to the subsequent high level trigger stage. The concept, design, control and simulation software as well as the expected performance of the system are described. Prototyping, production and tests are also summarized.

To be submitted to *JINST*

1 Introduction

The Compact Muon Solenoid (CMS) experiment at CERN, the European Organization for Nuclear Research, is designed to study physics at TeV scale energies accessible at the Large Hadron Collider (LHC). Muons are the most easily identifiable particles produced in proton-proton or heavy-ion collisions. They can be found among the decay products of many predicted new particles, for example the Higgs boson. The first selection of muons is performed on-line by the Level-1 Trigger (L1T) [1], which preselects the most interesting collisions for further evaluation and possible permanent storage by the High-Level Trigger (HLT) [2]. The L1T is a custom-designed, largely programmable electronic system, whereas the HLT is a farm of industrial processors.

The Barrel Regional Muon Trigger, also called Drift Tube Track Finder (DTTF), performs the Level-1 identification and selection of muons in the Drift Tube (DT) muon chambers [3] located in the central region of CMS. For each collision, every 25 ns for protons or every 125 ns for heavy ions, if muon candidates are present, it allocates transverse momenta, location and quality information. The latter reflects the level of confidence attributed to the parameter measurements, based on detailed knowledge of the detectors and trigger electronics and on the amount of information available. The candidates are sorted by rank, which is a function of transverse momentum and quality. The highest-ranked ones are transferred to the Global Muon Trigger (GMT) [4], which also receives muon candidates from the other two CMS muon detectors, the Resistive Plate Chambers (RPC) and the Cathode Strip Chambers (CSC) [3], and matches them with those from the DTTF. The GMT, using sophisticated algorithms, determines the four highest-ranked muon candidates in the entire CMS detector and transfers them to the Global Trigger [5], which takes the final Level-1 Trigger Accept (L1A) decision based on information from all trigger detectors.

The following section gives an overview of the Drift Tube Muon Trigger System, describing briefly the concept, the hardware and the software aspects. Section 3 gives more details on track finding and the associated track finder electronics boards, and Section 4 is dedicated to the sorting of muon candidates. Timing and synchronization are described in Section 5, followed by a section on readout. Section 7 summarizes development and testing procedures. Configuration aspects and software are described in Sections 8 and 9. The performance of the DTTF as obtained through Monte Carlo simulation is finally presented in Section 10, followed by conclusions.

2 Overview of the Drift Tube Muon Trigger System

The DTTF uses track segment information of candidate muons delivered by the local trigger electronics of the DT chambers [1,6] in the barrel region of CMS. The chambers are arranged in four muon stations (MB1, MB2, MB3, MB4 in increasing radius) embedded in the iron yoke surrounding the superconducting magnet coil. Similarly, in the endcap regions the CSCs are arranged in four stations, ME1, ME2, ME3 and MB4. ME1 is the innermost station seen from the interaction point. Each DT chamber consists of staggered planes of drift cells. Four planes are glued together and form a superlayer. The three innermost stations are made of chambers with three superlayers. The inner and outer superlayers measure the azimuthal coordinate φ in the bending plane transverse to the accelerator beams. The central superlayer measures η , the pseudorapidity coordinate along the beam direction. The fourth muon station has only φ -superlayers.

For triggering purposes, the DT chambers are organized in sectors and wedges. The layout is shown in Fig. 1. The CMS barrel iron yoke is made of five wheels along the detector axis in beam direction. Wheels are subdivided in twelve 30° sectors in azimuth. There are twelve horizontal wedges. Each wedge has five 30° -sectors in φ . For track finding purposes, the sectors of the central wheel are logically split into two parts.

In the forward regions, the CSC Track Finder (CSCTF) [7,8] determines Level-1 muon candidates. The DTTF and the CSCTF exchange information with each other from the overlap region between the barrel and the endcap muon chambers. RPCs complement both DT and CSC chambers. Compared to these chambers, they have an excellent timing resolution, but an inferior resolution in momentum and position due to large pad sizes in azimuth. They provide their own muon candidates to the GMT, which tries to match them with DT and CSC candidates, thus improving resolution and increasing the geometrical acceptance.

The local trigger electronics of the DT chambers delivers track segments (TS) in the φ -projection and hit patterns in the η -projection through optical links. It also identifies the bunch crossing (BX) to which these belong. A

Figure 1: Layout of the DTF system in the trigger chain.

segment is reconstructed if at least three out of four planes of drift cells have been hit and if the hits can be aligned. Segments in φ are first reconstructed separately in each of the two φ -superlayers. A track correlator (TRACO) then tries to match them and outputs a single φ -segment if the correlation was successful. From each muon chamber at most two segments with the smallest bending angles or, in other words, the highest transverse momenta, are forwarded to the DTF.

The tasks of the DTF are to reconstruct complete muon track candidates starting from the track segments, and to assign transverse momenta, φ - and η -coordinates, as well as quality information. The transverse momentum is calculated from the track bending in the φ -projection caused by the magnetic field in the return yoke. Using the information from this projection alone also allows a coarse assignment of η by determining which chambers were crossed by the track. The information from the η -superlayers is used to determine the η -coordinate with even higher precision. The refined η -assignment relies on track finding performed in the non-bending plane and on matching the found tracks with those of the φ -projection. The track finding in φ is performed by sector processors, also called Phi Track Finders (PTF), which use an extrapolation principle to join track segments. There are 72 PTF boards. As explained above, the sectors of the central wheel are logically subdivided into two parts. There are hence 24 boards covering the twelve sectors of the central wheel, and 48 boards covering the twelve sectors of each of the four outer wheels. Since muon tracks can cross sector boundaries, information has to be exchanged between sector processors, and a cancellation scheme to avoid duplicated tracks has to be incorporated. Tracks starting in the central wheel can either remain in it or continue to the positive-side or negative-side wheels. The twelve central wheel processors for the positive side of the interaction point are assigned to find tracks that do not exit the central wheel or that exit it on the positive side only. The twelve central wheel processors for the negative side deal with the tracks that exit at the negative side. Splitting the central sectors in two in such a way leads to a reduction in the number of required electrical connections.

The track finding in η and the assignment of refined η -values are performed by twelve η processors, also called Eta Track Finders (ETTF). Each ETTF covers one wedge. For each wedge, the combined output of the PTFs and the ETTFs, which consists of the transverse momentum including the electric charge, the φ - and η -values and quality for at most 12 muon candidates corresponding to a maximum of two track candidates per 30° sector, is delivered to a first sorting stage, the Wedge Sorter (WS). There are twelve of these sorters. The two highest-rank muons found in each WS are then transmitted to the final Barrel Sorter (BS). The latter selects the best four candidates in the

entire central region of CMS, which are then delivered to the Global Muon Trigger for matching with the RPC and CSC candidates. The DTTF data are permanently recorded by the data acquisition system. A special readout unit, the DAQ Concentrator Card (DCC) has been developed. It gathers the data from each wedge, through six Data Link Interface Boards (DLI) with two 5.7 mm diameter 8-wire halogen-free Ethernet-type cables per board. Each DLI serves two wedges. A Fast Signal Interface Card (FSC) has also been developed in order to prevent buffer overflows in case of too high trigger rates. Its functionality is explained in section 6. The FSC receives input from the DLI boards and the DCC through seven cables of the same type as the ones for the DLI-DCC connections.

All electronic modules of the DTTF are built using field programmable gate arrays (FPGA). They are located in three racks in the counting room adjacent to the CMS experimental cavern. Two racks contain six track finder crates (Fig. 2, left), which each house the electronics for two wedges as well as a CAEN VN2738 Crate Controller [9]. There are 25 input optical fibres per wedge, corresponding to 300 fibres in total. For timing purposes, there is also one Timing Module (TIM) in each of these crates. The third rack houses the central crate (Fig. 2, right) containing the BS, the DCC, the FSC, a TIM module, and electronics for interfacing with the LHC machine clock and the CMS Trigger Control System [10]. The crates are 9U $\times$ 400 mm Wiener 6023 crates, with 5V/3.3V/115A power supplies [11]. A 6U-crate for testing purposes, with a CAEN V2718 crate controller [12], is also located in this rack. For control purposes, three 2U-high Dell Power Edge 2850 PCs housed in a separate rack containing also PCs for other subsystems, are available. Each PC controls one rack.

Figure 2: Track finder crate (left), central crate (right).

On-line and off-line software to configure, operate and test the DTTF has been developed. Extrapolation and assignment look-up tables (LUTs) for the PHTFs, and η -patterns for the ETTFs have initially been generated by Monte Carlo simulation. As soon as the LHC starts its operation, they will be tuned using real muon tracks. The configuration parameters are loaded into the FPGAs using the Trigger Supervisor framework [13], a software system that controls the CMS trigger components. General monitoring software for the DTTF is provided within the CMS monitoring framework. Detailed hardware monitoring is available through a spy program, which allows to collect data independently of the central data acquisition (DAQ) system. Pattern test programs are accessible through the Trigger Supervisor for routine checks during normal operation. Specific programs for the commissioning of all electronics modules have also been developed for local use.

Some pre-production DTTF modules as well as software have been evaluated in a muon beam test at the CERN Super-Proton-Synchrotron [14, 15], and with cosmic muons during the CMS “Magnet Test / Cosmic Challenge” (MTCC) [16].

3 Track Finding

3.1 Phi Track Finder

The tasks of the Phi Track Finder system are to join compatible track segments to complete muon tracks and to assign transverse momentum, charge, location and quality parameters. The individual PHTF sector processors receive the track segments from the local trigger of the DT chambers through approximately 60 m long optical links operated in Gbit-Ethernet mode with 8bit/10bit encoding and a gross transmission rate of 1.6 Gb/s [17, 18]. The deserializers are embedded in Altera StratixGX devices [19]. Parity checking is automatic. All PHTF boards except the ones for the negative side of Wheel 0 receive four links. The DT local trigger delivers at most two TS per chamber in the φ -projection. The number of physical chambers is 250, but the trigger information for the top and bottom sectors, which are covered by two MB4 chambers in each wheel, is combined, so that there are only 240 logical chambers. A maximum of 480 TS may therefore be available. The TS information is composed of the relative position of a segment inside a sector (ϕ , 12 bits), its bending angle (ϕ_b , 10 bits) and a quality code (3 bits) which indicates how many drift cells per superlayer have been used to generate the TS. The TS of muon station MB3 contain no ϕ_b -information as the bending is always close to zero at this station due to the magnetic field configuration. If there are two TS present in a chamber, the second TS is not sent at the bunch crossing from which it originated but at the subsequent one, provided that in this next BX no other segment occurred, the probability of which is negligible. A tag bit to indicate this "second TS" status is therefore necessary. Furthermore, the BX number and a calibration bit are part of the TS information. The PHTF sector processors attempt to join track segments to form complete tracks. Starting from a source segment, they look for target segments that are compatible with respect to position and bending angle in the other muon stations. The parameters of all compatible segments are pre-calculated. Extrapolation windows, which are adjustable, are stored in LUTs. Muon tracks can cross sector boundaries, therefore data have to be exchanged between sector processors. Fig. 3 explains the basic extrapolation scheme in φ .

Figure 3: Extrapolation scheme in φ .

Each PHTF is made of dedicated units, as shown in Fig. 4. The units operate in a pipelined mode. A total of 16 BX is needed to perform all steps of the track finding.

The Input Receiver and Deserializer Unit receives and synchronizes 110 bits of data from each optical link. In a first step clock phase corrections are made by oversampling four times at 160 MHz. Then the correct BX is determined, thus compensating the differences in input cable lengths. If two TS are present in a chamber, it is also necessary to deserialize them, since they are originally sent in subsequent crossings.

Figure 4: PHTF block and timing diagram.

The next step is the Extrapolator, which determines if TS pairs originate from the same muon track. From stations 1 and 2 extrapolations to all outer stations are performed (station 1 to stations 2, 3 and 4; station 2 to stations 3 and 4). It is not possible to start extrapolations from station 3 due to the too small bending angle ϕ_b , for all possible transverse momenta. Nevertheless, a backward extrapolation from station 4 to station 3 is performed. There is also an option to make an additional extrapolation from station 2 to station 1 if the first station has too many hits due to hadron punch-through or noise.

The PHTFs exchange TS information with neighbours since tracks can cross sector boundaries. Concerning the η -projection, the PHTFs get TS information only from the PHTFs of higher η -ranges, since tracks do not significantly bend or scatter in this projection. The processors of the outermost wheels exchange TS information with the sectors of the CSC ME1 chambers in both endcaps through 72 7.5 m long 68-wire standard parallel SCSI cables connected to dedicated DT/CSC transition boards. Concerning the φ -projection, a PHTF needs track segments from the neighbouring PHTFs of its own wheel and also from those of the next wheel at higher $|\eta|$. Fig. 5 shows the pattern of neighbour connections for a group of PHTFs, which is repetitive across the DT system.

Figure 5: PHTF neighbour connections.

Every PHTF processor, except those in the central wheel (Wheel 0), forwards its input to five other PHTFs, one

previous wheel neighbour at lower $|\eta|$ in its own wedge, and two sideways neighbours in its own wheel and the previous wheel. The neighbour connections are provided by 120 2.67 m long high-density flat cables with 68 wires of 0.6 mm diameter each. Due to the large number of required neighbour connections the tasks of the central wheel are shared by two PHTFs per wedge. One of them processes muons that either remain in Wheel 0 in all stations or leave Wheel 0 in the positive η -direction. The other PHTF processes muons that leave Wheel 0 in the negative η -direction. For each TS pair there is a LUT that contains the extrapolation window depending on the ϕ_b angle (Fig. 3). The sizes of the extrapolation windows, which depend strongly on the source and target stations, can be tuned to physics requirements or experimental conditions. More details are given in section 8, and typical window sizes are shown in Fig. 18. An extrapolation is successful if the ϕ -position in the target station is inside the window predicted by the LUT. The extrapolation results are stored in 12-bit and 6-bit tables. A bit set to 1 indicates a valid extrapolation. The 12-bit tables are the results of the TS pairs that have the source in the reference wheel of the PHTF processor. The 6-bit tables belong to the TS that have the source in the next wheel. A source TS in the reference wheel can have 12 potential targets, 6 in the reference wheel and 6 in the next wheel. A source TS in the next wheel can, however, only have 6 targets in that next wheel, because a muon that left the reference wheel never returns. The total bit count of all extrapolation result tables per sector processor is 180 bits. The Extrapolator also has the possibility to filter out low quality TS, which can occur if the BX could not be correctly assigned by the local DT trigger electronics.

The next step after extrapolation is to determine which TS originate from a single muon track. It is performed by the track assembling unit, which links compatible TS into complete tracks. It starts by searching for the longest possible track. All TS used for this track are then cancelled. The procedure is repeated with the remaining TS, until no more TS can be joined. Tracks are linked by combining AND connections of extrapolation results according to a priority scheme. The output of the track assembling unit contains the addresses of each TS participating in the found track. The track address, also called an index, indicates whether a TS is coming from the same wheel as the PHTF processor or from the next wheel. The output data are sent to the pipe and selection unit. Subsets of output data are also sent to the parameter assignment unit described below, the ETTF processors and the wedge sorters. The pipe and selection unit keeps all input TS until the addresses of the two longest tracks are found. When the addresses are available, a multiplexer at the end of the pipeline selects the TS parameters of the found tracks and forwards them to the parameter assignment units.

Based on the TS parameters belonging to a track, the parameter assignment units attribute physical quantities to a track. In particular, the transverse momentum (5 bits), the sector local φ -value at muon station 2 (6 bits), the electric charge (1 bit) and the track quality (3 bits) are assigned. The p_T - and the charge assignment are based on the ϕ -value difference in the two innermost stations participating in the track. The local φ -values are obtained using the six most significant bits of the input ϕ -values at station 2. If no TS is present at station 2, the φ -value is obtained through extrapolation from the innermost TS present. The quality parameter reflects the number and combination of muon stations participating in a track [1]. The maximum quality of 7 is assigned if a muon was reconstructed from TS in all four stations.

3.2 Eta Track Finder

The geometry and the magnetic field configuration method make it impossible to derive from the η -information the physical parameters of a muon track in a standalone way. A pattern matching rather than an extrapolation method was chosen [20], since for muon stations 1, 2 and 3 the η -information coming from the DT local trigger through 60 optical links [17] is contained in a 14-bit pattern representing seven adjacent chamber areas plus one quality bit per area. If all four planes of an η -superlayer are hit, a quality bit of 1 is assigned. If only three out of four planes are hit, the quality bit is set to 0. If fewer than three planes are hit, no η -segment is considered to be found and the corresponding pattern bit is set to 0. Predefined track patterns - basically straight-line patterns, but taking into account multiple scattering and the actual magnetic field configuration - are compared with the actual hit pattern (Fig. 6). If a track in the η -projection is found, a matching with the information from the φ -projection is attempted. If a matching is possible, the rough η -value deduced from the track finding in φ is replaced by the more precise value found in η .

The patterns of possible tracks are grouped according to the geometrical features determined by the output η -values. A group contains all possible patterns belonging to the same output η -value, ordered by quality. The

Figure 6: Pattern matching scheme in η .

patterns of muons crossing more stations have higher priority. To create the patterns the ETTF hardware sets up AND conditions for the corresponding hit and quality bits. The combinations with the same priority are ORed afterwards. The highest priority pattern for a muon in each group is selected by a priority encoder. The outputs of this first level priority selection are also grouped by their positions in the η -category delivered by the PHTF units. Inside each category a new priority list is generated using the same principles as in the previous priority setup. The result of this selection is used for matching if one of the PHTFs of a wedge also found a muon in the corresponding category group. If a matching is possible, a high-precision or “fine” global η -value is assigned to the muon. It is defined at station 2 if a hit in that station belongs to the used pattern. If not, stations 3 or 1 are used, in preferential ordering. If the ETTF does not find any muon in the group where the PHTF found one, it assigns a rough global η -value and sends a “rough” tag in the output to indicate how the η of the muon was generated. If the PHTF delivers more than one muon inside one group, no matching is performed and the rough η -value is delivered. The ETTF delivers the η -values at the same time as the PHTF delivers the physical parameters of the found tracks to the Wedge Sorter. The WS can therefore handle them as a single entity.

4 Sorting

The task of the muon sorting stage is to select the four highest-rank barrel muon candidates among the up to 144 tracks received from the PHTF sector processors and to forward them to the Global Muon Trigger. Suppression of duplicate candidates found by adjacent sector processors is also performed by the sorters. Due to the partitioning of the system it is possible that more than one PHTF reconstructs two copies of the same muon candidate, which would lead to a fake increase in the rate of dimuon events. This background has to be suppressed at least to below the real dimuon rate, which amounts to about 1% of the single muon rate. The sorting and the fake track cancellation is performed in two stages: twelve Wedge Sorter boards select up to two muons out of the at most twelve candidates collected from a wedge of the barrel. One single Barrel Sorter board performs the final selection of four tracks out of the up to 24 candidates collected from the WS boards.

4.1 Wedge Sorter

As it is shown in Fig. 7, if a muon track crosses the boundaries between wheels, two neighbouring PHTFs can build the same track, since they operate independently within their own sectors. Thus, a single muon can be reconstructed twice and two muons could be forwarded to the subsequent stages of the trigger.

The Wedge Sorter receives encoded information about the position of local track segments used by the PHTF to build the tracks. Moreover, each track has a reconstruction quality attached. If two muons from consecutive sectors are found to be built with common segments, the Wedge Sorter cancels the member of the pair with the worst reconstruction quality. After the suppression of fake tracks the WS has to sort out the best two tracks among

Figure 7: Examples of duplicate track generation.

the received sample. This is done according to 8-bit ranking words, made of reconstruction quality (3 bits) and transverse momentum values (5 bits). A fully parallel one-step sorting algorithm is used.

As it is illustrated in Fig. 8, the WS receives, through a custom backplane, two muon candidates from each of the six PHTFs with their parameters coded as 24-bit words. The corresponding η -information for each candidate is received from the ETTF coded as 7-bit words through a front panel connector. The fake track suppression and the sorting are performed in two BX or 50 ns, after which the two highest-rank muon candidates are sent to the Barrel Sorter, through two low-voltage differential signaling (LVDS) links. The algorithms used for the fake suppression and the sorting can be configured in different modes through internal registers. The WS operates at 40 MHz.

Figure 8: Data flow and processing units inside the Wedge Sorter.

4.2 Barrel Sorter

The Barrel Sorter has to suppress fake tracks and select the best four candidates over the full barrel region and forward them to the Global Muon Trigger. It receives up to two muon candidates from each of the twelve Wedge Sorters through 24 2.4 m long SCSI parallel twin flat 68-wire in hose cables. The muon tracks delivered by the Wedge Sorters to the Barrel Sorter still contain the information about the track segments used in the reconstruction by the PHTF. Just like in the Wedge Sorter along a wedge, each of two adjacent PHTFs can build a candidate if a muon track crosses the boundaries between wheels or neighbouring sectors in φ . From the local φ -positions and the sector position of the PHTF output a CMS global φ -value (8 bits) is calculated. The Barrel Sorter cancels the track with the worst reconstruction quality if two muons from adjacent sectors in a wheel are found to be built with common segments.

Simulations of single muon events show that the combined fake track cancellation algorithms performed by the WS and the BS allow to limit the fake dimuon rate to a level of 0.3%, as is shown in Tab. 1.

Table 1: Dimuon fake rate after duplicate track suppression performed in WS and BS.

	PHTF output	WS output	BS output
Dimuon fake rate	27%	8%	0.3%

After suppression of fake tracks the BS has to sort the four highest-rank tracks out of the possible 24 candidates received from the twelve Wedge Sorters. This is again done according to 8-bit ranking words made of reconstruction quality and transverse momentum values.

The input track data consist of 31 bits, while the output track data are 32-bit words. The full algorithm input/output bit count is 872. The latency for duplicate track cancellation, sorting and multiplexing operations is limited to 3 BX or 75 ns by running at 80 MHz internally. The best four candidates are then sent through LVDS links to the Global Muon Trigger. Different modes of operation are configurable through internal registers, and the data flow inside the BS can be monitored through spy registers. For test purposes, a standalone trigger signal can be delivered by the BS, according to trigger conditions that can be set through configuration registers.

5 Timing and Synchronization

The LHC machine broadcasts its 40.08 MHz bunch-crossing clock and 11.246 kHz orbit signals through the TTC (Timing, Trigger and Control) network [21] with high-power laser transmitters over single-mode optical fibres to the experiments. Each of the DTTF crates has a timing module to distribute the clock to the individual DTTF boards, which are equipped with clock receivers and a multichannel clock distribution system. The core of this system is a sophisticated phase-locked loop (PLL) clock chip (Cypress RoboClockIITM CY7B994V [22]) with several grouped clock outputs. The sub-units of the boards are individually clocked by these clock output lines. The PLL clock chip allows to determine different clock phase and delay values for each group, which makes it possible to choose optimal values for the input links and also for the data transmission between system sub-units. The clock chip output groups are controlled by the clock control lines of the Controller FPGA, driven by the clock control registers. Their delay and phase values are programmable.

In order to check that muon tracks are correctly assigned to the bunch crossing from which they originated, the bunch crossing zero (BC0) signal is sent together with the data. The orbit gap position is compared to the BC0 signal contained in the data. In the Barrel Sorter it is also possible to detect any synchronization misalignment among the twelve Wedge Sorters. A VME error register can be read out. In addition to the clock, the TIM modules also send the BC0 signal, the bunch counter reset (BCRes) signal and the Level-1 Trigger Accept decision to the DCC board and to the DAQ and spy modules contained in many of the blocks of the DTTF.

The overall latency of the DTTF system, from the input to the optical receivers to the output of the BS is 29 BX. An additional 3 BX are needed to transfer the data from the BS to the Global Muon Trigger. Changes in latency should not occur due to the rigid pipelining. However, such a change will be immediately discovered from the data themselves, through the monitoring. The data of the bunch crossings previous and subsequent to the triggered BX are also stored in the DAQ record for diagnostic purposes.

6 Readout

The DTTF sends data to the CMS DAQ system for readout. Each PHTF and ETTF FPGA contains local DAQ blocks, from which the data are sent as a TTL serial bit stream through an LVDS interface to the Data Link Interface board of each track finder crate, and then forwarded to the DTTF readout board, the DAQ Concentrator Card, via Channel LinkTM connections [23]. This board houses an interface, from which the data are sent to the central DAQ system. The interface, link and transmission protocol S-link64 have been developed at CERN [24]. The DTTF readout scheme is shown in Fig. 9. The PHTFs have three local DAQ blocks, which contain the input track segments, the track addresses, and the output track parameters, respectively. The ETTFs have only two local DAQ blocks, which contain the input pattern bits and the output η values.

Figure 9: DTTF readout scheme.

The DTTF data record is composed of all input and output signals from the triggered BX, its predecessor and its successor. Headers and trailers, including a cyclic redundancy code to detect data transmission errors and the record length, are added. Each triggered event contains 53244 bits of input and output data, composed of 3×214 bits per PHTF and 3×195 bits per ETTF. At the maximally allowed Level-1 trigger rate of 100 kHz this would amount to a data rate of 5.32 Gbit/s or 665 MByte/s. The DAQ system only allows 2 kByte of data for each DTTF Level-1 event on average, which is equivalent to a bandwidth of 200 MByte/s. Therefore a data compression has to be performed. Simulations have confirmed that a simple zero suppression scheme is adequate. The input data blocks are split into data words according to the contents of the local DAQ blocks explained above. If the input and output data in both Track Finders contain zero quality, the data word is considered as null data and is suppressed. If the PHTF track address word has all bits set to one it will also be considered as a null track. The ETTF input hit pattern data words are treated as null data if no hit appears, which is equivalent to all bits set to zero. The DCC compresses the data blocks in real time. The final event size per triggering muon ranges between 512 and 640 bits.

A mechanism has been developed in order to prevent buffer overflows in case of too high trigger rates. The derandomizer buffer depths of the local DAQ blocks are dimensioned such that on average an overflow would occur not more often than once every 27 hours. The DCC board emulates the status of these buffers. If it finds that 75% of buffer space is filled, a warning signal is issued to the Fast Signal Interface Card, which in turn sends it to the central Trigger Control System. The latter then initiates the application of predefined trigger throttling rules to avoid the loss of events, such as applying prescale factors.

7 Development and Tests

7.1 Prototyping

The development of the DTTF system was constantly followed by prototyping [25]. Since the main and most complicated module in the DTTF is the PHTF processor, it was the subject of most of the prototyping activity. The development of prototypes of the other units was easier due to smaller degree of complexity, and also due to the experience gained during PHTF prototype building and testing.

The long development era of the PHTF unit comprised four prototypes: (i) Simulation Mapping (SM) prototype; (ii) Technology Evaluation (TE) Prototype; (iii) Functionality Evaluation (FE) prototype; and (iv) Pre-Production (PP) prototype.

The SM prototype was designed after performing the first physics simulations. The goal was to understand if the simulated features of the PHTF could be materialized in hardware. Even if this first prototype did not fulfill the

requirements, it established the most important principles for a realistic implementation.

The TE prototype, finished in the year 2000, was an intermediate step towards the next generation of prototypes. It was not foreseen to perform track finding functions, but clarified several important open questions. The goals were:

- to test the available FPGAs and their handling;
- to test inter-FPGA data paths and their usability for track finding purposes;
- to use it as a test bench for data link tests;
- to perform tests to establish a feasible data spy structure;
- to perform control system tests for efficient system control setup.

The TE prototype was designed as a 6U standard VME board with two FPGAs. The first one was programmed as VME controller and interface; the second one contained a moderate size on-chip memory and access to two mezzanine cards through two 64-bit wide strobed data buses.

During the TE prototype tests, all design goals were achieved. The conclusions were:

1. Modern FPGAs are powerful and flexible enough to serve as a base for the final DTF implementation. The configuration is straightforward and easy. Their on-chip memory can be programmed for LUTs.
2. Parallel LVDS links can serve as internal connections between the different boards.
3. VME control can be achieved with a single J1 Connector A24/D16 connection.
4. The board-internal data buses can be driven at 40.08 MHz and can be as wide as 128 bits.
5. The boards need separate JTAG (Joint Test Action Group, IEEE 1149.1 Standard Test Access Port and Boundary-Scan Architecture) chains for FPGA configuration and other purposes.
6. Spy functions can be programmed in FPGAs and can be accessed via a secondary JTAG chain.
7. The clock distribution requires special attention.

The FE prototype exclusively used VHDL behavioral code to generate the FPGA programs. The VHDL code had to be partitioned in several blocks in order to fit the size of the available FPGA chips: six Input Receivers, two Extrapolators, the Track Assembling Unit, three Pipe and Selection Units, and two Parameter Assignment Units. The main constraint was the number of available I/O pins, limited by the soldering technology that was accessible. The 9U board housed 15 FPGAs, 9 EPC2 and 3 EPC8 configuration circuits. Except for the controller FPGA, all these participated in the primary JTAG chain, which had 27 members.

The FE prototype, finished in 2002, demonstrated the track finding functionality for the first time, using simulated TS inputs. This is shown in the oscilloscope screenshot in Fig. 10, where Trace #1 is the 40 MHz clock. Pulses in Trace #2 (#3) represent the input TS signal (the output of the Extrapolator, respectively). Pulses in Trace #4 represent the PHTF output muons, arriving 425 ns (17 BX) after the TS input.

However, the construction, debugging, and tests of this prototype revealed a number of problems, most of them connected to the complexity of the board, the number of internal connections, the amount of components, and their development.

For the PP prototype a new approach was taken. Most of the FPGAs were merged into one modern very highly integrated circuit. This approach was named “System-on-Chip” or SOC. The design of the PP prototype started with the development of the VHDL code for a large scale FPGA. I/O constraints showed that it was not practical to include the Input Receiver functionality in the SOC FPGA. The SOC FPGA, an Altera Stratix EP1S40 with 1020

Figure 10: Oscilloscope screenshot demonstrating the track finding functionality of the FE PHTF prototype. Signals are explained in the text.

pins, was soldered on a mezzanine board plugged onto the 9U VME motherboard. This allows exchanging the mezzanine board in case of problems with the FPGA and simplifies considerably the layout of the motherboard. Overall risk and cost were significantly reduced due to this technique. Apart from these changes, the PP prototype design preserved all the FE prototype solutions that proved reliable and straightforward. The PP was developed within ten months and available in July 2004.

7.2 Production

The first PHTF boards of the final version were available in summer 2005. Production and quality control of all DTTF boards were completed by April 2006. Figs. 11 (left), 11 (right), 12 (left), and 12 (right) show pictures of the final version of one PHTF, one ETTF, one WS, and one BS, respectively.

Figure 11: PHTF (left) and ETTF (right) production boards.

Fig. 13 shows one DTTF rack with its three track finder crates fully populated with production boards, installed at the CMS cavern. The central crate is shown in Fig. 14. Installation at the CMS underground counting room, commissioning with cosmic muons, and integration with the rest of the CMS L1 Trigger are well advanced.

Figure 12: WS (left) and BS (right) production boards.

Figure 13: Three DTTF crates fully populated with production boards, without most of the front panel cables and connectors attached: PHTF (yellow), ETTF (green), WS (gray), TIM and DLI (blue), and Crate Controllers (red).

Figure 14: Central crate.

7.3 Tests

In October 2004 the DTTF behavior was studied in a test beam at CERN. All PHTF features, except track assembling, were validated. The experimental configuration consisted of two DT chambers (one MB1 and one MB3), equipped with readout and trigger electronics, one DT Sector Collector, one PHTF, one TIM, and one WS prototype boards. Input and output PHTF information were recorded at 40 BX slots around the trigger using two Pattern Units [26]. Full information about all steps of the PHTF track finding process in 10 BX slots around the trigger was accessed using Spy DAQ (Section 9.1).

A full account of the results of the 2004 beam test analysis may be found in [14, 15]. The most important results are outlined below.

Figure 15: Several PHTF occupancy levels as a function of the BX number: PHTF input TS (black histogram and triangles), coincidences (red or light-grey), and PHTF output tracks (solid blue or dark-grey). The figure originates from data taken at the muon test beam in 2004.

In Fig. 15 the black histogram and triangles show the PHTF input occupancy as a function of BX for the MB1 and MB3 chambers, respectively. The correct trigger BX number is 24. One can observe high efficiency at the right BX number, at the cost of a significant (order 10-15%) component of ghost TS at the wrong BX number. The red or light-grey distribution shows the effect of a conventional coincidence analysis; the level of ghost coincidences at the wrong BX numbers has been reduced by about a factor 2.5.

The difference between the red or light-grey and the solid blue or dark-grey histograms illustrates the importance of the original PHTF extrapolation approach. The ghost component is reduced by an additional factor 6, to always below the 1% occupancy level. The flat input component at the level of 0.2% corresponds to real out-of-time test beam muons.

Fig. 16 (left) illustrates the PHTF extrapolation principle in the $\phi_b - \Delta\phi$ plane for test beam muons. The reconstructed PHTF tracks (red or light-grey points) lie in the narrow window allowed by the extrapolation LUTs (blue or dark-grey lines). The black points are wrong associations and are not reconstructed by the PHTF extrapolation algorithm. Fig. 16 (right) illustrates the PHTF p_T -assignment technique in the $p_T - \Delta\phi$ plane. The black points are test beam muon events. The green or light-grey (red or dark-grey) segments represent the high- p_T (low- p_T) p_T -assignment LUT values. Units of p_T are CMS trigger p_T bins. The switch over between the two regions is at

Figure 16: PHTF extrapolation (left) and p_T -assignment (right) principle illustration, using test beam muons. Symbols are explained in the text.

about $p_T = 17$ GeV/c, corresponding to a physical $p_T = 20$ GeV/c.

The data analysis results showed in all cases excellent agreement with the design performance requirements, in particular 98% efficiency to reconstruct tracks at the right BX, and the expected ghost rejection power at the wrong BX [14, 15].

Figure 17: (Left) Oscilloscope screenshot showing input and output PHTF signals at MTCC. Signals are explained in the text. (Right) Typical cosmic muon extrapolation correlation (MB1 to MB2) at the MTCC.

In August 2006, the DT Trigger, including the DTF, has provided a 3-sector cosmic trigger for CMS at the MTCC. The DTF hardware included 3 PHTFs, one TIM, one WS, and one BS production board. The BS TTL output line defined the CMS L1A signal. The software setup included Spy DAQ, on-line monitoring, and a C++ bit level emulator program (Section 9).

Fig. 17 (left) shows an oscilloscope screenshot of one DTF triggered cosmic event at the MTCC. Pulses in Tracks #1, #2, and #3 represent synchronized PHTF input TS at stations 1, 2, and 3, respectively, in Wheel 2 Sector 10. The pulse in Track #4 represents the BS L1A signal, after 29 BX (from the PHTF input to the BS output).

At the MTCC, a total of 25 million events, at 0 T and 4 T magnetic field, with the electromagnetic and hadronic

calorimeters and the tracker in the readout, have been offered to the Collaboration. Also, integration tests of the DTTF with the CSCTF and the GMT, at the electronics and physics levels, were performed successfully.

From the DTTF point of view, the MTCC was the first opportunity to validate triggers coming from long tracks (track assembling), tracks changing wheels and/or sectors, and dimuons. For this purpose a sample of 2.5 million events was accumulated using Spy DAQ. Fig. 17 (right) shows the extrapolation correlation for cosmic muons in stations 1 and 2 of Wheel 2 Sector 10. The analysis of the data has shown perfect agreement of the hardware performance with the expected behavior in all cases.

8 Configuration

In order to fulfill the requirements, the DTTF boards have to be configured appropriately. The connection between the DTTF electronic model and physics is implemented via look-up tables in the PHTF boards, and via η -patterns in the ETTF boards.

There are two instances of LUTs in the PHTF boards: extrapolation LUTs and assignment LUTs. The PHTF extrapolation LUTs implement the PHTF extrapolation scheme discussed in Section 3.1. Extrapolations between the following station pairs are performed: MB1 to MB2, MB3, MB4, and ME1; MB2 to MB1, MB3, MB4, and ME1; and MB4 to MB3. The ϕ_b and ϕ resolution is 8 bits, corresponding to the 8 most significant bits of the 12-bit φ -values. For every extrapolation there are two LUT files, one containing the upper limit of the extrapolation window, and the second the lower limit. The total size of a one extrapolation LUT file is 1.5 kB. The total size of the extrapolation LUTs stored in one PHTF is 27 kB.

Figure 18: Typical extrapolation (left) and p_T -assignment (right) LUTs generated with ORCA simulated muons.

Two kinds of extrapolation windows have been generated and used up to now. The first set was generated using ORCA [27] simulated single muon events with a transverse momentum distribution between 3 and 100 GeV/c, and flat distributions in azimuth and pseudorapidity. LUT windows were calculated at fixed extrapolation (99%) efficiency for events with a TRACO-correlated TS in the source station. Since the size of the extrapolation window should represent a balance between muon trigger efficiency and background rejection, it was found convenient to explicitly link the calculation to a meaningful physics parameter such as efficiency. Fig. 18 (left) shows a typical extrapolation LUT in the $\phi_b - \Delta\phi$ plane (MB1 to MB2). Points are ORCA simulated muons, the white (blue) area is the allowed (forbidden, respectively) extrapolation region. At CMS, also TRACO-uncorrelated TS will be used. The overall extrapolation efficiency design goal was targeted at 94%. The ghost background rejection power should be of the order of 10.

Simulated LUTs downloaded into the hardware have been used at all instances of prototype testing and production quality control, checking the hardware performance with large samples of single- and di-muon simulated events. In addition, the same LUT windows were used at the 2004 beam test.

A second set of LUTs with maximally opened extrapolation windows has been produced and used to test special configurations of the DTTF system. The open LUTs are important in order to accumulate unbiased samples of muon and background data. A first important example has been the CMS DT cosmic trigger at the MTCC. Even more important will be the ability to accumulate unbiased samples of guaranteed muon data, from J/ψ , W , and Z decays, during the first months of CMS data taking. The first DTTF trigger configuration is expected to be “trigger on every muon”, implemented via open LUTs. Using these unbiased samples and given the size of the measured background levels, one of the first tasks will be to compute the physics extrapolation LUTs.

The PHTF boards use LUTs for muon output φ position and p_T assignment. The p_T LUTs assign transverse momentum at fixed trigger transverse momentum cut efficiency (90%). For p_T -assignment, six extrapolations are used in preferential ordering (MB1 to MB2, MB3, MB4; MB2 to MB3, MB4; MB4 to MB3). For every extrapolation, two p_T -assignment LUT files are calculated: one for low- p_T muons and the other for high- p_T muons. The separation between the two is decided according to the value of ϕ_b at the source station. The p_T granularity is reported on a 5-bit non-linear scale. The size of one p_T -assignment LUT file is 6 kB. The total size of the p_T -assignment LUT files stored in a PHTF is 72 kB.

The φ -assignment LUTs map the individual PHTF local ϕ -coordinate at station MB2 into the CMS global φ -coordinate. For PHTF muon candidates with a TS in station MB2, the mapping is direct using 10-bit ϕ resolution. For muon candidates with no actual TS in station MB2, extrapolation from stations MB1 or MB4 (in preferential ordering) has to be performed first. The size of one file is 6 kB. The total size of the three PHTF φ -assignment LUT files is 18 kB.

Fig. 18 (right) shows a typical high- p_T -assignment LUT in the $p_T - \Delta\phi$ plane (MB1 to MB2). Simulated p_T LUTs are expected to be used during the first months of CMS running. The DTTF is not cutting directly using the assigned p_T values. In CMS, all trigger cuts are implemented at the Global Trigger. Using unbiased samples of guaranteed muons and the p_T values measured at the CMS Tracker, the actual physics LUTs will be calculated.

In principle, the PHTF granularity implies that different extrapolation and assignment LUTs can be used for every different DT wheel and sector. Simulation studies have shown that the effect of a misaligned muon detector will be negligible at the DT trigger level [28], therefore it is not expected that the sector degree of freedom will have to be used. Use of different sets of extrapolation and, especially, p_T -assignment LUTs as a function of the wheel number is not excluded, as effects of backgrounds, dead material, and magnetic field depend mostly on pseudorapidity. The final need of this will be evaluated using the data.

The ETTF boards use η -patterns to find tracks in the CMS non-bending plane. The ETTF η -patterns are not stored in LUTs, but directly embedded in the VHDL η -pattern finding logic. The total number of possible η -patterns in the hardware is 2346. The assigned η granularity is 6 bits, in the pseudorapidity region -1.2 to 1.2. If no η -pattern is found, the DTTF muons are still assigned a rough value of η , based on the position where the muon crossed wheels. There are 12 different such cases per wheel. The rough value of η in each case is at the center of the corresponding pseudorapidity window.

Simulation results show that the probability that an actual η -pattern does not appear in the list is smaller than 0.1%; this number is negligible compared to the fraction of muons with rough η -assignment due to DT inefficiencies (intrinsic or geometrical), which is 10%. The effect on the allowed set of η -patterns from muon detector misalignment is also negligible.

9 Software

9.1 On-line Software

The DTTF on-line software provides applications to control, monitor and test all modules of the system. The software is developed in the CMS XDAQ on-line software framework [2], which features a client-server architecture. Besides other things the XDAQ framework provides applications with services to communicate amongst each other in a distributed environment. Control of the applications is provided via a SOAP server embedded in the framework. On reception of SOAP control messages the server invokes specific callback methods in the applications, which then execute the code necessary to control and operate the hardware. In CMS VME hardware is

generally accessed from Linux PCs via a PCI to VME bridge. The XDAQ framework contains libraries providing high level access to VME hardware. These libraries make the XDAQ applications independent of the specific hardware used to access the VME modules. On the DTTF client side, the software has evolved from customized Java/XML Graphical User Interfaces to integration in the CMS Trigger Supervisor framework [13], along with all the other CMS trigger subsystems.

The control and testing programs of the DTTF system include four generations of growing complexity: test bench, beam test, MTCC and DT commissioning, and, finally, CMS. Access to all input, output, and intermediate DTTF signals is mandatory. At the hardware level, the solution adopted was serial (JTAG) access. The PHTF and ETTF boards contain three JTAG chains. The first two chains are the standard built-in ones, for FPGA programming and Boundary Scan purposes, respectively. In addition, a third JTAG chain implements the more sophisticated spy system. Spying allows synchronous and triggerable access to the data. After triggering, sequential access to VME registers allows reading the data stored in every spy block. From the software point of view, the spy system provided the basis of the dynamic functionality tests, and, ultimately, of the local DAQ system of the DTTF.

On top of the hardware layer, on-line control and test software has been developed. During the PHTF and ETTF prototyping phases, several generations of control and test software were produced. The complexity of the actions performed on the system and the information accessed grew in parallel to the software complexity.

The first generation software coincided in time with the tests of the PHTF FE prototype. It allowed static tests of the board integrity, and of the muon-finding functionality using Boundary Scan. In addition, the simplest dynamical functionality tests were performed.

The second generation of software was developed for the ETTF FE prototype, and included a first version of a local DAQ, called Spy DAQ. Dynamic functionality tests were based on large samples, using simulated input data. The User Interface included routines for visualization and statistical treatment of the data. A thorough study of the performance allowed full debugging of the VHDL behavioral description.

Software of the third generation was developed for the PHTF PP prototype. The test bench studies of the second generation were repeated, resulting again in full debugging of the VHDL behavioral model. A standard protocol to test the production cards was established and executed. This generation contained also test programs for VME readout of the DCC unit (emulating a simplified version of the CMS Central DAQ).

Software of the fourth generation included, in addition, on-line operation and monitoring of the DTTF system as a whole during the beam test in 2004, the MTCC in 2006, and DT commissioning in 2007. This generation has allowed integrated control of several DTTF processors, as well as the Sorters, and all other DTTF modules. The Barrel Sorter spy information was incorporated into the local DAQ and monitoring programs. Working versions of all the programs expected to run at the CMS Control Room have been prepared, debugged, and operated for many months.

9.2 Simulation software

A full C++ simulation of the DTTF system has been produced, using object-oriented programming.

The goal of the simulation is twofold. First, it allows event-by-event comparison of the hardware results versus the C++ simulated ones. In this mode, the simulation works as a DTTF emulator. The C++ emulator performance was validated at the DTTF design phase against the VHDL simulation, allowing debugging and matching of both descriptions of the system. During the prototyping and production phases, thorough testing and quality control of the produced hardware were based on C++-simulated events. Further debugging at a finer level of the C++ emulator has been performed using real data at the 2004 beam test and MTCC. At CMS, the emulator will be used as a most useful monitoring tool, checking on an event-by-event basis the decision of the DTTF trigger against the emulated one, based on real DT Local Trigger input data. At least at the beginning of the experiment, the goal is to check a sizable fraction of L1-accepted triggers as part of the data quality monitoring.

Fig. 19 shows a comparison of several characteristics of the actual and emulated PHTF output at the MTCC: first output track quality (left) and its assigned p_T (right). The comparison agreement in the variables of the first output track is 100%. The general agreement found is 99.98%.

Figure 19: Comparison of several characteristics of the actual and emulated PHTF output at the MTCC: first output track quality (left) and its assigned p_T (right).

Second, the C++ description has been used to generate first versions of the DTF extrapolation and assignment LUTs; to compute the expected performance of the DTF trigger for generic single- and di-muon events [29], and in feasibility studies of the most relevant physics channels at the LHC [30]. Last but not least, the simulation will be an important part of the CMS physics analysis at the LHC.

The DTF C++ simulation has been integrated in the CMS simulation and reconstruction packages ORCA [27] and CMSSW [31].

10 Muon Trigger Performance

In this section, the expected DTF performance highlights, computed using simulated events, are presented [15]:

- **Efficiency:** The combined effects of efficiency and transverse momentum resolution are summarized in the 'turn-on' curves in Fig. 20, where the DTF single muon trigger efficiency is represented as a function of the muon transverse momentum, for different transverse momentum thresholds. The plateau efficiency is 95%, only limited by the DT detector geometrical acceptance.

Figure 20: The DTF trigger efficiency as a function of the muon transverse momentum, for different p_T thresholds.

- **Momentum resolution:** From the turn-on curves an overall transverse momentum resolution of 14% is obtained, dominated by dead material effects in the barrel iron yoke. The DT chamber intrinsic resolution contribution, measured at the 2004 beam test, is 5%.
- **Position resolution in φ :** Resolution in φ is dominated by the PHTF output bin size (0.02 rad). The dimuon resolving power is limited by the TS position resolution.
- **Position resolution in η :** Resolution in the η -coordinate is 0.03 (0.08) for fine (rough) assignment.

Figure 21: Integrated inclusive muon DTF trigger rates at nominal LHC luminosity, as a function of the muon p_T threshold.

- **Trigger rate:** The accumulated inclusive muon DTF trigger rate at nominal LHC luminosity, as a function of the transverse momentum threshold, is shown in Fig. 21.

The black curve shows the p_T spectrum for muons in the pseudorapidity region $|\eta| < 1.04$, generated in PYTHIA [32] minimum bias events; the red (dark-grey) curve shows the rate at the DTF input. Muons with $p_T = 0$ GeV/c correspond to calorimeter punch-through events and represent about 30% of the total rate at the DTF input. Almost all muons with $p_T < 4$ GeV/c are stopped before reaching the MB2 station of the DT detector, and muons with only MB1 hits cannot be reconstructed as DTF segments. For $p_T > 5$ GeV/c the size of the ghost component is similar to the one for real muons. The green (light-grey) points represent the DTF output rate as a function of the real muon p_T . The trigger is fully efficient for muons with $p_T > 5$ GeV/c. The punch-through and ghost components have been reduced to negligible levels. Finally, the blue (dashed) curve shows the DTF output rate as a function of the DTF assigned p_T . Output transverse momenta are assigned for 90% cut efficiency. At low-, intermediate-, and high- p_T the rate is dominated by real muons (decays in-flight, heavy quark and vector boson decays, respectively). For a transverse momentum threshold of 15 (25) GeV/c the expected rate is about 6 (2) kHz.

11 Conclusions

The CMS Drift Tube Trigger Track Finder has been presented. Its rationale, modus operandi, hardware implementation and software tools have been described. Finally, the system performance has been summarized – the actual one achieved at beam tests and at the MTCC, and the one expected at the LHC.

Production of the full DTTF system electronics is complete. Installation at the CMS underground Counting Room, commissioning with cosmic muons, and integration with the rest of the CMS L1 Trigger are well-advanced, to be ready for the first LHC data.

Acknowledgments

The support by the Austrian Federal Ministry for Education, Science and Culture and the Ministerio de Educación y Ciencia of Spain, under project number FPA2005-00780, is acknowledged. Ch. Deldicque, J. Erö, J. Fernández de Trocóniz, I. Jiménez, H. Rohringer, H. Sakulin and C.-E. Wulz are grateful for grants by the Austrian Exchange Service and the Ministerio de Educación y Ciencia through the Acciones Integradas program, under project numbers 21/2004 and HU 2003-0012. Contributions from the electronics laboratory of INFN Bologna and from RWTH Aachen are gratefully acknowledged. Initial work on the DTTF concept has been performed by doctoral students A. Kluge and T. Wildschek. The CMS-internal reviewers D. Acosta, Ch. Schwick, D. Holmes and S. Erhan are thanked for their detailed advice.

References

- [1] CMS Collaboration; The TriDAS Project – The Level-1 Trigger Technical Design Report, **CERN/LHCC 2000-38** (2000).
- [2] CMS Collaboration; The TriDAS Project – Data Acquisition and High-Level Trigger Technical Design Report, **CERN/LHCC 2002-26** (2002).
- [3] CMS Collaboration; The Muon Project, Technical Design Report, **CERN/LHCC 97-32** (1997).
- [4] H. Sakulin et al., “Implementation and Test of the First-Level Global Muon Trigger of the CMS Experiment”, Proc. of the 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany, 12-16 Sep. 2005.
- [5] M. Jeitler et al.; “The Level-1 Global Trigger for the CMS Experiment”, Proc. of the 12th Workshop on Electronics for LHC Experiments, Valencia, Spain, 23-27 Sep. 2006, **JINST 2 P01006** (2007).
- [6] P. Arce et al.; Nucl. Instr. and Meth. **A534** (2004) 441.
- [7] D. Acosta et al.; “Performance of a Pre-Production Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon Systems”, Proc. of the 10th Workshop on Electronics for LHC and Future Experiments, Boston, USA, 13-17 Sep. 2004.
- [8] D. Acosta et al.; Nucl. Instr. and Meth. **A496** (2003) 64.
- [9] <http://www.caen.it/nuclear/product.php?mod=VN2738>.
- [10] The CMS Trigger and Data Acquisition Group, CERN CMS-Note **2002-033** (2002).
- [11] <http://www.wiener-d.com/>.
- [12] <http://www.caen.it/nuclear/product.php?mod=V2718>.
- [13] I. Magrans de Abril, C.-E. Wulz, and J. Varela; IEEE Trans. Nucl. Sci. **Vol. 53 Nr. 2** (2006) 474.
- [14] M. Aldaya et al.; Nucl. Instr. and Meth. **A579** (2007) 951.

- [15] I. Jiménez; Diploma Thesis, Universidad Autónoma de Madrid, 2006.
- [16] CMS Collaboration; CERN CMS-Note **2007-005** (2007).
- [17] L. Guiducci et al.; “DT Sector Collector electronics design and construction”, Proc. of the Topical Workshop on Electronics for Particle Physics, Prague, Czech Republic, 3-7 Sep. 2007.
- [18] <http://proj-gol.web.cern.ch/proj-gol/>.
- [19] <http://www.altera.com/products/devices/stratixgx/sgx-index.jsp>.
- [20] M. Brugger, M. Fierro, and C.-E. Wulz; Nucl. Instr. and Meth. **A482** (2002) 254.
- [21] B. G. Taylor, “Timing Distribution at the LHC”, Proc. of the 8th Workshop on Electronics for LHC and Future Experiments, Colmar, France, 9-13 Sep. 2002.
- [22] <http://www.cypress.com/products/index.jsp?fid=27&rpn=CY7B994V>.
- [23] <http://www.national.com/pf/DS/DS90CR215.html> (transmitter),
<http://www.national.com/pf/DS/DS90CR216.html> (receiver).
- [24] H. C. van der Bij et al.; <http://hsi.web.cern.ch/HSI/s-link/>.
- [25] H. Bergauer et al., J. Erö (editor); Report submitted to GT/GMT/DTTF Review Committee, Vienna, Austria, 2004; http://wwwhephy.oeaw.cat.at/p3w/cms/trigger/Review/Information/DTTF_Description_1.0.pdf.
- [26] G. M. Dallavalle et al.; “Pattern Unit For High Throughput Device Testing”, Proc. of the 4th Workshop on Electronics for LHC Experiments, Rome, Italy, 21-25 Sep. 1998.
- [27] CMS detector reconstruction software group; <http://cmsdoc.cern.ch/orca>.
- [28] I. Belotelov et al.; CERN CMS-Note **2006-017** (2006).
- [29] CMS Collaboration; CMS Physics Technical Design Report, Volume I, Detector Performance and Software, **CERN/LHCC 06-01** (2006).
- [30] CMS Collaboration; CMS Physics Technical Design Report, Volume II, Physics Performance, J. Phys. G: Nucl. Part. Phys. **34** (2004) 995.
- [31] CMS detector reconstruction software group; <https://cmsdoc.cern.ch/cms/cpt/Software/html/General>.
- [32] T. Sjöstrand, S. Mrenna, and P. Skands; JHEP **0605** (2006) 026.