The Level-1 Global Trigger for the CMS Experiment at LHC

Presented at the 12th Workshop on Electronics for LHC Experiments and Future Experiments

M.Jeitler, A. Taurok,
H. Bergauer, C. Deldicque, J.Erö, M. Ghete, P. Glaser, K. Kastner, I. Mikulec, T. Nöbauer, B. Neuherz, M. Padrta, H. Rohringer, T. Schreiner, J. Strauss, C.-E. Wulz
Institut für Hochenergiephysik der Österreichischen Akademie der Wissenschaften,
Nikolsdorfergasse 18, A-1050 Vienna , Austria

taurok@hephy.oeaw.ac.at

Abstract

The electronics of the First Level Global Trigger for CMS electronics is described. It is the last stage of the Level-1 trigger system and decides for every LHC bunch crossing whether to reject or to accept a physics event for further evaluation by the High Level Trigger. The Global Trigger receives trigger objects from the Global Calorimeter Trigger and the Global Muon Trigger and applies in parallel up to 128 physics trigger requirements, so-called 'Algorithms'. In addition, up to 64 so-called 'Technical Trigger' signals can be used to either accept or reject events. The Algorithm and Technical Trigger bits are then combined to a Final_OR signal to start the readout procedure of an event.

I. OVERVIEW

The Global Trigger (GT) is the final step of the CMS Level-1 Trigger [1, 2, 3]. It consists of several VME boards mounted in a VME9U crate together with the Global Muon Trigger boards (GMT) and the central Trigger Control System (TCS) [3, 4, 5].

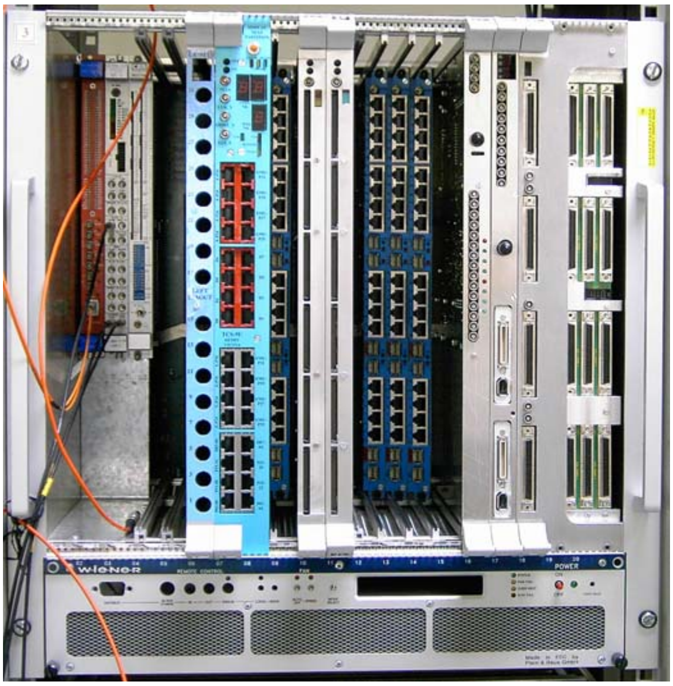

Figure 1 Global Trigger crate with prototype modules

The rightmost 4 VME slots (21-18) contain the GMT boards, slot 17 the GTFE readout board, slot 17 the TIM board, which broadcasts the clock and fast control signals to all boards. In slots 15-13 PSB boards receive trigger data from the Global Calorimeter Trigger. Slot 12 is free and in slot 11 the GTL Logic module calculates up to 128 Algorithms, which are combined to Final_OR signals on the FDL board in slot 10. In slot 9 a PSB board receives 'Technical Trigger' signals and sends them to the FDL board. The Final_OR signals go to the central Trigger Control board (TCS) in slot 8 and are transmitted via 'L1AOUT' modules in slot 6 and 7 to the TTC system of CMS(see also Fig. 1).

For every LHC bunch crossing the GT decides to reject or to accept a physics event for subsequent evaluation by the High Level Trigger. During normal physics data taking the decision is based on trigger objects, which contain information about energy or momentum, location and quality. In addition special trigger signals – so-called Technical Triggers - delivered by the subsystems are also used. The trigger objects are received from the Global Calorimeter Trigger (GCT) and the Global Muon Trigger (GMT). The input data coming from these subsystems are first synchronized to each other and to the LHC orbit and then sent via the crate backplane to the Global Trigger Logic module, where the trigger algorithm calculations are performed. For each quadruplet of "particle-like" input channels (4 μ, 4 non-isolated and 4 isolated e/ γ , 4 central and 4 forward jets, 4 τ -jets) Particle Conditions are applied. A condition for a group of up to 4 particles of the same type may require that E_T or p_T is above a threshold, that the particles are within a selected window in η or in ϕ or that the absolute difference in η or/and ϕ between two particles is within a required range. In addition, so-called 'Delta Conditions' can calculate relations in η and ϕ between two particles of different kinds. Conditions can also be applied to the trigger objects total E_T, missing E_T and H_T, the sum of the transverse energies of the highest-E_T jets. There is also a possibility to trigger on jet multiplicities.

Several Particle and Delta Conditions are then combined by a simple combinatorial logic (AND-OR-NOT) to form Algorithms. Of course, each Particle Condition bit can be used either as a trigger or as a veto condition. Each of the 128 possible algorithms applied during a given data taking period represents a complete physics trigger requirement and is monitored by a rate counter. As a last step, the Algorithms are combined by a final OR function to generate a 'L1_Accept' signal that starts the Data Acquisition System and the Higher Level Trigger software. All Algorithms can be prescaled to limit the overall Level-1 trigger rate. Eight final ORs are provided in parallel to operate sub-systems independently for tests and calibration.

In case of a readout request ('L1A' signal) the Global Trigger is read out like any other subsystem. The L1A signals arrive via the TTC network and are broadcast by the Timing board (TIM) to all other boards, including those of the Global Muon Trigger, where the arrival time of the L1A signal is translated into the corresponding Ring Buffer address. On each board a Readout Processor circuit extracts data from the Ring Buffers, adds format and synchronization words and sends the event record to a readout module, the Global Trigger Front-end board (GTFE). The incoming data are checked there, combined with GMT data to a GMT-GT event record and sent via an S-Link64 interface to the CMS Data Acquisition.

II. SYNCHRONISATION OF INPUT SIGNALS

The Global Calorimeter Trigger (GCT) sends calorimeter trigger objects over fast 1.28 Gbps serial links to three PSB input boards.

A PSB board contains 4 *Infiniband* connectors and 8 DS92LV16 Serializer/Deserializer chips from National Semiconductor to convert the serial data back to 80 MHz, 16 bits wide data streams each carrying data of 2 calorimeter channels that are multiplexed in time. Four calorimeter channels are combined to one 'quadruplet' of 32 bits that means that one Infiniband cable sends data of one 'quadruplet'. A PSB board therefore transmits data of 4 quadruplets to the logic board (GTL) board via the backplane.

As the precise arrival time of the data bits is unknown the SYNC chip on the PSB board first samples the input bits 4 times per 12.5 ns tick to find the switching point of the input data. Normally the sample furthest away from the switching time is selected and transmitted. [2] Then the SYNC chip delays the trigger data for a programmable time and sends the data as 80 MHz GTL+ signals over the backplane to the GTL board.

Phase selection and delay adjustment is done separately for each 16-bit stream to compensate for any time skew between cables and link chips. The SYNC chip also writes the input data into Ring Buffers and, in parallel, into SPY memories. The Ring Buffers keep data for some time until a L1A signal arrives. Then the Readout Processor (=ROP) moves data belonging to the L1A signal from the Ring Buffer into a Derandomizing Memory and transfers them embedded in a formatted record to the GTFE board.

A counter provides the write address for the Ring Buffer and the common BCRES signal resets the counter. The Ring Buffer has been synchronized correctly to the LHC orbit when the first data word of the first bunch crossing is written into the first memory address. The program of the synchronization procedure uses an 8k SPY/SIM memory running in parallel, which accepts the data of a full LHC orbit. It starts the SPY/SIM memory to acquire data of one complete orbit and checks if the data of the first bunch crossing were really written into the first address. If not, the delay for BCRES has to be adjusted accordingly. The BC0-data are flagged by a special sequence in bit 15 of the trigger objects.

During data acquisition a 'private' monitoring program can force the SPY memory to run continuously and to stop in case of an L1A signal to check the history of the input data.

In test mode, software can load the SPY/SIM memory with test data or simulated input data to send them instead of real data.

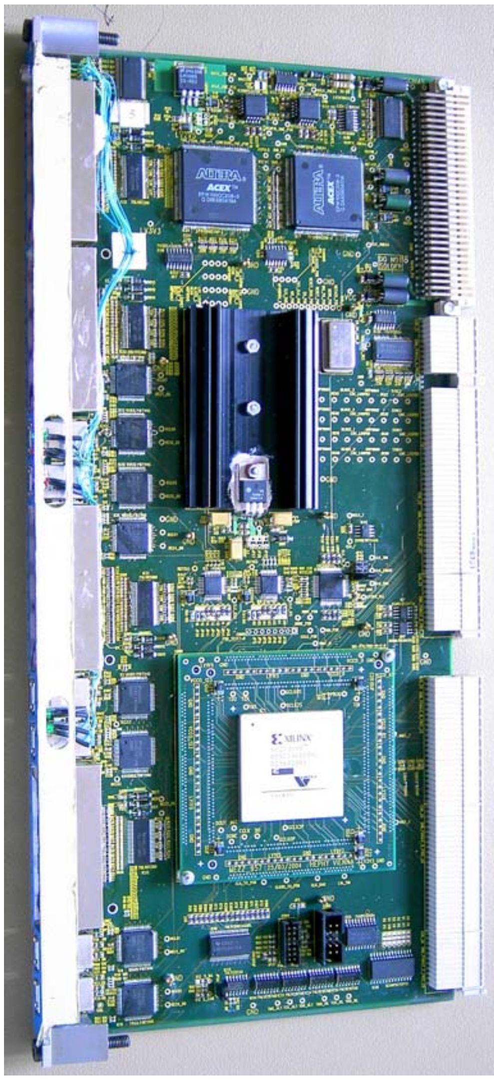

Figure 2 PSB board Version 1

Alternatively to using the two serial receiver chips, a PSB module may accept up to 64 parallel LVDS input signals via RJ45 connectors at 40 MHz frequency. Up to 16 bits are reserved for trigger signals of the TOTEM detector to include it into the CMS data acquisition. The parallel data are sampled 4 times per bunch crossing to synchronize them to the local clock signal. Then they are interlaced into an 80-Mhz data stream and transmitted and monitored instead of one of the quadruplets. The synchronization circuit exists for each group of 4 parallel input bits.

One dedicated PSB module receives Technical trigger bits as parallel LVDS data and sends them directly to the Final_OR circuit in the Final Decision Logic board (FDL).

III. TRIGGER LOGIC

On the 'Global Trigger Logic' board (GTL) the three programmable receiver chips accept the 80-MHz trigger data

and distribute them to two Condition Chips (COND). Each Condition chip receives all input data, converts them to 40-MHz objects, applies Trigger Conditions and combines the results to up to 64 Algorithms. The Algorithm bits are sent as parallel signals via short flat cables to the Final Decision Logic board (FDL) located in the adjacent slot.
As each COND chip receives all trigger bits, all kinds of logical relations between the trigger data could be implemented. Only latency requirements and chip resources restrict the number and type of triggers. But resources could be increased by replacing the Stratix chip EP1S40 by EP1S60 from Altera. [8]

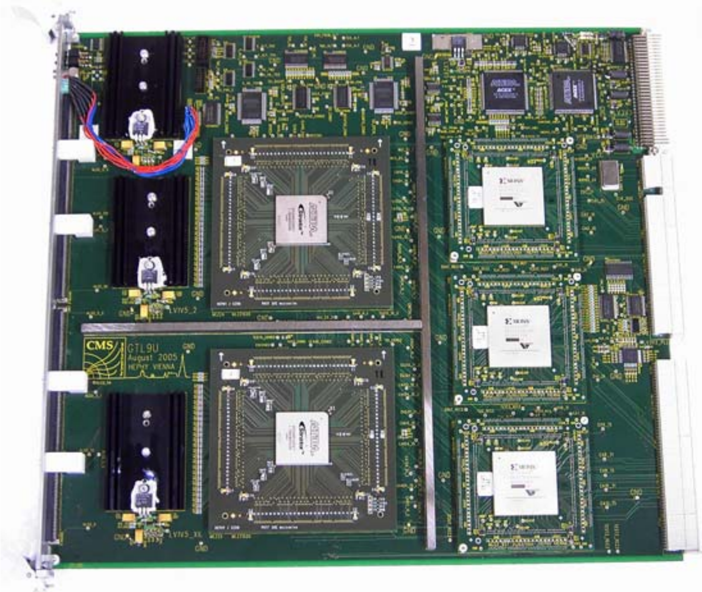

Figure 3 GTL Logic board

A. *Algorithms and Conditions*

To implement the Algorithm logic, small predefined VHDL modules are used to compose more complex trigger requirements. 'Single Particle Templates' and 'Correlation Templates' were defined for 'particle' groups (muons, electron/gamma showers, jets). A **Single Particle Template (*SPT*)** compares p_T or E_T against thresholds and checks if the particle is inside an η and/or ϕ window. For muons the required Isolation-, MIP- and Quality bits are checked in addition, and another p_T threshold can be set for isolated muons. A **Correlation Template (*CT*)** compares the differences $|\Delta\eta|$ and $|\Delta\phi|$ between two particles of the same type against thresholds and checks the charge bits for muons.
To make a '**Condition'**, the required SPT is instantiated four times to apply them to all four 'particles' and - if asked for - also the CT is instantiated as illustrated in Fig. 4. Then the results go to a combinatorial logic circuit to find 'n out of 4' particles fulfilling the requirements set by the SPTs and CT.
Four Conditions types for each 'particle' group are defined:
1s …to find one particle out of 4
2s …to find two particles out of 4
2wsc …to find two particles out of 4, correlated in η and ϕ
4s …to find four particles out of 4
If three objects are required for a particular algorithm the unused sub-condition is set to trivial values (e.g.: E_T =0 GeV, $0° < \phi < 360°$ etc.).

Conditions for the total transverse energy, the hadron transverse energy, the missing transverse energy and 12 numbers of jets above different thresholds consist only of comparators. As a last step the Condition bits are combined by a simple combinatorial logic to form a trigger **Algorithm**. All Condition bits can be used either as trigger or as veto bits.
To run the trigger Algorithms, the p_T or E_T thresholds of existing Conditions are loaded into registers using VMEbus instructions.

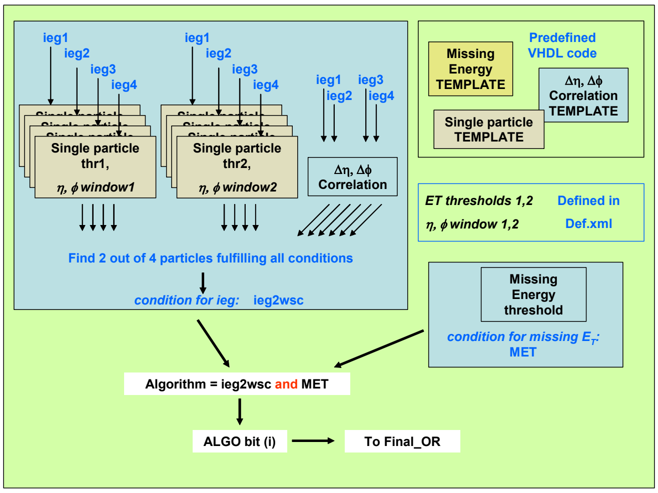

Figure 4 Algorithm composed by Conditions

When designing a new trigger setup first the Algorithms and Conditions are defined with a Java program ("gt_gui ") that runs on all machines. Its output (file "def.xml ") is used by a C++ program ("gts") that generates the variable VHDL files and a file ("vme.xml") that contains addresses and contents for all threshold registers.
A second set of thresholds will be defined for lower luminosity periods. The new VHDL files are merged with the fixed code and used by the "Quartus" software from Altera company to generate a new firmware version. The new firmware must then be loaded to run the new trigger setup. Several firmware versions will be defined to handle data taking as well as calibration and testing periods.

IV. FINAL_OR LOGIC

The 'Final Decision Logic' board (FDL) receives 128 ALGORITHM bits from the GTL board and 64 Technical Trigger bits from a dedicated PSB board. Rate counters monitor each trigger bit and pre-scalers reduce the average rate if required.
The CMS data acquisition system (DAQ) can be divided into 8 DAQ-partitions to test and calibrate parts of the readout and trigger electronics in parallel. Therefore the FDL board combines all or a subset of the Algorithm and Technical trigger bits to 8 Final_OR signals, one for each DAQ-partition, to trigger the DAQ-partitions independently from each other. Mask bits are used to include the Algorithm and Technical Trigger bits into the Final OR gates. For the Technical Trigger bits there exist also veto-mask bits to inhibit Final_OR signals.
The Final_OR signals go to the central trigger control board (TCS) that forwards them - when allowed – as 'Level 1

Accept' (=L1A) signals to the front-end electronics to read the data of the bunch crossing that has generated the trigger signal and also the data of one bunch crossing before and one after.

The FDL board can be read-out like any other front-end electronics module and contains also Ring Buffer memories which store all trigger bits. When an L1A arrives, a Readout Processor (ROP) copies data of the correct bunch crossing into a Derandomizing Buffer, embeds them into a formatted record and sends the record via a Channel Link interface and the backplane to the Global Trigger readout (GTFE) board.

As on the other boards, so-called SIM/SPY memories allow either to spy all Algorithm-, Technical Trigger- and Final_OR bits or insert simulated bits for tests. In spy mode the SIM/SPY memories run in parallel to the Ring Buffer so that latency and synchronization to the LHC orbit can be checked and adjusted.

V. DATA ACQUISITION

The FDL and the PSB input boards move all trigger input data into Ring Buffers to store them until a L1Accept signal arrives. The Ring Buffers are implemented as dual port memories inside the FPGA chips and accept the data of one full orbit. On the one side a constant write enable signal writes the trigger data of every bunch crossing into the memory. At the end of an orbit the write address returns to the first location, overwriting old data but keeping the history until a L1A signal arrives after the local latency. The local latency is the time between trigger data passing through and the time when the L1A generated by these data returns.

A delayed BCRES signal resets the counter that provides the write address so that data of BC=0 (=BC0 data) are written into location 0. To adjust the BCRES delay correctly, the software can read the SPY memory which runs in parallel to see if the BC0 data were written into the first memory word. On the other side of the Ring Buffer a counter provides the read address which lags behind the write address by the amount of the local latency minus 1 BC, so that a L1A signal reads the correct data words from the BC before until the BC after the event's BC. The L1A signal is extended to 3 BC and is applied as a write-enable signal to the Derandomizing Buffer FIFO, which extracts the data of 3 BC per event. When for debugging purposes 5 BC per event are read, the reset signal for the read counter is delayed by 1 BC less and the L1A signal is extended to 5 BC.

A readout processor (ROP) designed as a state machine reads the FIFO data of one event and wraps them with format words to create event records. The ROP is located either in the same chip or in a ROP chip if the board contains multiple FPGAs.

First the ROP sends a 'read FIFO' command to all FIFOs on the board or in the chip, respectively, to store all data words and their BC-numbers in registers. Then it sends the format words (24-bit event number, board identifier, …) to the Channel Link and fetches one 16-bit word after the other from the FIFO registers to transmit them also to the Channel Link. The ROP sends the next 'read FIFO' command to the FIFOs and repeats this procedure for the next two BC data. Finally the ROP sends an 'End_of Record' to the Channel Link and then switches to an 'IDLE' code to keep the link alive.

A. Readout board (GTFE)

The Global Trigger Readout Board (GTFE = *Global Trigger Front End)* receives event records via the backplane from the boards in the crate. The readout processor chip (ROP_DAQ) receives event records from the GMT, the FDL, the TCS and all PSB boards, checks the incoming format, combines them to a Global Trigger event record and sends it to the SLINK64 mezzanine board [7]. The ROP_EVM chip uses an identical control logic and receives event records from the TCS and FDL boards, adds GPS time - received via a TTCrq mezzanine board - and sends the compiled record via a second SLINK64 to the Event Manager of CMS. Both ROP chips use a Xilinx XC2V2000 FPGA that is mounted on a mezzanine board. [9]

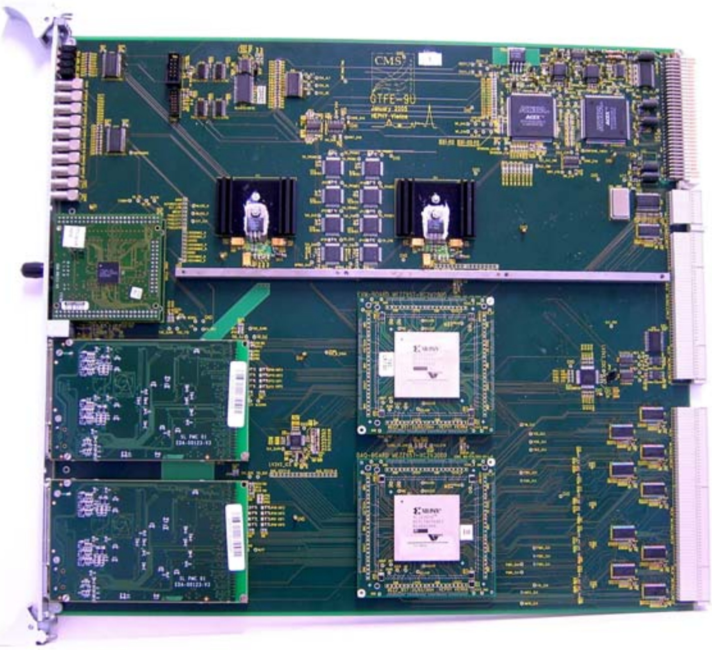

Figure 5 GTFE readout board

The GMT and the GT boards use 28-bit Channel Links to send the readout records to the GTFE board. The Channel Link bits 15-0 carry trigger data going into the data FIFOs, bits 23-16 could carry private monitoring data going into separate Monitoring memories, and bits 27-24 carry control bits going to the control logic that detects the begin and end of records. As long as IDLE data arrive the FIFOs remain inactive. The FIFOs are configured so that the output width is 4 times the input width, reordering the trigger data into 64-bit words for the SLINK64 and thus replacing a 4-to-1 multiplexer. A synchronous reset input enables the common L1Reset signal to erase all events in the FIFOs. All FIFOs can keep more than 20 events, are written with 40 MHz and are read with an 80 MHz clock. When the FIFOs become 75% full, a 'Warning' flag is sent to the Trigger Control board to reduce the trigger rate. The capacity of the FIFOs could be doubled by replacing the Xilinx XC2V2000 by a XC2V3000 chip. [9]

Both GTFE chips receive also the common signals L1A, BCRES and Event Counter Reset and create for each event a local Event and BC-number used as reference.

The Crate Readout Processor (Crate-ROP) is implemented as a state machine that reads the FIFOs of all active boards. When the first active FIFO has received an 'end-of-record' flag, the ROP applies the standard HEADER word to the

SLINK64 and reads all the data of one event on a board, then switches to the next active board FIFO and continues until the last connected board.

A comparator circuit checks if the number of events in each channel since the last 'Event Number Reset' signal agrees with the reference number. Any difference is flagged as error bit in the EVENT_STATUS byte. Nevertheless the event transmission continues until the end. Such errors will show up until the synchronization of all boards has been done correctly.

Finally, the Crate ROP appends as the last word the Event Status, the updated CRC number and the Event length. During the transmission the CRC and Event status are updated but the Event length is preloaded via VME because it is constant and depends on the number of bunch crossings per event and the boards which contribute data.

The Crate-ROP transmits data as long as there are records in the FIFOs and as long as the SLINK64 is ready. When the SLINK64 returns a 'full' flag, the Crate-ROP simply waits until the SLINK64 becomes ready again. When the off-time is too long the board FIFOs will be filled. When the 75% level is reached, the Crate-ROP sends a 'Warning' flag to the central trigger control system to reduce the trigger rate. When running with an 80-MHz clock, the Crate-ROP transfers a normal GT/GMT event (200 64-bit words) within 2.5 μs. Even when running only with a 40 MHz clock the event is transferred within 5.0 μs, thus still exceeding the required 100-kHz event rate.

When the SLINK64 for the Event Manager goes into status 'not ready', the corresponding status signal is sent directly to the TCS board to stop all DAQ-partitions.

In both ROP chips a dual port memory spies all event data which are sent to the SLINK64. The SPY memory can also be used to insert test data instead of readout data to test the reliability of the SLINK64. The other side of the SPY memory is accessed by VME-software.

VI. SUMMARY

The Global Trigger boards have been built, and all except the GTFE readout board are being integrated into CMS. Four of the PSB boards have been tested and production of the others has started. The complete trigger chain has been tested with cosmic muon data and the Global Trigger is functioning according to specification.

VII. LIST OF ACRONYMS

BC bunch crossing
BCRES common bunch crossing counter reset signal
DAQ data acquisition system
Event Manager: controls data flow of events in DAQ
FIFO First-In/First-Out memory
GCT Global Calorimeter Trigger of CMS
GMT Global Muon Trigger of CMS, mounted in GT-crate
GTL+ Gunning Transceiver Logic Plus : JEDEC JESD 8-3
L1A trigger signal to read all Front End buffers of CMS
LVDS Low Voltage Differential Signals
Quartus Altera company's Software to design and implement firmware
S-LINK64 Fast data Link extended to 64 bits (see references)
RJ45 Connector type, used by Ethernet
TOTEM Experiment measuring Total Cross Section, Elastic Scattering and Diffraction Dissociation at the LHC
VHDL VHSIC Hardware Description Language
VHSIC Very High Speed Integrated Circuit

VIII. ACKNOWLEDGEMENTS

The authors would like to thank R. Eitelberger and R. Stark for designing and machining all mechanical parts of the Global Trigger electronics. We also thank Wesley Smith from the University of Wisconsin and Joao Varela from IST, Lisbon for many helpful discussions and Fritz Szoncso (now Cern) for starting this project.

IX. REFERENCES

[1] Concept of the CMS First Level Global Trigger for the CMS Experiment at LHC, *Claudia-Elisabeth Wulz* Nucl. Instr. Meth. A473/3 (2001) 231-242

[2] Implementation and Synchronisation of the First Level Global Trigger for the CMS Experiment at LHC, *A. Taurok, H. Bergauer, M. Padrta* Nucl. Instr. Meth. A473/3 (2001) 243-259

[3] C M S, The TriDAS Project, Technical Design Report, Volume 1: The Trigger Systems, CERN/LHCC 2000 – 38, CMS TDR 6.1, December 15, 2000

[4] The Level-1 Global Muon Trigger for the CMS Experiment
H. Sakulin, A. Taurok, 9th Workshop on Electronics for LHC and Future Experiments, 29 Sept-3 Oct 2003, Amsterdam

[5] Implementation and Test of the First-Level Global Muon Trigger of the CMS Experiment H. Sakulin, A. Taurok
11th Workshop on Electronics for LHC Experiments, 29 Sept-3 Oct 2003, Amsterdam

[6] Datasheet: DS90CR287/DS90CR288A, +3.3V Rising Edge Data Strobe LVDS 28-Bit Channel Link-85 MHz, May 2002, National Semiconductor

[7] The S-LINK 64 bit extension specification: S-LINK64, Attila Racz, Robert McLaren & Erik van der Bij, EP Division, CERN.

[8] Stratix Device Handbook, Altera Corporation, San Jose, http://www.altera.com

[9] Xilinx Virtex-II Platform FPGAs: Complete Data Sheet, DS031 (v3.4) March 1, 2005 http://www.xilinx.com