



The Compact Muon Solenoid Experiment
Conference Report

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



30 October 2013 (v2, 12 November 2013)

The CMS Modular Track Finder Boards, MTF6 and MTF7

D. Acosta, G. Brown, A. Carnes, M. Carver, D. Curry, G. P. Di Giovanni, I. Furic, A. Kropivnitskaya, A. Madorsky, M. Matveev, P. Padley, D. Rank, C. Reeves, B. Scurlock, S. Wang

Abstract

To accommodate the increase in energy and luminosity of the upgraded LHC, the CMS Endcap Muon Level 1 Trigger system has to be significantly modified. To provide the best track reconstruction, the Trigger system must now import all available trigger primitives generated by Cathode Strip Chambers and by other regional subsystems, such as Resistive Plate Chambers. In addition to massive input bandwidth, this also requires a significant increase in logic and memory resources. To satisfy these requirements, a new Sector Processor unit for muon track finding is being designed. This unit follows the micro-TCA standard recently adopted by CMS. It consists of three modules. The Core Logic module houses the large FPGA that contains the processing logic and multi-gigabit serial links for data exchange. The Optical module contains optical receivers and transmitters; it communicates with the Core Logic module via a custom backplane section. The Look-up Table module contains a large amount of low-latency memory that is used to assign the final transverse momentum of the muon candidate tracks. The name of the unit - Modular Track Finder - reflects the modular approach used in the design. Presented here are the details of the hardware design of the prototype unit based on Xilinx's Virtex-6 FPGA family, MTF6, as well as results of the conducted tests. Also presented are plans for the pre-production prototype based on the Virtex-7 FPGA family, MTF7.

Presented at *TWEPP-13 Topical Workshop on Electronics for Particle Physics*

The CMS Modular Track Finder Boards, MTF6 and MTF7

D. Acosta^a, G. Brown^a, A. Carnes^a, M. Carver^a, D. Curry^a, G. P. Di Giovanni^a, I. Furic^a, A. Kropivnitskaya^a, A. Madorsky^{a*}, M. Matveev^b, P. Padley^b, D. Rank^a, C. Reeves^a, B. Scurlock^a, S. Wang^a

^a *University of Florida, Gainesville, FL, USA*

^b *Rice University, Houston, TX, USA*

E-mail: madorsky@phys.ufl.edu

ABSTRACT: To accommodate the increase in energy and luminosity of the upgraded LHC, the CMS Endcap Muon Level 1 Trigger system has to be significantly modified. To provide the best track reconstruction, the Trigger system must now import all available trigger primitives generated by Cathode Strip Chambers and by other regional subsystems, such as Resistive Plate Chambers. In addition to massive input bandwidth, this also requires a significant increase in logic and memory resources.

To satisfy these requirements, a new Sector Processor unit for muon track finding is being designed. This unit follows the micro-TCA standard recently adopted by CMS. It consists of three modules. The Core Logic module houses the large FPGA that contains the processing logic and multi-gigabit serial links for data exchange. The Optical module contains optical receivers and transmitters; it communicates with the Core Logic module via a custom backplane section. The Look-up Table module contains a large amount of low-latency memory that is used to assign the final transverse momentum of the muon candidate tracks. The name of the unit – Modular Track Finder – reflects the modular approach used in the design.

Presented here are the details of the hardware design of the prototype unit based on Xilinx's Virtex-6 FPGA family, MTF6, as well as results of the conducted tests. Also presented are plans for the pre-production prototype based on the Virtex-7 FPGA family, MTF7.

KEYWORDS: CMS, Endcap, Muon, Trigger, FPGA, Track Finder, micro-TCA

Contents

1. Upgrade motivation	1
1.1 Transverse Momentum assignment	1
1.2 Trigger primitive bandwidth and FPGA resources	2
2. MTF6 hardware prototype	2
2.1 Modular structure	2
2.2 The Core Logic module	2
2.3 The Optical module	3
2.4 The p_T LUT module	4
3. Testing	4
3.1 Optical communication	4
3.2 PCI express	5
3.3 p_T LUT tests	6
3.4 Integration test	6
3.5 Other tests	6
4. Plans for design based on Virtex-7 (MTF7)	7

1. Upgrade motivation

The current Muon Endcap (ME) Level 1 Trigger system at the CMS experiment at CERN demonstrated a performance adequate for the LHC luminosities before the upgrade [1]. The following sections cover the most important improvements needed in order for the ME trigger system to function properly after the LHC is upgraded.

1.1 Transverse Momentum assignment

A flexible and powerful way to assign transverse momentum (p_T) to the muons that have been identified by the ME trigger system is to use a Look-up Table (LUT). This approach allows for complete algorithmic flexibility as well as fixed latency, which is independent of the algorithm used. In the currently deployed system, the size of the p_T assignment LUT (p_T LUT) is 4 MB, which offers just 22 bits of address space. The parameters that have to be supplied to the LUT for the proper p_T assignment include the angular φ difference between track stubs, the η coordinates of track stubs, station presence information (mode), bending direction, etc. Fitting all that information into 22 bits of address space turned out to be a challenging task; certain non-trivial approaches had to be used, such as non-linear parameter scales, dynamic bit field reassignment based on track mode, etc.

A proper p_T assignment in the upgraded system is expected to be even more complex due to increased expected background. This requires implementing a p_T assignment LUT with significantly bigger address space to receive more track data.

1.2 Trigger primitive bandwidth and FPGA resources

The currently deployed ME trigger system filters the trigger primitives generated by each 60° azimuthal Muon Endcap sector. Only the better-reconstructed 15 primitives out of 90 in total are sent to Sector Processor boards. This reduces the efficiency for events with multiple muons in a small geometrical region (approx. 10°) and will lead to inefficiency in extremely high pile-up conditions [2]. To solve that problem, we need to import all trigger primitives into the upgraded Sector Processor boards. A much bigger FPGA is needed to process that massive amount of information. Additionally, we may need to import data from other regional subsystems; that increases the required FPGA logic size even more.

2. MTF6 hardware prototype

2.1 Modular structure

CMS has recently adopted the microTCA hardware platform [3] as a standard for new equipment development. The prototype Sector Processor has been constructed using that standard. The design consists of the following modules (Figure 1):

- Core Logic module
- Optical module
- p_T LUT module

Custom backplane

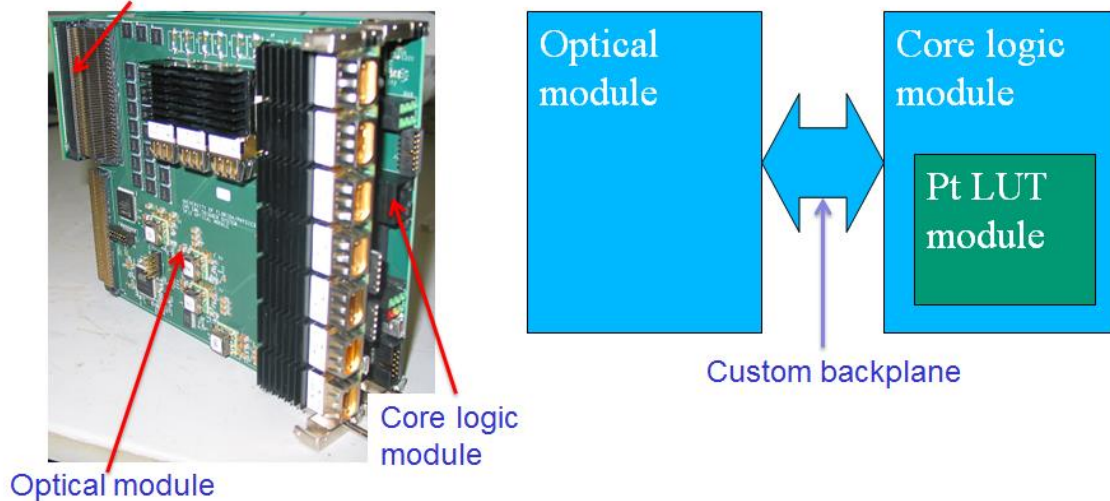


Figure 1: MTF6 prototype

2.2 The Core Logic module

The most important element that the Core Logic module contains is the large FPGA for trigger data processing. The current prototype is based on the Virtex-6 family of Xilinx FPGAs [4]. The core logic FPGA used in that design is XC6VHX565T-FFG1924. In addition to that, the Core Logic module contains, as shown in Figure 2, a smaller control FPGA, a Module Management Controller (MMC), configuration memory for both FPGAs, power supplies, and clock

management circuitry. The MMC design is provided by the University of Wisconsin, Madison [5]. The module is able to receive trigger data on 53 GTX links (up to 4.8 Gb/s each), and 8 GTH links (up to 10 Gb/s each). It can output trigger decisions and other data using 12 GTX links and 2 GTH links.

PCI express (PCIe) was selected as the main control interface solution for the upgraded Sector Processor design. This choice is dictated by the bandwidth requirements, specifically downloading p_T LUT memory contents. PCIe has high bandwidth (that can be scaled if necessary by using up to 4 lanes), low latency, and low overhead. Each module is provided with direct access to the host computer memory.

MTF6 is also compatible with IPbus [6], a control interface that is accepted as a standard in CMS.




Figure 2: MTF6 Core logic module

2.3 The Optical module

The Optical module (Figure 3) contains 7 12-channel optical receivers, 3 12-channel transmitters, MMC and control circuitry, and power supplies. All optical components are 10 Gb/s parts (AFBR-810EZ transmitters and AFBR-820EZ receivers) manufactured by Avago [7]. The optical receivers and transmitters are connected to FPGAs on the Core Logic module via a short custom backplane section. This module is compatible with the next-generation prototype based on Virtex-7 FPGAs.

2.4 The p_T LUT module

The p_T LUT module (Figure 4) is implemented as a mezzanine card that sits on top of the Core Logic module. It contains 1 GB of Reduced Latency DRAM (RLDRAM) memory, manufactured by Micron [8] (part number MT44K32M18RB-093E:A). This type of memory, while retaining all the advantages of DRAM (large capacity, low power consumption, low price), has been specifically designed to reduce latency for random address accesses. The address bit count usable for the p_T assignment is 30.




Figure 3: MTF6 Optical module (compatible with MTF7)

RLDRAM installed on the p_T LUT module is working with the 200 MHz clock. Each address and control bit is running at a rate of 200 Mbps, and each data bit is written and read at a rate of 400 Mbps. Even though RLDRAM can tolerate much higher clock frequency and data rates, this does not reduce the latency for random address access. Running at lower clock frequency, however, translates into lower power consumption and simplified FPGA logic.

3. Testing

3.1 Optical communication

Optical communication has been tested by transmitting data from multiple sources: MTF6's own transmitters (loopback), Muon Port Card¹ (MPC) prototype, and earlier Sector Processor 2010 prototype. 47 GTX and 8 GTH input channels have been run in parallel at nominal bitrates (3.2 Gb/s for GTX and 10 Gb/s for GTH), using Pseudo-Random Bit Sequence (PRBS) as test

¹ Muon Endcap trigger primitive transmission card, designed and provided by Rice University.

patterns. With the exception of two GTH channels that are known to have a defective layout on this prototype, all other channels demonstrated a Bit Error Rate (BER) lower than 10^{-13} . A 10Gb/s eye pattern is shown in Figure 5.




Figure 4: MTF6 pT LUT module




Figure 5: 10 Gb/s eye pattern

3.2 PCI express

The PCIe setup used for tests included the NAT-MCH unit with the PCIe switch from NAT Europe [9], AMC113 adapter card from Vadatech [10], OSS-PCIe-HIB35-x4-F PC adapter card (OneStopSystems) [11], PCIEO half-cable from Samtec [12], and VT892 uTCA chassis from Vadatech [13]. The connection from the control PC to the uTCA chassis was implemented with a 50 m optical fiber. MTF6 has a single PCIe lane implemented.

The performance has been evaluated to be 2.3 Gb/s for reading from MTF6 and 2.88 Gb/s for writing into MTF6.

3.3 p_T LUT tests

The tests performed included writing to and reading from consecutive addresses, and reading from randomly generated addresses. It was demonstrated that the memory array on the p_T LUT module could be read up to 5 times in each BX, so the p_T could be assigned to up to 5 muons in each BX.

3.4 Integration test

An integration test was performed with the two MPCs [14] sending data to MTF6 (Figure 6). Each MPC was clocked with its own Clock and Control Board (CCB) [14], and MTF6 was clocked from the AMC13 board [15], to simulate the layout of the real Endcap Muon Trigger system, in which each MPC resides in its own VME crate. Both CCBs and AMC13 were receiving the clock from a common source. Transmitted data included the PRBS sequence and randomly generated trigger data sent via test FIFOs. No errors have been detected.




Figure 6: Integration test setup

3.5 Other tests

The other components tested included the IPbus interface, and clock circuitry tests. Tests with the AMC13 clock distribution and DAQ collection board designed by Boston University are still in progress.

4. Plans for design based on Virtex-7 (MTF7)

To provide optimal trigger performance, the Level 1 Endcap Muon Trigger may need to use not only data from Cathode Strip Chambers (CSC), but also from other subsystems, such as Resistive Plate Chambers (RPC) and Gas Electron Multiplier detectors (GEM). The full bandwidth required for this cannot be precisely determined at this time. Taking this into account, MTF7 is being designed to provide maximum flexibility using presently available devices from the Xilinx Virtex-7 family. The Core Logic FPGA has 80 GTH receivers; all of them are dedicated for trigger data reception via optical links at up to 10 Gb/s rates (Figure 7). In addition to that, 4 GTX receivers on the Control FPGA can be used for trigger data reception as well. The total count of GTH transmitters that can send data via optical links is 28. Each serial link has two clocking options: a programmable frequency clock derived from the LHC 40.08 MHz clock, or a programmable clock derived from a fixed 250 MHz oscillator. The modular approach allows for easier partial upgrades between large LHC shutdowns.

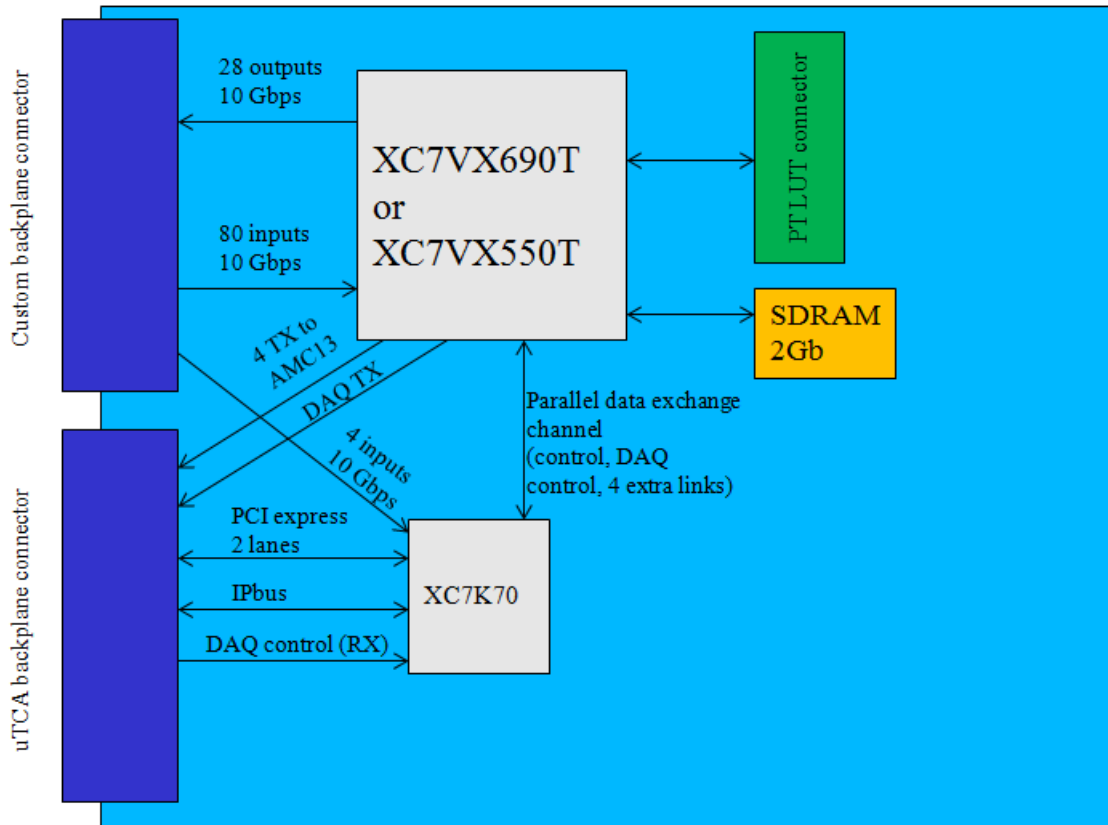


Figure 7: MTF7 Core Logic Board interconnections

Acknowledgments

We would like to thank T. Gorski (University of Wisconsin, Madison) for providing firmware, software, and support for the MMC design. We also wish to thank the Boston University Electronic Design Facility for providing AMC13 and TTT boards, and support for them. We thank IPbus team (multiple universities) for providing the IPbus firmware and software. We thank the RPC team (University of Warsaw) for their helpful suggestions.

References

- [1] J. Brooke, Performance of the CMS Level-1 Trigger, <http://arxiv.org/pdf/1302.2469v1.pdf>
- [2] CMS collaboration, CMS Technical Design Report For the Level-1 Trigger Upgrade <https://cds.cern.ch/record/1556311/files/CMS-TDR-012.pdf>
- [3] Micro-TCA specification: <http://www.picmg.org/v2internal/microtca.htm>
- [4] Virtex-6 FPGA family <http://www.xilinx.com/products/silicon-devices/fpga/virtex-6/index.htm>
- [5] University of Wisconsin MMC design http://www.hep.wisc.edu/ecad/SLHC/MMC/AMC13_MMC_v1p2_download/
- [6] IPbus <https://svnweb.cern.ch/trac/cactus>
- [7] Avago parallel optics http://www.avagotech.com/pages/en/fiber_optics/parallel_optics/
- [8] RLDRAM 3 memory <http://www.micron.com/products/dram/rldram-memory#fullPart&468=1>
- [9] NAT-MCH <http://www.nateurope.com/products/NAT-MCH.html>
- [10] Vadatech's AMC113 adapter <http://www.vadatech.com/products/?show=product&product=121>
- [11] OneStopSystems switch-based adapter http://www.onestopsystems.com/pcie_over_cable_z7.php
- [12] Samtec's PCIe optical cable <http://www.samtec.com/technical-specifications/Default.aspx?SeriesMaster=PCIEO>
- [13] Vadatech's VT892 uTCA chassis <http://www.vadatech.com/products/?show=product&product=186>
- [14] CCB and MPC manuals (Rice University) <http://bonner-ntserver.rice.edu/cms/projects.html>
- [15] AMC13 DAQ and control board (Boston University) <http://bucms.bu.edu/twiki/bin/view/BUCMSPublic/HcalDTC>