

The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2012 JINST 7 C12024

(<http://iopscience.iop.org/1748-0221/7/12/C12024>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 155.198.211.93

The article was downloaded on 13/03/2013 at 13:21

Please note that [terms and conditions apply](#).

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2012,
17–21 SEPTEMBER 2012,
OXFORD, U.K.

The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS

K. Compton,^d S. Dasu,^d A. Farmahini-Farahani,^d S. Fayer,^a R. Fobes,^d R. Frazier,^c T. Gorski,^d G. Hall,^a G. Iles,^a J. Jones,^b P. Klabbers,^d D. Newbold,^c A. Perugupalli,^d S. Rader,^d A. Rose,^{a,1} D. Seemuth,^d W. Smith^d and J. Tikalsky^d

^aBlackett Laboratory, Imperial College, London,
Prince Consort Road, London, SW7 2BW, U.K.

^bIceberg Technology,
Helland Mill, Mabe Burnthouse, Penryn, Cornwall, TR10 9JQ, U.K.

^cH.H. Wills Physics Laboratory, University of Bristol,
Tyndall Avenue, Bristol, BS8 1TL, U.K.

^dDepartment of Physics, University of Wisconsin-Madison,
1150 University Ave., Madison, WI 53706, U.S.A.

E-mail: awr01@imperial.ac.uk

ABSTRACT: Test results are presented for two AMC cards, the “CTP6” and “MP7”. The two cards take different approaches to connectivity: the CTP-6 has fully-populated backplane connectivity and a 396 Gbps asymmetric, optical interface, whilst the MP7 instead favours a 1.4 Tbps, symmetric, all-optical interface.

The challenges of designing the MP7 card necessitated the development of several test cards; the results of which are presented.

KEYWORDS: Digital signal processing (DSP); Trigger concepts and systems (hardware and software); Modular electronics; Digital electronic circuits

¹Corresponding author.

Contents

1	Introduction	1
2	The CTP-6	1
2.1	Serial links	2
2.2	Custom backplane	2
2.3	Testing the CTP-6	2
3	The MP7	3
3.1	Optical interface	3
3.2	The FPGA	4
3.3	PCB	5
3.4	Testing of the MP7	5
4	Test boards	5

1 Introduction

To maintain or improve the physics performance of the CMS calorimeter trigger under full-energy, high-luminosity conditions, the existing mixed ASIC/FPGA-based system will need to be replaced. Two trigger architectures have been proposed, one by each group that built the current calorimeter trigger.

The University of Wisconsin, Madison, favours a conventional trigger architecture, requiring backplane interconnections to handle the boundaries between regions. The CTP-6 is a prototype of the processor card needed for this architecture.

Imperial College, London, is exploring the possibility of a spatially-pipelined, time-multiplexed trigger architecture, which requires no backplane interconnection and streams all data from an event into a single FPGA for processing. Because such a trigger architecture throws no data away prior to the final processing stage, it requires a larger optical bandwidth than a conventional trigger. The MP7 is an all-optical, data-stream processor of the final design needed for such a trigger.

A collaboration has been formed between Imperial College and University of Wisconsin to build a system which could be used as either trigger architecture.

2 The CTP-6

The CTP-6, figure 1, is a prototype Calorimeter Trigger Processor Card by the University of Wisconsin. It is an AMC card with two Virtex-6 (XC6VHX250T) FPGAs [1] with GTX links, dual SDRAM for dedicated DAQ and TCP/IP buffering, a Module Management Controller (MMC) of the Wisconsin design, and custom-designed power modules providing 1.0 V, 1.5 V and 2.5 V.




Figure 1. Top (right) and bottom (left) surfaces of the CTP6. The CTP6 is based around two, high-end Xilinx Virtex-6 FPGA, being fed by four PPOD optical receiver modules and feeding one PPOD optical transmitter module, providing a total optical bandwidth of approximately 310 Gbps in and 78 Gbps out. The board also has a further twelve serial links going to and coming from the backplane.

2.1 Serial links

On the front panel, the CTP-6 uses 4 Avago AFBR-820B Receiver Modules, and one AFBR-810B Transmitter module to provide 48 optical inputs and 12 outputs running at up to 6.4 Gbps [2]. There are 24 inter-FPGA links, also running at 6.4 Gbps. The trigger-architecture favoured by Wisconsin also requires back-plane communication. The CTP-6 has 12 bidirectional links to the backplane which can be run at up to 4.8 Gbps.

2.2 Custom backplane

A backplane has been manufactured by VadaTech to a design specified by the University of Wisconsin and is now commercially available with part number VT894 [3]. This backplane has the same standard connections as the VadaTech VT892 backplane [4], but adds 88 differential pairs for the data sharing requirements of the Wisconsin trigger architecture. This is done using the AMC connector ports which are unconnected on the VT892. The additional connections use the same link technology as the standard backplane connections. Figure 2 shows the connection topology. Slots 2–5 and 8–11 have bi-directional links in a 2×4 array for handling boundaries within a processing region. Slots 6 and 7 have links to the upper or lower ranks of the processing array and, in the Wisconsin trigger architecture are used for collating data for inter-crate communication. For completeness, the remaining links are routed to slots 1 and 12, but these links are unnecessary for the Wisconsin trigger architecture.

2.3 Testing the CTP-6

Testing of the CTP-6 is underway at the University of Wisconsin. The parallel FPGA flash memories allow both FPGAs to load in approximately 6 seconds. An SI5338 programmable clock synthesizer [5] allows multiple options for the reference clocks of the GTX links. Gigabit Ethernet

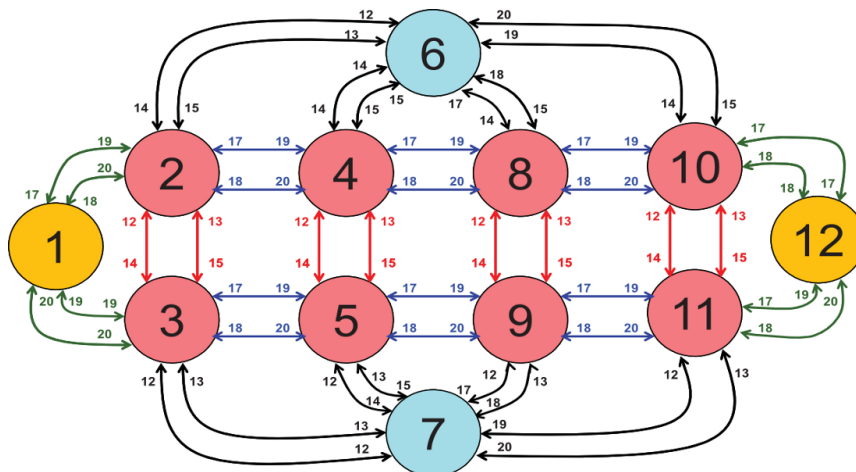


Figure 2. The link topology of the Wisconsin-specified custom backplane. The numbered circles represent backplane slots and the arrows indicate the custom connections between those slots, with the numbers next to the arrows indicating through which AMC-connector fabric pair the connection is made.

is successfully running on a Microblaze embedded processor on the back-end FPGA. The MMC is working with phased power-supply enables, analogue voltage-sensors, and sensors for “power good” and “FPGA load done”. The inter-FPGA serial links are operating at 6.4 Gbps. The Wisconsin trigger design also requires data transfer through the standard μ TCA fat-pipe. This has been tested, routing through a cross-point tongue, based on a Vitesse VSC3172 72×72 -way switch [6], mounted on a VadaTech MCH, and is working at 4.8 Gbps.

3 The MP7

In contrast to the typical design methodology employed in particle physics, the MP7, figure 3, was not designed to fulfil a specific role, but rather, to be a generic stream-processing engine. This design philosophy manifests itself most clearly in the distinctive data interface, which is all-optical and has symmetric transmit and receive bandwidths.

3.1 Optical interface

Using a single interface for the sending and receiving of all data offers several distinct advantages: primarily, as stated above, the board ceases to have a specific role and becomes a truly generic stream-processing engine. Application of the board to a specific task is no longer restricted by the bandwidth and compatibility of each type of interface, but only by the total bandwidth, whilst all specialization required for a task is contained within the programming of the board and the interconnections between boards.

The MP7’s optical interface is provided by up to six Avago MiniPOD transmitters and six Avago MiniPOD receivers [7]. Each MiniPOD device provides 12 optical links running at up to 10.3 Gbps, giving the MP7 a total optical bandwidth of up to 740 Gbps in each direction. The optical outputs from the MiniPOD devices are transmitted on Corning unruggedized, non-peelable




Figure 3. Top (right) and bottom (left) surfaces of the MP7. The central component of the MP7 is a single, mid-range Xilinx Virtex-7 FPGA, being fed by six MiniPOD optical receiver modules and feeding six MiniPOD optical transmitter modules, providing a total optical bandwidth of approximately 740 Gbps in each direction. The optics are not mounted in these images.




Figure 4. The MP7 card with MiniPOD optics, optical fibres, front panel and a heatsink mounted. This configuration of the optics can simultaneously transmit and receive nearly 0.5 Tbps. The optical ribbons used here are samples: the spare connectors and excess length will not be present in the final fibre design.

optical ribbons [8] to 48-way MTP connectors (of which only 36 channels are used) mounted on the front-panel. The optics may be seen mounted on the board in figure 4.

3.2 The FPGA

The MP7's processing capability is provided by a Xilinx Virtex-7 FPGA [9] in a FFG1927 package. Xilinx provide four pin-compatible parts in this package with varying amounts of logic and numbers of serial links; the VX415T, VX485T, VX550T and the VX690T. The VX550T and VX690T parts both provide a full complement of 72 optical links in each direction, with the latter part sim-

ply offering more logic resources than the former. The MP7 has been designed such that, when using the VX485T part, the 48 optical links available in each direction are all routed to the rear 8 MiniPODs, so that system design is simplified, and a cost saving made, by not having to deal with partially-used MiniPODs.

3.3 PCB

The MP7 has 144 differential-pairs running at 10 Gbps and many more running at the order of 1–2 Gbps. To ensure signal integrity in the high-speed traces, a ground-plane has been placed between each layer containing high-speed traces, and the presence of these ground planes, necessitated the use a 16-layer stack-up. The number of layers, coupled with the 1.6 mm board thickness specified by the μ TCA standard, and the sheer number of high-speed traces required the use of Nelco N4000-13 EP SI [10] as a board substrate, due to its excellent high-frequency performance and low dielectric constant.

3.4 Testing of the MP7

The main source of risk in the design of the MP7 was the sheer number and high speed of the optical links and, as such, the main test of the MP7 is the integrity of the optical links. The links have been validated by loop-back test, both as two sets of 24 links running PRBS-31, and all 48 links simultaneously running 8B10B. Running the links at 10 Gbps the MP7 transfers 0.48 Tb per second in each direction. To date, the MP7s have transferred in excess of an Exabit of data without an observed error, giving a limit on the per-board bit-error rate of approximately 3×10^{-17} .

4 Test boards

The unprecedented (in our experience) high speed and sheer density of the signal traces for the optical links on the MP7 board led us to believe it prudent to develop a test board for the purpose of examining the effects of various design factors on the integrity of 10 Gbps signals. Of particular interest were the effects of stripline and micro-stripline structures, layer transitions using both vias and blind vias, and stubs associated with both full and blind vias.

The most unexpected result from the initial test board was a 50 ps skew between members of a 6 inch differential pair (figure 5a), far greater than that expected from the trace length-matching tolerance, rendering a 10 Gbps signal unusable. The cause of the skew was determined to be the different impedance seen by traces aligned along the fibres in the substrate from that seen by trace aligned between the fibres in the substrate (figure 6). The suggested solution for this problem was to have the PCB manufactured at an angle of 22.5° relative to the weave of the substrate, so that no trace, whether at 0° , 45° , 90° or 135° , could run parallel to the fibres and, by so doing, the impedance seen by all traces would average out to approximately the same value. To test this solution, the second test board was manufactured using this technique, and the trace-trace skew was seen to vanish (figure 5b).

During layout of the MP7, it became apparent that to meet both the tight spatial constraints on the board and the strict conditions on the Virtex-7's multi-gigabit transceiver power supplies, some of the switch-mode power supply modules would need to be located directly beneath the high-speed traces. Concerns about the risk to the high-speed traces of EMI from the power modules prompted

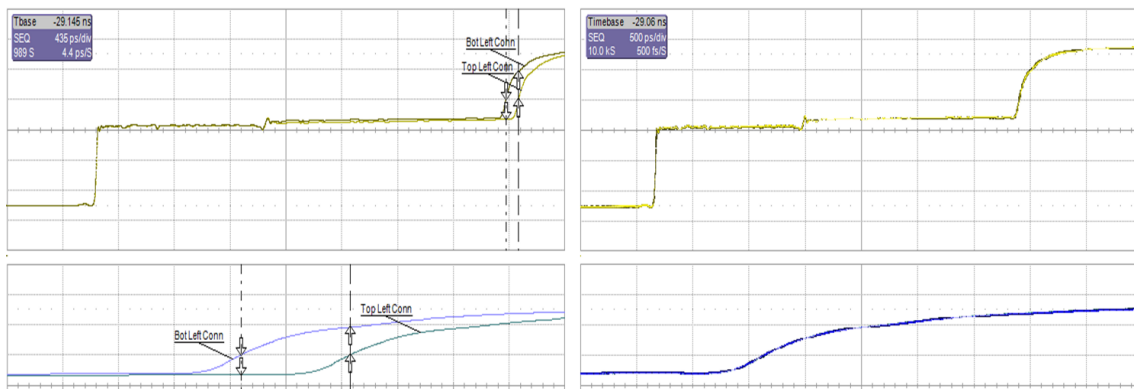


Figure 5. The trace-trace skew measured on the MP7 test board. The left-hand plot shows the 50 ps infra-pair skew measured on the first test board, where the traces were aligned with the weave of the substrate. The right-hand plot shows that, by rotating the traces to an angle of 22.5° relative to the weave, the second test board did not suffer from the same infra-pair skew.




Figure 6. Figure demonstrating a trace aligned along the fibres in the substrate and one aligned between the fibres in the substrate. The different dielectric constant of the fibre and the epoxy result in different impedances for each member of the differential pair and an infra-pair skew. On a 6 inch trace, an inter-trace skew of 50 ps was observed, rendering a 10 Gbps unusable. By rotating the trace to an angle of 22.5° relative to the weave, the dielectric constant is averaged over epoxy and fibre and the source of skew removed.

the inclusion on the second test board of a high-speed trace running under the chosen regulator to test for such effects. Although the test structure was a considerably harsher design than that on the MP7, being longer, having more bends, having more passes under the power module and having fewer power-planes between the regulator and the high-speed traces, no more significant degradation of the signal was seen than was expected for a trace of that length.

As well as directly testing various trace structures, the test boards were also used to test the accuracy of the Mentor Graphics HyperLynx PCB analysis software without full EM simulation. It was found that, overall, the software gave very accurate predictions, with one exception: because our licensed version of the software uses an approximation of vias, rather than a full EM simulation, these approximations would, in some cases, give questionable results. By combining the results from the test boards with the predictions made by the simulation software, a good understanding of the 10 Gbps traces was gained prior to the manufacture of the MP7, and the performance of the final board is as good as, or better than, expected.

Acknowledgments

Imperial College would like to acknowledge the Science and Technology Facilities Council (STFC) for funding the U.K. contribution to this research. Also, we would like to thank Sarah Greenwood for her exceptional layout work on the MP7 test cards and on the MP7 itself.

The University of Wisconsin would like to acknowledge the U.S. Department of Energy (DOE) and the National Science Foundation (NSF) for funding the US contribution to this research.

References

- [1] Xilinx Incorporated, *Virtex-6 family overview product specification*, (Jan 2012)
http://www.xilinx.com/support/documentation/data_sheets/ds150.pdf.
- [2] Avago Technologies, *AFBR-810BxxxZ and AFBR-820BxxxZ: twelve-channel transmitter and receiver pluggable, parallel-fiber-optic modules data sheet*, (Feb 2010)
<http://www.avagotech.com/docs/AV02-1955EN>.
- [3] VadaTech Incorporated, *μ TCA chassis with 12 double-width AMC slots: VT894*, (Jan 2012)
<http://www.vadatech.com/files/pdfs/VT894%20SpecRev.pdf>.
- [4] VadaTech Incorporated, *μ TCA chassis with 12 double-width AMC slots: VT892*, (Apr 2011)
<http://www.vadatech.com/files/pdfs/VT892%20SpecRev.pdf>.
- [5] Silicon Laboratories Incorporated, *I²C-programmable any-frequency, any-output quad clock generator*, (2012) <http://www.silabs.com/Support%20Documents/TechnicalDocs/Si5338.pdf>.
- [6] Vitesse Semiconductor Corporation, *VSC3172 6.5 Gbps 72 \times 72 asynchronous crosspoint switch datasheet*, (Jan 2011).
- [7] Avago Technologies, *MicroPODTM and MiniPODTM 120G transmitters/receivers*, (Feb 2011)
http://www.avagotech.com/pages/minipod_micropod.
- [8] Corning Cable Systems LLC, *Non-peelable fiber optic ribbon, 12 fiber. Specification sheet EVO-996-EN*, (Nov 2009).
- [9] Xilinx Incorporated, *7 series FPGAs overview advance product specification*, (May 2012)
http://www.xilinx.com/support/documentation/data_sheets/ds180_7Series_Overview.pdf.
- [10] Park Electrochemical Corporation, *Nelco[®] N4000-13 EPTM, Nelco[®] N4000-13 EPTM SI[®] high-speed multifunctional epoxy laminate and prepreg*, (Sep 2012)
<http://www.parkelectro.com/parkelectro/images/n4000-13ep.pdf>.