Published by IOP Publishing for Sissa Medialab

Received: *October 21, 2022*
Accepted: *December 5, 2022*
Published: *January 17, 2023*

Topical Workshop on Electronics for Particle Physics
Bergen, Norway
19–23 September 2022

Lessons learnt from the first vertical slice of the CMS Outer Tracker

G. Fedi,[a,*] S. Fiorendi,[b] M. Holmberg,[c] A. Howard,[a] G. Iles,[a] D. Monk,[a] M. Pesaresi[a] and K. Whalen[c] on behalf of the CMS collaboration

[a] *Imperial College London, London, U.K.*
[b] *University of Tennessee, Knoxville, TN, U.S.A.*
[c] *Rutherford Appleton Laboratory, Chilton, Oxfordshire, U.K.*

E-mail: giacomo.fedi@cern.ch

Abstract: A vertical slice of the CMS Outer Tracker has been tested at the tracker integration facility and at the M2 muon beam facility at CERN. It includes the final prototype of the 2S module with an optical link to the back-end ATCA system. The performance of the system will be described, including cooling limits of racks, robustness of 25 Gbit/s trigger optical links, and readiness of firmware.

Keywords: Front-end electronics for detector readout; Detector cooling and thermo-stabilization; Optics; Trigger concepts and systems (hardware and software)

*Corresponding author.

https://doi.org/10.1088/1748-0221/18/01/C01041

Contents

2023 JINST 18 C01041

1 Introduction

The Outer Tracker [1] for the CMS Phase-2 upgrade is a new silicon charged particle detector that is fundamentally different from its predecessor. For every collision, it is able to perform in situ data filtering to select hits generated by high transverse momentum charged particles coming from the interaction region of the detector. These hits, known as "stubs", are used to identify up to 100 particle trajectories per collision with a latency of ~4 µs. The resulting tracks are used to substantially improve the trigger performance, which at present relies only on data from the muon and calorimeter sub-detectors. The new Outer Tracker also performs its existing primary task of supplying all detector hits upon receipt of a Level-1 Accept (L1A) trigger [2].

The new back-end (BE) electronics, situated in an underground cavern adjacent to the detector, consist of two parts: the data, trigger and control or DTC boards, and the track finder processor or TFP boards. The DTCs communicate to the front-end (FE) electronics, using ~30 k optical fibres with a data throughput of 5–10 Gbit/s. It receives both the L1A and trigger data on the same optical link. In both cases the data needs to be extracted, merged with other FE modules, and then sent either to the data acquisition system (L1A data) or the TFP for track identification (trigger data) via 25 Gbit/s optical links.

The DTC is based on a Serenity [3] ATCA board configured with the appropriate optical modules installed and firmware/software that comprises two components: an algorithm block to perform track processing and an infrastructure tasked with handling the different optical link protocols.

The prototype Serenity board used in these tests is a dual FPGA board with up to 12 pairs of Samtec FireFly optical transceivers. It hosts a ComExpress CPU module and is controlled by an open source IPMC mezzanine (OpenIPMC [4]). To augment flexibility during the prototyping phase, the FPGAs are mounted on daughter cards and connect to the main Serenity PCB through a Samtec ZA1 interposer.

This paper reports on measurements made with prototype Serenity boards that were necessary to validate design choices before embarking on production. Cooling studies at both board and rack level are reported in sections 2 and 3 respectively. The qualification of Samtec FireFly 25 Gbit/s, 12-channel, optical modules is reported in section 4. In the last section 5 the vertical integration of the system at CERN is reported; both at the tracker integration facility (TIF) and at the M2 muon beam facility.

2 Board thermal measurements

Each 19′ 56U electronic rack is expected to host two ATCA shelves, each with up to 12 application cards. Rack power is limited to 10 kW, which requires maximising the power delivered to the electronics (as opposed to fans) while still keeping the electronics cold enough for a reliable ten-year lifespan. To put this in perspective, the ATCA fans alone can consume more than 30% of allotted shelf power at full speed. The optics transceivers (FireFly) are particularly challenging as their temperature is recommended to be kept below 50 °C for long term operations [5], while the FPGA die temperature should not exceed 100 °C.

The thermal performance of the Serenity prototype was studied for a range of different heat-sinks, see figure 1, mounted on the two Kintex Ultrascale 15 Plus (KU15P) FPGAs. A specialised piece of firmware, using high-speed oscillators, allowed the FPGA power consumption to be varied between 40 W and 80 W under software control, though later studies with larger FPGAs consumed up to 150 W. The thermal cooling performance was measured using the internal FPGA temperature sensor. FPGA heat sinks were tested with different geometries, materials (aluminium and copper), integrated heat-pipes and vapour-chambers. A 10×10 cm^2 planar heat sink is sufficient to cool the FPGA; any larger heat sink brings little advantage and reduces the space available for board components. The ATCA standard limits components on the main side of the board to a height of 21.33 mm (no protective cover), which makes cooling parts that consume more than 100 W difficult. This is one of the reasons why the FPGA is directly mounted on the PCB in the new version of the Serenity board, rather than on the daughter card. This allows the cooling fin length to be increased from 9.7 mm to 12.3 mm. Heat pipes were installed in trenches on the bottom part of a pair of aluminium heat sinks with two different geometries: in one case four heat pipes were installed radially with one end in contact with the FPGA package, in the other case two heat pipes were installed horizontally with the centre of the heat pipes in contact with the FPGA package. In both cases, the performance did not show any particular performance improvement. Tests with a copper planar vapour chamber combined with a copper heat sink showed promising results. The power dissipation of a 9×9 cm^2 copper finned heat-sink operating at an extrapolated FPGA die temperature of 100 °C increased from 141 W to 203 W.

The final ATCA board for Phase-2 will be based on the large Virtex Ultrascale 13 Plus (VU13P) FPGA that can consume up to 200 W and it is a factor of 3.3 larger than the KU15P. It also has a larger package (52.5×52.5 mm^2) compared to the KU15P (42.5×42.5 mm^2) and is composed of four separate dies rather than a single die. This helps distributing the heat and improves the thermal performance by 10–20%. From extrapolations of measurements made with the dual KU15P configuration and more recent studies with a VU13P configuration we expect a finned

2023 JINST 18 C01041

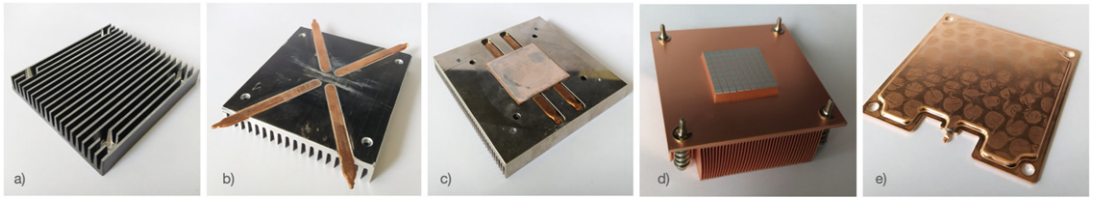

Figure 1. Some of the tested heat sinks: (a) aluminium finned; (b) aluminium finned with radial heat pipes; (c) aluminium finned with horizontal heat pipes; (d) copper finned; (e) vapour-chamber.

copper heat sink, with a planar size of 10×10 cm^2, to be sufficient to maintain the FPGA die temperature below 100 °C. In the case of a board with two VU13Ps the FPGA power consumption per FPGA is limited to 120 W per FPGA because of card electrical power limitations (400 W). A vapour-chamber, when coupled to a VU13P, improves the thermal performance by at least 10%, but with a non-recurring engineering cost of ~5 k USD.

3 Rack thermal measurements

In contrast to many commercial ATCA shelves, those used in the LHC experiments have a purely vertical air flow. Heat is extracted from the air flow by water-cooled heat exchangers. At the top of the rack a turbine diverts the flow into the sides of the rack until the air reaches an air deflector that redirects the air flow once more into the main rack space. The shelves have two fan packs in the top and bottom sections of the shelf. The fan speed can be automatically or manually set in a range between 0 and 15. The power consumption of the fans depends exponentially on the fan speed setting, with a maximum of 2 kW, which is significant when compared with all the ATCA cards power consumption (maximum 6.3 kW, though typically ~3–4 kW). In order to stay within the rack power limit of 10 kW and allow for two fully populated shelves (4 kW each) the fan speed is limited to speed 10, corresponding to 0.5 kW per shelf. A side benefit is that the noise level falls to a more manageable level (~76 dBA instead of >90 dBA). There is only a small reduction in cooling capability due to this restriction.

The rack configuration was tested with ATCA heater boards that allowed the total rack power consumption (including fans) to be varied. A number of sensors (e.g. temperature, power, and water flow) were installed to monitor the system. A water flow rate of 18 l/min at 16 °C was shared amongst the three heat exchangers. Temperature sensors in the fan pack of each shelf measured the inlet and outlet air temperature of the ATCA shelves. For the top shelf, the inlet air temperature, which is used to cool the ATCA boards, increased linearly from 20 °C to a maximum of 30 °C as the total rack power increased from 1.2 kW to 10 kW. The results shown in section 2 were obtained assuming that the heat exchangers would supply air at approximately 20 °C, providing a headroom of 30 °C for the FireFly and headroom of 80 °C for the FPGA. The higher inlet air temperature of 30 °C has a disproportionate impact on the thermal headroom of the FireFly devices (i.e. reduced to just 20 °C), which is too small to dissipate the power, particularly from the 12-channel 25 Gbit/s optical devices. At present, the only viable solution is to double the water flow rate to ~36 l/min,

2023 JINST 18 C01041

which should be sufficient to reach an inlet air temperature of 22 °C. Studies are underway to measure the water flow rate in the CMS BE racks and to improve the water flow rate at the TIF.

4 25 Gbit/s optical tests

The 25 Gbit/s BE trigger and DAQ links are a key component of CMS. They do not operate with a FEC to minimise latency unlike the commercial 100G-SR4 Ethernet standard and are therefore relatively niche products. Up until recently only limited studies have been presented on the robustness of the 25 Gbit/s BE trigger links. New studies were undertaken to provide insight into the performance of all the different sections of the link; i.e. the FPGA to the opto-electronic device, the optical transmitter, the optical receiver, and the opto-electronic to FPGA receiver performance.

Eight pairs of pre-production (beta release) 25 Gbit/s FireFly 12-channel transmitter-receivers have been tested using the Serenity board. The tests are performed with custom tcl scripts that interact with a Xilinx IBERT design. Unless otherwise stated, all links are configured in optical loop-back at 25.6 Gbit/s with a PRBS31 test sequence.

The electrical path from the FPGA to the FireFly was tested by a wide-ranging parameter scan of the FPGA transceiver parameters (signal strength, pre- and post-emphasis) and the FireFly transmitter equalisation. All the FireFly devices showed similar results with some margin on working parameters, see figure 2 (left).

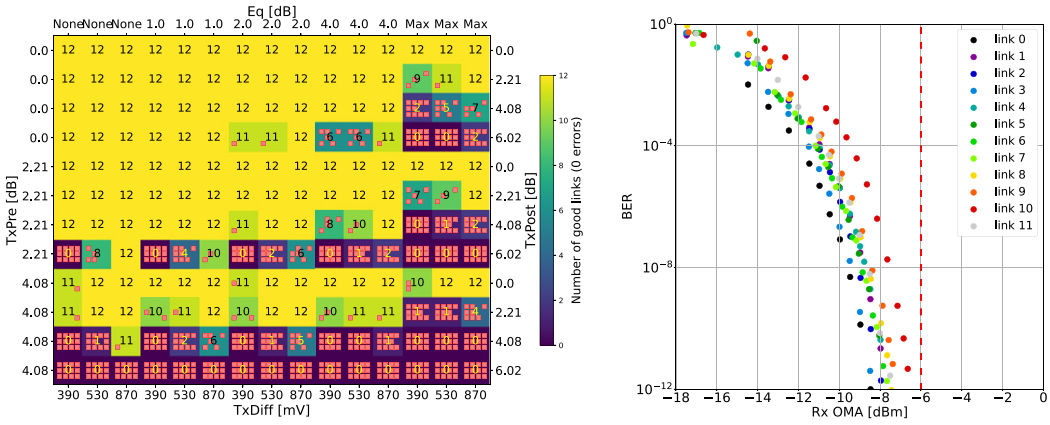

Figure 2. Left: BER parameter scan for the same FireFly. The four axes show all the parameter combinations measured. In each bin the number of links with zero errors in 10^{12} bits are given. The red squares within a bin represent the links with errors; right: BER versus receiver OMA for all 12 links of the 25 Gbit/s FireFly.

The optical transmitter performance was measured for each of the 12 channels by connecting the fibre individually to a Keysight N1092C optical sampling scope. The Extinction Ratio (ER) and the Optical Modulation Amplitude (OMA) were measured and found to be within specifications (ER > 3 dB and OMA between −3 and 3 dBm). One transmitter channel failed during testing, though this may be due to an incorrect initial power sequence. In figure 3 (left) the eye diagram of a typical FireFly is shown. The optical receiver performance was measured by inserting a

2023 JINST 18 C01041

Noyes VOA5 optical attenuator into the optical path of the channel under test while the remaining channels were left in direct loop-back. To create the receiver performance plot, see figure 2 (right), the Rx OMA was estimated from the measured Tx OMA minus the attenuation setting, which included a constant term. The receiver sensitivity at a BER $< 10^{-12}$ is within the specification of −6 dBm.

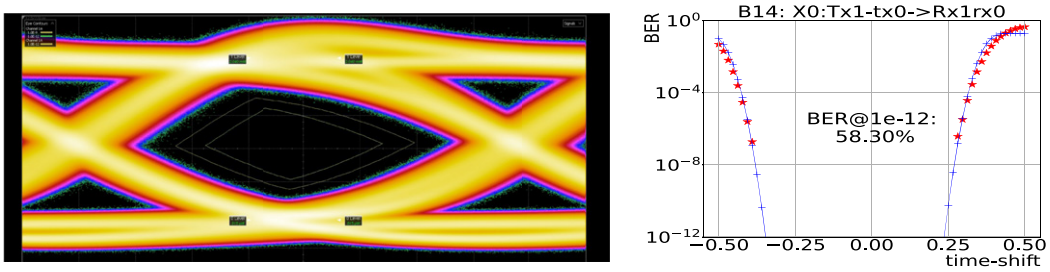

Figure 3. Left: optical eye of FireFly transmitter at 25.6 Gbit/s; right: a 1-D bathtub for a single link showing the dual-Dirac fit (blue) to the lower error values (red) which gives an extrapolated opening at 10^{-12} BER.

The performance of the electrical path from FireFly to FPGA was evaluated by measuring a bath-tub with a Xilinx IBERT. To speed up the process the IBERT dwell time was limited to 10^{-6}, and the BER performance at 10^{-12} was extrapolated with dual-Dirac fit [6], see figure 3 (right). Note that only data points on the edge of the bath-tub (i.e. low BER) were used for the fit.

5 Vertical integration

The integration of Outer Tracker FE and BE electronics is being carried out in the TIF. So far, 2S modules have been connected to a Serenity board with DTC firmware loaded onto the FPGA. Stubs and clusters, either internally generated on the modules or from cosmic muons, are driven by the VTRx+ [7] to the FireFly optical receiver on the Serenity through a 100 m long optical fibre. Both the stub path and the L1 data path have been successfully verified and the test stand is being used for continuous development of the DTC firmware and software. The DTC to TFP path was tested between Serenity and Apollo [8] boards using multiple 4Rx/4Tx FireFly interfaces (40 links) running at 25 Gbit/s. Independently, a narrow slice of the TFP algorithm reconstructing tracks from received stubs has been tested in an Apollo board, demonstrating end-to-end track finding in a small angular region of the Tracker.

With the current available prototype electronics we cannot build a full vertical slice as only a limited number of modules are available; not enough to generate the amount of stubs needed to reconstruct tracks in the TFP. A vertical slice, from the FE to the DTC, has been installed in the M2 muon beam facility at CERN as part of a beam test. The muon beam passes through four 2S modules and the generated stubs are sent to a Serenity board, running a version of the DTC firmware, over optical fibres running at 5 Gbit/s. The stubs are processed and routed to 10 Gbit/s network IP cores, which send the data using UDP packets to a computer server. The vertical slice

2023 JINST 18 C01041

has been shown to be stable with zero errors over runs lasting six hours or more, and over 50 TB of data has been stored to disk so far. Analysis of the data shows that the stubs from all four modules are synchronised to within 0.2 ns after tuning of the FE delay-locked loops. The system is able to successfully measure the properties of the muon beam, including track parameters, in real-time at the full 40 MHz stub rate and the slice will be expanded to handle more 2S modules in the near future.

6 Summary

The experience acquired during the integration tests that have been carried out in the TIF and other sites is reported. Thermal tests at ATCA board level and at rack level showed possible cooling issues when the system is pushed to its limits. FireFly 12-channel devices have been successfully qualified at 25 Gbit/s. The integration of the tracker FE and BE in the TIF showed excellent progress and the completion of the FE/BE vertical slice is expected in the next couple of years.

References

[1] CMS collaboration, *The Phase-2 Upgrade of the CMS Tracker*, CERN-LHCC-2017-009, CMS-TDR-014 (2017) https://doi.org/10.17181/CERN.QZ28.FLHW.

[2] CMS collaboration, *The Phase-2 Upgrade of the CMS Level-1 Trigger*, CERN-LHCC-2020-004, CMS-TDR-021 (2020).

[3] A. Rose et al., *Serenity: an ATCA prototyping platform for CMS Phase-2, Topical Workshop on Electronics for Particle Physics (TWEPP2018)*, *PoS* **343** (2018) 11.

[4] L. Calligaris et al., *A flexible and low-cost open-source IPMC mezzanine for ATCA boards based on OpenIPMC*, 2022 *JINST* **17** C03007.

[5] J.R. Cornelius et al., *Extended temperature performance of VCSEL based optical transmitter, 2016 IEEE Avionics and Vehicle Fiber-Optics and Photonics Conference (AVFOP)* (2016), pp. 279–280.

[6] R. Stephens, *Jitter Analysis: The dual-Dirac model, RJ/DJ, and Q-scale*, Agilent Technologies (2004).

[7] L. Olanterä et al., *Radiation hard optical link developments at CERN, 2017 17th European Conference on Radiation and Its Effects on Components and Systems (RADECS)* (2017), pp. 1–6.

[8] A. Albert et al., *The Apollo ATCA Platform, Topical Workshop on Electronics for Particle Physics (TWEPP2019)*, *PoS* (2020) 120.

2023 JINST 18 C01041