

**OPEN ACCESS**

## Lessons learned while developing the Serenity-S1 ATCA card

To cite this article: T. Mehner *et al* 2024 *JINST* **19** C02018

View the [article online](#) for updates and enhancements.



**PRIME**  
PACIFIC RIM MEETING  
ON ELECTROCHEMICAL  
AND SOLID STATE SCIENCE

HONOLULU, HI  
Oct 6–11, 2024

Abstract submission deadline:  
**April 12, 2024**

**Learn more and submit!**

**Joint Meeting of**

The Electrochemical Society  
•  
The Electrochemical Society of Japan  
•  
Korea Electrochemical Society

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS

GEREMEAS, SARDINIA, ITALY

1–6 OCTOBER 2023

## Lessons learned while developing the Serenity-S1

### ATCA card

T. Mehner<sup>1</sup>,<sup>a,\*</sup> L.E. Ardila-Perez<sup>2</sup>,<sup>a</sup> M. Balzer,<sup>a</sup> G. Fedi<sup>3</sup>,<sup>b</sup> M. Fuchs<sup>4</sup>,<sup>a</sup> A. Howard,<sup>b</sup> G. Iles<sup>5</sup>,<sup>b</sup> M. Loutit,<sup>6</sup> S. Mansbridge,<sup>b</sup> F. Palla,<sup>d</sup> D. Parker<sup>7</sup>,<sup>b</sup> M. Pesaresi<sup>8</sup>,<sup>b</sup> A. Rose,<sup>b</sup> M. Saleh,<sup>b</sup> O. Sander<sup>9</sup>,<sup>a</sup> M. Schleicher,<sup>a</sup> C. Strohman,<sup>f</sup> D. Tcherniakhovski,<sup>a</sup> T. Williams<sup>10</sup><sup>e</sup> and J. Zhao<sup>c</sup>

<sup>a</sup>Institute for Data Processing and Electronics, Karlsruhe Institute of Technology, Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Germany

<sup>b</sup>Physics Department, Blackett Laboratory, Imperial College London, Prince Consort Rd, London, SW7 2BW, U.K.

<sup>c</sup>Institute of High Energy Physics, Chinese Academy of Sciences, 19B Yuquan Road, Shijingshan District, Beijing, China

<sup>d</sup>Pisa Section, National Institute of Nuclear Physics, 56124 Pisa, Italy

<sup>e</sup>Particle Physics Department, STFC Rutherford Appleton Laboratory, Harwell Campus, Didcot, OX11 0QX, U.K.

<sup>f</sup>Department of Physics, Cornell University, 616 Thurston Ave., Ithaca, NY 14853, U.S.A.

<sup>g</sup>School of Physics, University of Bristol, Queens Road, Bristol, BS8 1QU, U.K.

E-mail: [torben.mehner@kit.edu](mailto:torben.mehner@kit.edu)

**ABSTRACT:** The Serenity-S1 is a production-optimised Advanced Telecommunications Computing Architecture (ATCA) processing blade based on the AMD Xilinx Virtex Ultrascale+ device. It incorporates many developments from the Serenity-A and Serenity-Z prototype cards and, where possible, adopts solutions being used across CERN. Due to the shortage of components during the recent semiconductor crisis, commonly used components in the prototypes had to be replaced by new ones after qualification. In this work, we discuss various improvements to simplify manufacturing, the performance of new components, some of the more difficult aspects of procurement, the performance of production-grade Samtec 25 Gb/s optical firefly parts, and concerns regarding the rack cooling infrastructure.

**KEYWORDS:** Detector control systems (detector and experiment monitoring and slow-control systems, architecture, hardware, algorithms, databases); Digital electronic circuits; Trigger concepts and systems (hardware and software)

ARXIV EPRINT: [2311.02222](https://arxiv.org/abs/2311.02222)

\*Corresponding author.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Board overview</b>	<b>1</b>
<b>3</b>	<b>Component selection challenges</b>	<b>2</b>
<b>4</b>	<b>PCB design</b>	<b>3</b>
<b>5</b>	<b>Cooling system design and testing</b>	<b>4</b>
<b>6</b>	<b>FireFly long trace performance</b>	<b>4</b>
<b>7</b>	<b>Testing and production plans</b>	<b>5</b>
<b>8</b>	<b>Conclusion</b>	<b>6</b>

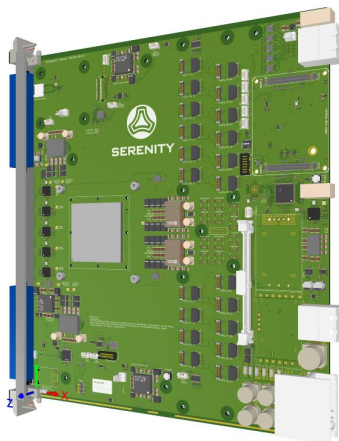
---

## 1 Introduction

The Serenity-S1 ATCA card is an integral component of the High-Luminosity Large Hadron Collider (HL-LHC) upgrade [1, 2]. A considerable number of 721 Serenity-S1 cards will be integrated into various sub-detector systems within the Compact Muon Solenoid (CMS) experiment. As a result, a versatile design is required, featuring optical high-speed data input and output capabilities as well as a large processing FPGA. The Serenity-S1 card evolved from its predecessors, namely the Serenity-A and Serenity-Z [3, 4] prototype cards, which were developed at Karlsruhe Institute of Technology (KIT) and Imperial College London, respectively. These prototype cards are already in use for firmware development and system integration in several CMS sub-detector working groups. The Serenity-S1 is the result of a synergistic partnership within the Serenity consortium, encompassing multiple institutes across numerous countries. This cooperative design process was facilitated by a Git-controlled Altium project, enabling concurrent work on the schematic and sequential work on the layout. Enabling this workflow was a key reason why we shifted to Altium for the development of Serenity-S1. Furthermore, the real-time monitoring of component availability that Altium offers was also used intensively, especially during the semiconductor crisis.

## 2 Board overview

The board design as seen in figure 1 incorporates segmentation into two distinct sections, namely the payload area towards the front panel and the service area towards the backplane. This segmentation was motivated by an initial plan to develop both a single and a double FPGA version. At this time, only the single FPGA version is foreseen within CMS. The division into payload area and service area is nevertheless advantageous, as the service area, which is dedicated to the slow control of the board, can be reused for future ATCA developments. Within the service area, an AMD Xilinx Kria



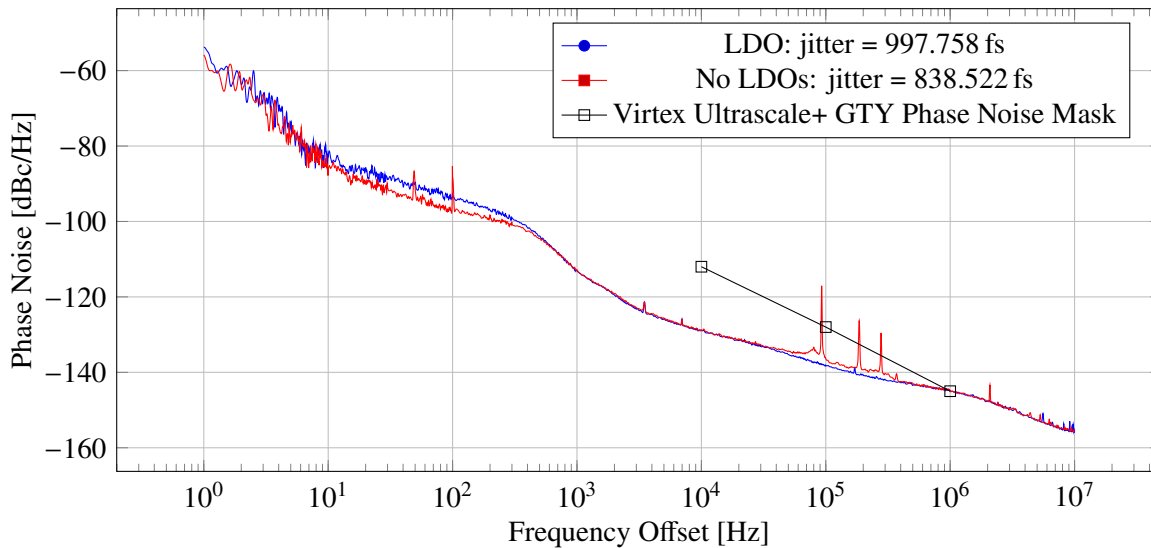
**Figure 1.** 3D render of the Serenity-S1 ATCA board.

K26 System-on-Module (SoM) is running a Linux operating system to manage the slow control tasks [3, 5]. This section also houses the required power supplies, an Ethernet PHY, an SD card slot, SSD storage, and clocking resources tailored for the Kria module. The second portion of the service area contains what is needed to meet the specific requirements of the ATCA standard. This includes an Intelligent Platform Management Controller (IPMC) named OpenIPMC in the DIMM form-factor [6], Zone-2 backplane connections, the necessary input power modules designed to harness the dual  $-48$  V shelf power source, and an Ethernet switch that allows connecting up to seven endpoints. The payload area is responsible for processing data with a high throughput and low latency. The main processing unit in this domain is an AMD Xilinx Virtex Ultrascale+ VU13P FPGA. Data transmission is handled by ten pairs of 12-channel unidirectional FireFly connectors and one additional 4-channel bidirectional connector. Furthermore, the payload area houses a synchronous and asynchronous clock tree, providing reference clocks to all multi-gigabit transceivers (MGTs) of the FPGA.

### 3 Component selection challenges

The initial choice for the clock distribution network, the Skyworks Si5395A zero-delay jitter cleaning phase-locked loop (PLL), became unavailable without any indication of its future availability. Subsequently, a suitable alternative was identified, the Microchip ZL30274, which features two internal PLLs and similar phase-noise performance. An evaluation card was developed and sent to the EP-ESE group at CERN for assessment of compatibility with the CMS requirements. Extensive testing, including measurements with and without the on-board low-dropout regulators to filter the switched-mode power supplies, shows that the ZL30274 met the demands of the Virtex Ultrascale+ FPGA [7] and the  $\leq 10$  ps jitter specification of CMS in general [8] (see figure 2). As a consequence, the ZL30274 emerged as a viable replacement for the Si5395A in the design.

Similar concerns arose regarding the core power supply, which is delivering 170 A at 0.85 V to the FPGA. While the initial choice involved two parallel LTM4700 power modules, their limited availability and a subsequent 70 % price increase motivated an exploration of alternatives. Due to the availability relaxing, the LTM4700 was kept in the design, but the option to use a multi-phase power supply that can be built using industry-standard components was added. The multi-phase supply can be mounted



**Figure 2.** Phase noise measurements for ZL30274 with and without LDO (1 Hz to 10 MHz).

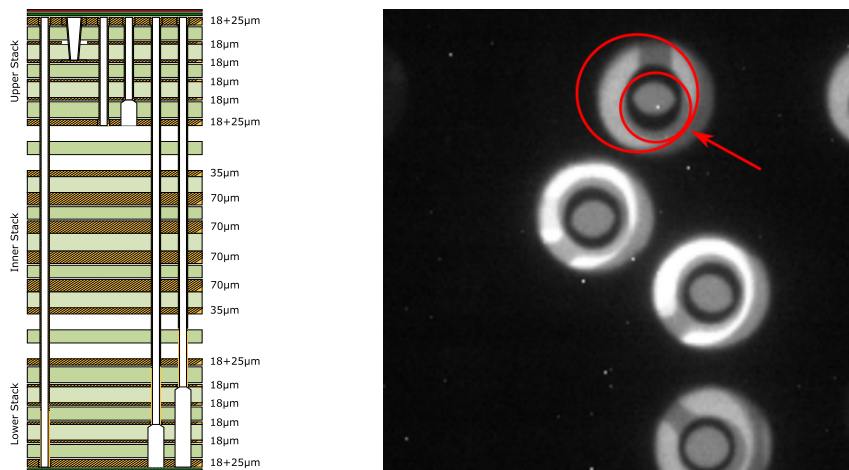
using Würth Electronic REDCUBE terminals that can transmit 50 A each. As soon as the newly designed multi-phase power supply has been successfully tested, it can replace the LTM4700 modules, which currently cost \$ 300 while the multi-phase power supply will be priced at \$ 40. This underscores the importance of planning for alternative solutions when component availability is uncertain.

Recently, the production of the standardised power connector for ATCA cards has been halted by multiple manufacturers. While Positronic and Conec still produce this connector, there is a worrying trend of diminishing support for ATCA connectors. In response, we adopted a multi-vendor-compliant approach by adapting our footprints and accounting for alignment pin positions. Additionally, the total number of ATCA connectors required for Serenity-S1 boards in CMS, plus a reserve of 15 %, has already been acquired.

#### 4 PCB design

The PCB design presented three significant challenges. First, the 124 high-speed links from the FPGA to the FireFly connectors require a minimum of three impedance-controlled, ground-shielded layers that are accessible from the top through stub-less vias. Second, to limit thermal losses due to high currents in the power planes, the copper thickness was increased from the standard 18  $\mu\text{m}$  used for the signal layers to 70  $\mu\text{m}$ . To further reduce power distribution losses, additional layers were used where feasible. Lastly, the stipulation by the experiment to use halogen-free materials posed a challenge due to the limited experience of PCB manufacturers with high-speed halogen-free materials, prompting close collaboration with manufacturers to define a suitable layer stack. The result is shown in figure 3 (left).

The resulting PCB contains three sub-stacks to satisfy layer stack symmetry while incorporating a copper-heavy inner stack and two impedance-controlled outer stacks. The outer two sub-stacks include four impedance-controlled layers accessible from the top via back-drilling and microvias. The inner stack conducts high currents across four 70  $\mu\text{m}$  copper layers, with two allocated for core supply current. To save cost, the PCB supplier initially intended to use a low-speed material for the inner layer but encountered complications due to mismatched expansion parameters between



**Figure 3.** Serenity-S1 layer stack (left) and x-ray picture showing a short of a via pad (small circle) with the copper flood (large circle) due to center stack expansion (right).

high-speed and low-speed materials. This resulted in shorts around the edge of the PCB, as seen on the right side of figure 3.

## 5 Cooling system design and testing

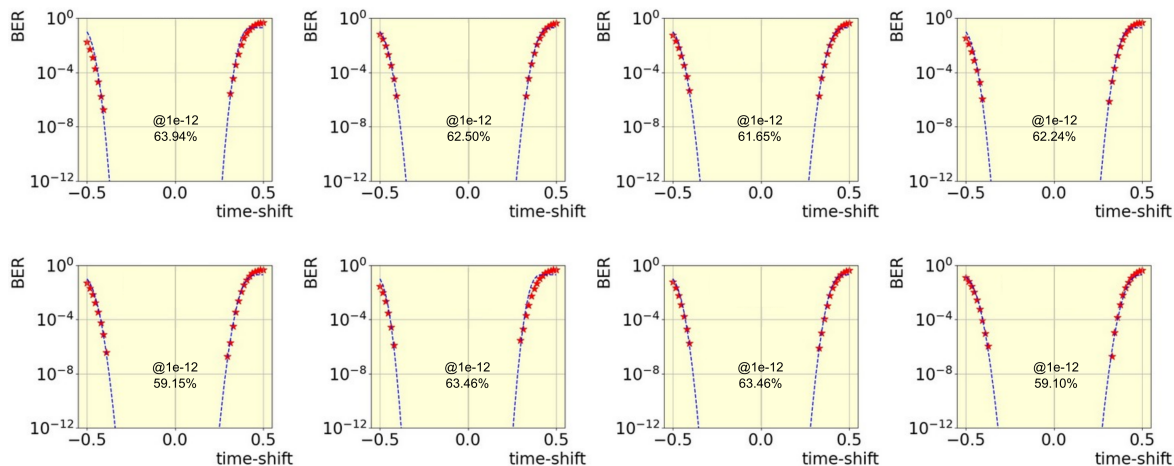
Cooling the card requires a custom heatsink solution to maintain a low profile, as the ATCA standard only allows components to elevate to a maximum of 21.33 mm above the PCB. Additionally, the ATCA crates, which house the cards, require heat extraction to achieve the desired operating temperatures. In the CMS setup, each rack accommodates two ATCA crates, cooled by internal air loops, and three water-cooled heat exchangers to remove the heat. Given the maximum allowed power usage of 10 kW in a full rack, an equivalent amount of heat must be removed. Maintaining FireFly optical devices below 50 °C is crucial for their longevity over the targeted 10 years of operation. This results in the boundary conditions that the cooling system must meet. To assess cooling efficiency, measurements were conducted using heater cards at the tracker integration facility (TIF) and at the CMS service cavern USC55 [9]. With the rack operating at maximum power and the heat exchangers supplied with water at 16 °C, it was initially only possible to cool the circulating air down to 30 °C. Increasing the rack water flow from 18 L/min to 30 L/min improved this to 26 °C. Larger-diameter pipes to the rack should yield a rack water flow of 34 L/min and a resulting air temperature of 25 °C. This should provide adequate headroom to cool the optical devices. Further improvements are unlikely because the water temperature cannot be lowered for legacy reasons, and it is also not possible to use higher pressure to increase water flow.

## 6 FireFly long trace performance

To safeguard the FireFly transceivers from the FPGA-generated heat, they are placed within separate air passages. However, this arrangement necessitates extended traces running from the FPGA to the FireFly connectors. On the Serenity-S1, these traces span 270 mm, prompting concerns regarding the ability to achieve 25 Gb/s links with a low bit error rate of  $10^{-12}$ . Benefiting from the similar



structure in Serenity-A, these long traces could be assessed in advance. A tool was developed to scan the optical performance for different values of the parameters Tx pre-emphasis, Tx post-emphasis, Tx main driver, and Rx Equalizer. The tool allows the evaluation of the operational regions of a device in parameter space and determines if all links work without error. The bathtub plots resulting from these evaluations (see figure 4) demonstrated an opening ranging between 59.10% and 63.94%, indicating that the long traces on the Serenity-S1 will likely be functional. This underscores the value of functional evaluation boards for testing purposes. The tests were conducted using prototype FireFly devices, with the first production devices anticipated at the end of October 2023.



**Figure 4.** Bathtub plot of Serenity-A long trace (260 mm).

## 7 Testing and production plans

The first set of 12 boards will imminently arrive from the PCB manufacturer at KIT and undergo assembly there. These initial boards will not have an FPGA populated to facilitate preliminary tests on the power supplies. Subsequently, once the FPGA is socketed or soldered onto the boards, further testing will be conducted before dispatching them to CERN and Imperial College for temperature cycling tests, optical assessments, and performance characterisations of the high-speed lanes.

The full-scale production is scheduled to commence in October 2024, following the tenders starting in February. A pre-series consisting of 50 boards will be the first to be manufactured, with the remaining boards planned for 2025. Separate tenders are being issued to PCB manufacturers and assemblers, with critical components provided to support pre-production, including all ATCA connectors, which will be provided for the entire production run. The high cost of the components necessitates a high yield. In order to achieve this, the board is being designed for ease of manufacture. Very small components are avoided where possible. Dedicated test points will allow in-circuit testing of power nets before power is applied. All power supplies are under I2C control, providing diagnostics across the board. Lastly, large through-hole components will not directly connect to power planes so that they are thermally isolated, allowing safe reflow and rework.

## 8 Conclusion

In conclusion, the Serenity-S1 ATCA card developed for the HL-LHC upgrade demonstrates adaptability in overcoming component selection challenges, complex printed circuit board construction, and cooling system complexities. The evaluation of alternative components and multi-vendor-compliant approaches enhances the resilience of the project. Additionally, by engaging in collaborative efforts with multiple PCB manufacturers, PCB production becomes more dependable and consistent. The construction of prototype cards facilitates early hardware and firmware evaluation, which is crucial for system integration studies at every sub-detector system. The Serenity-S1 project represents a collaborative and innovative approach, well-prepared to meet the demands of the HL-LHC upgrade.

## Acknowledgments

The authors would like to acknowledge the support from the CERN EP-ESE group and specially Philippa Hazell and Sophie Baron for the measurements of the phase-noise of the ZL30274 jitter cleaner. Marvin Fuchs acknowledges the support by the Doctoral School “Karlsruhe School of Elementary and Astroparticle Physics: Science and Technology”.

## References

- [1] CMS collaboration, *Architecture and prototype of the CMS Global Level-1 Trigger for Phase-2*, 2023 [JINST 18 C01034](#).
- [2] A. Ryd and L. Skinnari, *Tracking Triggers for the HL-LHC*, *Ann. Rev. Nucl. Part. Sci.* **70** (2020) 171 [[arXiv:2010.13557](#)].
- [3] SERENITY collaboration, *ZynqMP-based board-management mezzanines for Serenity ATCA-blades*, 2022 [JINST 17 C03009](#).
- [4] A. Rose et al., *Serenity: An ATCA prototyping platform for CMS Phase-2*, *PoS TWEPP2018* (2019) 115.
- [5] M. Fuchs, L.E. Ardila-Perez, T. Mehner and O. Sander, *Split Boot — True Network-Based Booting on Heterogeneous MPSoCs*, *IEEE Trans. Nucl. Sci.* **70** (2023) 1157 [[arXiv:2301.05642](#)].
- [6] L. Calligaris et al., *A flexible and low-cost open-source IPMC mezzanine for ATCA boards based on OpenIPMC*, 2022 [JINST 17 C03007](#) [[arXiv:2112.12888](#)].
- [7] AMD Xilinx, *Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics*, [DS923](#) (2021).
- [8] C.M.S. Collaboration, *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*, [CERN-LHCC-2021-007](#), [CMS-TDR-022](#), CERN, Geneva (2021).
- [9] CMS collaboration, *Lessons learnt from the first vertical slice of the CMS Outer Tracker*, 2023 [JINST 18 C01041](#).