

PREPARED FOR SUBMISSION TO JINST

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS TWEPP2024
 SEPTEMBER 30 - OCTOBER 4, 2024
 ONLINE

Design, Construction, and Testing of the APOLLO ATCA Blades for Use at the HL-LHC

**Alp Akpinar,^a Aymeric Blaizot,^a Serhii Cholak,^a Gianfranco de Castro,^a Zeynep Demiragli,^a Alec Duquette,^b Jonathan Richard Fulcher,^{a,1} Dan Gastler,^a Kristian Hahn,^c Eric Shearer Hazen,^a Si Hyun Jeon,^a Peace Kotamnives,^b Alexander Madorsky,^a David Monk,^c Sheena Noorudhin,^c Michael Oshiro,^b James Rohlf,^a Charles Ralph Strohman,^b Emily Minyun Tsai,^d Peter Wittich,^b Siqi Yuan,^a and Rui Zou^b on behalf of the CMS
 Collaboration**

^a*Department of Physics, Boston University, Commonwealth Ave, Boston, MA, U.S.A.*

^b*Department of Physics, Cornell University, Sciences Dr, Ithaca, NY, U.S.A.*

^c*Department of Physics, Northwestern University, Sheridan Rd, Evanston, IL, U.S.A.*

^d*Department of Physics, Northeastern University, Forsyth St, Boston, MA, U.S.A.*

E-mail: jonathan.fulcher@cern.ch

ABSTRACT: The Apollo Advanced Telecommunications Computing Architecture (ATCA) platform is an open-source design consisting of a generic “Service Module” (SM) and a customizable “Command Module” (CM), allowing for cost-effective use in applications such as the readout of the inner tracker and the Level-1 track trigger for the CMS Phase-II upgrade at the HL-LHC. The SM integrates an intelligent IPMC, robust power entry and conditioning systems, a powerful system-on-module computer, and flexible clock and communication infrastructure. The CM is designed around two Xilinx Ultrascale+ FPGAs and high-density, high-bandwidth optical transceivers capable of 25 Gb/s. Crates of Apollo blades are currently being tested at Boston University, Cornell University, and CERN.

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

ARXIV EPRINT: [2501.03702](https://arxiv.org/abs/2501.03702)

¹speaker and corresponding author

Contents

1	Introduction	1
2	Overview and CMS Applications	1
3	Revision 2 SM and CM Performance Tests and Results	2
3.1	Hardware and Firmware Tests	2
3.2	Software Framework Tests	4
4	Revision 3 SM Performance Tests and CM Plans	4
5	Revision 3 Preproduction and Plans	5
6	Summary	6

1 Introduction

High-performance Advanced Telecommunications Computing Architecture (ATCA) blades [1] are becoming essential in high-energy physics experiments, particularly for the High-Luminosity LHC. The Apollo platform, aptly named after the Apollo space program[2], addresses challenges in power, cooling, optical fiber management, and communication interfaces by providing a streamlined hardware environment along with firmware and software tool kits. It features a versatile “Service Module” (SM) that manages ATCA communications, power distribution, and clocking, paired with an application-specific “Command Module” (CM) equipped with FPGAs and numerous optical links for high data rates (see section 3).

The SM is a standard-sized ATCA blade that supports a CM within a 7U × 180 mm cutout. It includes an Enclustra Zynq Ultrascale+ System-on-Module (SoM) with an embedded Linux OS for control and monitoring, along with an Intelligent Platform Management Controller (IPMC), OpenIPMC [3], and a Wisconsin Ethernet Switch Module [4]. The SM provides flexible interfaces such as gigabit Ethernet, additional Ethernet to the switch, and four 10 Gb/s bidirectional serial links to the CM, along with diagnostic and programming interfaces like asynchronous serial, I²C, and JTAG. Power delivery is managed by commercial modules, supplying up to 12 V_{DC} at 30 A, while the platform supports more than 100 optical links at speeds up to 25 Gb/s to meet the high data throughput demands of experiments.

2 Overview and CMS Applications

The CMS experiment will undergo major upgrades for the High-Luminosity LHC [5], including a new silicon tracker [6]. The Apollo platform will be used in four CMS subsystems: the Level-1 Track Finder (TF), the Inner Tracker Data Trigger and Control system (IT-DTC), the Precision Proton Spectrometer (PPS), and the high-precision luminosity monitoring system (BRIL) [7].

The Level-1 TF reconstructs tracks from the Outer Tracker and calculates track parameters for the Level-1 trigger system. Although the algorithm requires substantial FPGA resources, the total power consumption should remain within operating limits (< 100 W). It requires two Virtex Ultrascale+ FPGAs and 62 optical links at 25 Gb/s and a total of 162 Apollo blades.

The IT-DTC reads hits from the front-end pixel ASIC chips and converts the data into a compact format. It also converts the trigger signal from the CMS trigger to tokens and transmits them to the front-end. It requires up to 72 LpGBT [8] optical links at 10 Gb/s, 16 optical links at 25 Gb/s for central data acquisition (CDAQ), 4 optical links on each of the TEPX DTCs at 25 Gb/s for BRIL "LUMI" and totals 36 Apollo boards (figure 1).

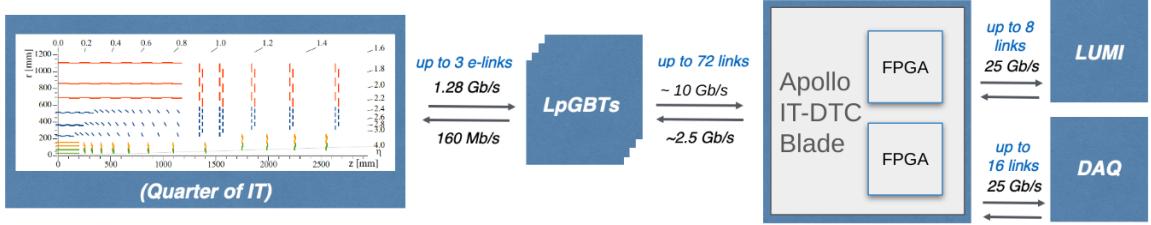


Figure 1. Example CMS Application IT-DTC.

The BRIL Lumi system receives event fragments on 4 to 8 25 Gb/s transceivers per blade from IT-DTC with dedicated luminosity processing hardware. These 25 "luminosity blade" Apollos perform real-time pixel clustering and cluster counting on CM FPGAs. The data are then decoded and processed. Additionally, the PPS also made a recent request for three Apollos.

To satisfy the requirements of these applications, the CMS version of the Apollo CM can accommodate one or two large Xilinx Ultrascale+ FPGAs with 4-channel transceiver and 12-channel transmitters and receiver Firefly optical engines that total 112 optical links with speeds up to 25 Gb/s. Two revisions have been tested, and a third revision is under development, of which the first three blades should be assembled in Q4 of 2024. Revision 2 (Rev2) was produced in Q3 of 2021 and has undergone thorough testing and evaluation. This paper describes the evaluation, shares what we learned, presents the performance of the Rev2 design (section 3), and discusses the major updates and the test status of the revision 3 (Rev3) design (section 4).

3 Revision 2 SM and CM Performance Tests and Results

The CMS Rev2 Apollo blade (figure 2) has two large FPGAs (XCVU13P, A2577 package) and sites for four- and twelve-channel Firefly optical transceivers with up to 124 links supporting speeds up to 25 Gb/s. A comprehensive set of performance tests was conducted on Rev2 to validate data transfer integrity, clock distribution, and microcontroller (MCU) code performance. Thorough thermal tests first presented in [9] were also repeated. This section summarizes the testing procedure and the final results.

3.1 Hardware and Firmware Tests

Link integrity tests focused on high-speed data transfer across FireFly optical links and FPGA interconnects, achieving bit error rates (BER) less than 10^{-16} at 25 Gb/s and 10 Gb/s. Extensive

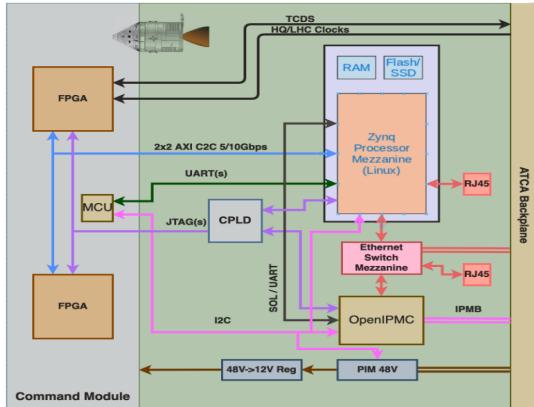


Figure 2. Apollo schematic.

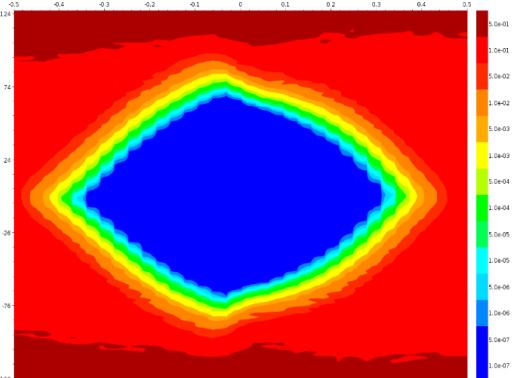


Figure 3. Typical eye with large opening at 10 Gb/s over TCDS lines from backplane on the Rev3 SM.

testing of the CMS Standard Protocol (CSP)[10] links was performed. Copper links were tested in single-FPGA loop-back and cross-FPGA configurations at both Boston University (BU) and CERN. Similarly, the 4-channel-transceiver FireFly optical modules were tested in single-FPGA loop-back and cross-FPGA configurations, along with 12-channel 25Gb/s FireFly modules with single-FPGA-loop-back and reverse-loop-back tests at CERN. All tests used both pseudo-random binary sequence (PRBS), and CSP with the CMS Clock and Trigger Distribution System (TCDS2)[11] modes. The CSP tests, which required correct timing alignment, were monitored and controlled through the CLI and Apollo Herd interfaces (see section 3.2). Chip-to-chip (C2C) links between the Zynq and FPGAs were also validated at 5Gb/s. Advanced eXtensible Interface (AXI) C2C link tests were conducted from the Zynq Linux OS to endpoints in the SM and CM, with the goal of evaluating data integrity and write speed. Although the current AXI C2C links do not employ Direct Memory Access (DMA), the test utility supported a variety of block size configurations, and dedicated firmware with Block-RAM (BRAM) was developed to support block write tests. Extended tests showed no errors over 96 hours (343 GB) and 14 hours (47 GB), respectively. However, the maximum write speed achieved was 24 Mb/s with block sizes larger than 4096 words. Future plans include developing an AXI DMA driver to increase transfer speeds closer to the selected physical link rate. The clock distribution test was performed using TCDS2, integrated within the CMS back-end system firmware framework (EMP)[10]. The test involved receiving clock and data stream signals from a DTH[12] via the backplane to FPGA1 on a CM, with a secondary relay to the second FPGA. Version 0.1.1 of the TCDS2 firmware successfully verified the lightweight endpoint on both FPGAs and demonstrated proper recovery of the 40 MHz LHC clock signal from the TCDS2 stream along with successful operation of the TCDS2 relay. More advanced testing with the latest version 0.2.0rc1 introduced the full endpoint and confirmed recovery of the 40 MHz clock stream on both FPGAs, with successful testing of both the internal and redistributed LHC 40 MHz and 320 MHz reference clocks. Further development is ongoing for an upgraded inter-FPGA relay compatible with the latest TCDS2 firmware.

Thermal performance tests were conducted at CERN to evaluate power and cooling efficiency (see figure 4). The test setup involved fixing the fan speed to 10/15 (67% of maximum), maintaining an inlet air temperature of 23°C, and running the heater firmware[13] on both FPGAs. The heaters were gradually activated until both FPGAs reached 85°C. At this point, the power consumption of the FPGAs and the maximum temperature of the FireFly modules were recorded, with the latter

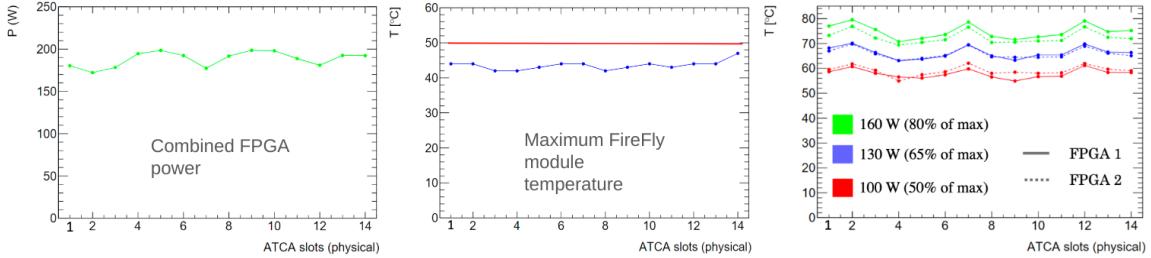


Figure 4. Temperature vs power as a function of ATCA slot for both VU13P FPGAs, and FireFly module compared to the recommended operating range of below 50 °C, of the Rev2 CM.

staying below 50°C (47°C maximum), demonstrating minimal thermal coupling to the FPGAs. Furthermore, FPGA temperatures were measured at three more realistic combined power levels, and the tests were repeated in different slots to verify uniform behavior throughout. The airflow impedance was found to be slightly high (0.26 in-H₂O @ 30 CFM). However, the system was able to accommodate up to 200W of FPGA power with a split of approximately 55% to FPGA1 and 45% to FPGA2, confirming the platform’s ability to manage significant thermal loads effectively (see figure 4). The MCU code was also tested, ensuring the proper configuration and sequencing of power supplies, as well as the correct setup and monitoring of FireFly modules and clocking circuits. Furthermore, environmental data was successfully transmitted to both the IPMC and the Zynq, ensuring robust control and monitoring of the system.

Lastly, the platform passed a significant milestone by validating the LpGBT slow control and data acquisition front-end link, which will be used in the Inner Tracker upgrade, by exercising the link to write and read registers on a front-end module. This test demonstrated that the Apollo platform is capable of handling the communications necessary for future CMS applications.

3.2 Software Framework Tests

The Apollo platform integrates seamlessly with the CMS Online Software Framework, known as EMP Shep/Herd, via a plugin called Apollo Herd, which provides a set of explicit commands to control and monitor various aspects of the system. These commands include powering on/off and programming the FPGAs, and running algorithm or link integrity tests. The software also allows direct access to Apollo’s address tables through a custom command line interface (CLI) called “BUTool”, providing read/write register access and facilitating low-level control. Several abstracted objects are utilized within the framework, including the SM and MCU, for monitoring status tables from BUTool; the CM FPGAs, for programming and conducting link tests; and both copper links and FireFly optical modules, serving as link test endpoints and monitoring targets.

Monitoring functions include real-time status tables, FireFly module temperatures, and optical power levels, as well as key FPGA parameters such as firmware versions, build timestamps, clock frequencies, and temperatures; Ensuring a robust environment for both control and monitoring.

4 Revision 3 SM Performance Tests and CM Plans

The Rev3 SM underwent a series of tests after its initial receipt in July of this year. Changes in Rev3 relative to Rev2 are minimal and consist of streamlining the routing of the TCDS2 link to improve link quality by eliminating unnecessary switching. Initial testing focused on verifying

the basic functionality of three units, starting with resistance and power measurements to ensure proper electrical behavior. The Ethernet switch and IPMC were evaluated to confirm stability. Ethernet connectivity was also tested to ensure seamless network integration. Furthermore, the Xilinx Zynq SoC was tested, including successful boot of the operating system (OS). A link eyescan was performed to validate TCDS2 signal integrity, and the new over-voltage and under-voltage protection circuits were thoroughly tested to ensure they functioned correctly within the specified limits. The over-voltage condition was simulated by changing the ratio of the resistive divider in the sensor circuit. The under-voltage condition was tested by simply turning the payload DC-DC converter off and checking that the sensor had tripped at the appropriate voltage using a scope. The sensors readout from the SM was also tested to verify that all monitored parameters were accurately reported. In parallel, the Rev3 SM was tested in combination with the Rev2 CM. These tests included programming of both FPGAs via JTAG, as well as read/write operations over the C2C AXI interface, and similar read/write operations over IPBUS, to both FPGAs. C2C link tests were conducted to assess the integrity and performance of the connections between the SM and CM. The connection to the MCU was also tested, along with sensor readouts from the CM to ensure comprehensive monitoring capabilities. These tests demonstrated the compatibility and communication between the revised SM and the older CM.

Several steps were completed to ensure the SM’s power delivery system functions correctly. The testing process involved connecting the IPMC module and enabling payload power through a jumper. The Complex Programmable Logic Device (CPLD) was then programmed, followed by a power-up check to verify that the IPMC had correctly enabled payload power. After this, the SoC was connected, and power was checked to confirm that all required voltages were properly enabled. Additionally, the Rev2 CM was connected to facilitate further testing.

After installing the SM in the crate, Ethernet access was verified and the TCDS2 system was checked. Integrated Bit Error Ratio Test (IBERT) tests were conducted on the TCDS2 backplane link to assess the performance between the DTC and the Apollo SM, with the quality of the signal plotted against the crate slot. The worst-case scenario was evaluated with a long-term test in which the BER was measured as far as 2×10^{-15} , with no error detected. Eye-plot shapes were checked and no error was seen. The results confirmed reliable communication between the Rev3 SM and the DTH using PRBS-31 patterns, aided by Decision Feedback Equalization (DFE); see figure 3.

5 Revision 3 Preproduction and Plans

The strategy for Apollo preproduction is to synchronize across all systems to reduce procurement and production costs. A pilot production run of approximately 30 boards is anticipated, primarily for the TF, with expected completion in Q1 2025. Production quantities include 162 TF boards plus 10 spares, 36 IT-DTC boards plus 5 spares, and 16 BRIL boards plus 4 spares, with the bulk expected in Q3 of 2026. Following completion of the currently ongoing Rev3 SM testing, full pre-production numbers will be initiated. The plan includes production of all PCB boards for the CM. Three boards will be populated, one without FPGAs. These boards will be tested with the final production versions of the Samtec Firefly 12-channel, 25 Gb/s optical parts once they become available in early 2025. Thorough testing of the Rev3 SM and CM is needed due to layout changes, most notably the addition of two inter-FPGA copper links, bringing the total to 56, and

an improvement to the distribution and layout of data and clock signals for TCDS2, adding more choice options for the master LHC clock source, while enhancing the symmetry between the two FPGAs, thus minimizing firmware design differences. For Rev3 SM changes, see section 4.

6 Summary

The Apollo ATCA blade will be used for the IT-DTC, TF, PPS, and BRIL systems of the CMS experiment at the High-Luminosity LHC (HL-LHC). Rev3 of the designs for the SM and CM are now finalized for preproduction. The first two revisions have demonstrated excellent link integrity, thermal performance, and adaptability to various requirements. Currently, Apollo Rev2 is being used for development and testing at multiple locations, including Boston University (BU), Cornell, and CERN. The SM Rev3 has been produced and is being tested at BU, while the CM Rev3 is nearing manufacture. The validation of the SM board is complete for Rev2 and is ongoing for Rev3.

The SM Zynq firmware has undergone extensive testing, while testing of the CM TCDS2 clocking scheme is well advanced, leading to high confidence in the final design. The CMS application framework, EMP, is also maturing, with testing being conducted for IT-DTC, TF and BRIL. Initial testing of the IT-DTC DAQ chain is scheduled at CERN in Q4 of 2024. Software CLI tools are very mature and well tested, with the EMP board-level plugin nearly complete and currently under test. Application-specific EMP plugins will follow.

Acknowledgments

We would like to thank the authors of ref. [13] for the code used in our power tests. This work was supported by the National Science Foundation under grants NSF-PHY-1946735, NSF-PHY-2209443, and NSF-PHY-2046626 and by the Department of Energy under grants DE-SC0021174 and DE-SC0016021.

References

- [1] A. Albert et al., *The Apollo ATCA Platform*, PoS (TWEPP2019) 120 (2020), [doi:10.22323/1.370.0120](https://doi.org/10.22323/1.370.0120).
- [2] N.A.S.A., *Apollo Space Program*, www.nasa.gov/the-apollo-program.
- [3] L. Calligaris et al., *OpenIPMC: A Free and Open-Source Intelligent Platform Management Controller Software*, IEEE Transactions on Nuclear Science, Volume 68, No. 8, pp. 2105-2112 (2021), [doi:10.1109/TNS.2021.3092689](https://doi.org/10.1109/TNS.2021.3092689).
- [4] M. Vicente et al., *Next Generation ATCA Control Infrastructure for the CMS Phase-2 Upgrade*, PoS (TWEPP2017) 102 (2018), [doi:10.22323/1.313.0102](https://doi.org/10.22323/1.313.0102).
- [5] The CMS Collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, [CERN-LHCC-2015-010, LHCC-P-008 \(2015\)](https://cds.cern.ch/record/2015-010).
- [6] The CMS Collaboration, *The Phase-II Upgrade of the CMS Tracker - Technical Design Report*, [CERN-LHCC-2017-009, CMS-TDR-014 \(2017\)](https://cds.cern.ch/record/2017-009).
- [7] The CMS Collaboration, *High-precision luminosity instrumentation for the CMS detector at the HL-LHC*, [doi:10.1016/j.nima.2022.167024](https://doi.org/10.1016/j.nima.2022.167024).

- [8] P. Moreira et al., *The lpGBT: a radiation tolerant ASIC for data, timing, trigger and control*, PoS (TWEPP2019) 120 (2020), <https://indico.cern.ch/event/799025/contributions/3486153/>
- [9] A. Albert et al., *The Apollo ATCA Design for the CMS Track Finder and the Pixel Readout at the HL-LHC*, PoS (TWEPP2021) 122 (2022), [doi:10.1088/1748-0221/17/04/C04033](https://doi.org/10.1088/1748-0221/17/04/C04033).
- [10] The CMS Collaboration, *The Phase-II Upgrade of the CMS Level-1 Trigger*, [CERNLHCC-2020-004, CMS-TDR-004 \(2020\)](https://cds.cern.ch/record/2684473).
- [11] CMS collaboration, *Progress in Design and Testing of the DAQ and Data-Flow Control for the Phase-2 Upgrade of the CMS Experiment*, [IEEE Trans. Nucl. Sci. 70 \(2023\) 914](https://doi.org/10.1109/TNS.2023.3645001).
- [12] J. Hegeman et al., *First measurements with the CMS DAQ and Timing Hub prototype-1*, PoS TWEPP2019 (2020) 111, [doi:10.22323/1.370.0111](https://doi.org/10.22323/1.370.0111).
- [13] J. Henkel et al., *Design and architectures for dependable embedded systems*, PoS (CODES+ISSS '11) pp. 69-78 (2011), [doi:10.1145/2039370.2039384](https://doi.org/10.1145/2039370.2039384).