



# The APOLLO ATCA Blade for Use at the HL-LHC

Engineers & Developers:

J.Fulcher\*, A.Akpinar, G.de Castro, S.Cholak, Z.Demiragli, D.Gastler, E.Hazen, A.Madorski, J.Rohlf, S.Yuan  
*Boston University*

A.Duquette, P.Kotamnives, M.Oshiro, C.Strohman, P.Wittich, R.Zou *Cornell University*

K.Hann, D.Monk, S.Noorudhin *Northwestern*

on behalf of the CMS Apollo Team



# Outline

- Brief reminder of what Apollo is and its applications
- An example application for Apollo
- Slightly more in-depth view of Apollo's components
- A look at on-board clock distribution for CMS
- Discussion of some tests performed
- A window on production schedule and quantities
- Summary

# Apollo ATCA Blade



- The Apollo project aims to provide a common ATCA blade to be used for readout and triggering applications in LHC experiments
- The project was previously presented in TWEPP [2019](#) and [2021](#) and is currently working on **3rd Revision**
- It is named after the Apollo program CSM spacecraft that comprised Command and Service modules and hence itself comprises the ATCA Service Module and various application-specific Command Modules
- Current Apollo applications:
  - CMS
    - [Inner Tracker \( CMS IT-DTC \) \[1\]](#)
    - [Track Finder \( CMS TF \) \[2\]](#)
    - [Lumi \(BRIL\) \[3\]](#)
  - ATLAS [LOMDT \[4\]](#)
  - With requests for others...
- This presentation is oriented towards the CMS applications



[1] [The Phase-2 Upgrade of the CMS Tracker](#)

[2] [The Level-1 Track Finder for the CMS High-Luminosity LHC Upgrade](#)

[3] [High-precision luminosity instrumentation for the CMS detector at the HL-LHC](#)

[4] [The MDT Trigger Processor upgrade for ATLAS Muons at the HL-LHC](#)

# Apollo: Hardware Features


Apollo Hardware has a modular design with two main components: [Service Module](#), and [Command Module](#).

## Service Module (Boston U.):

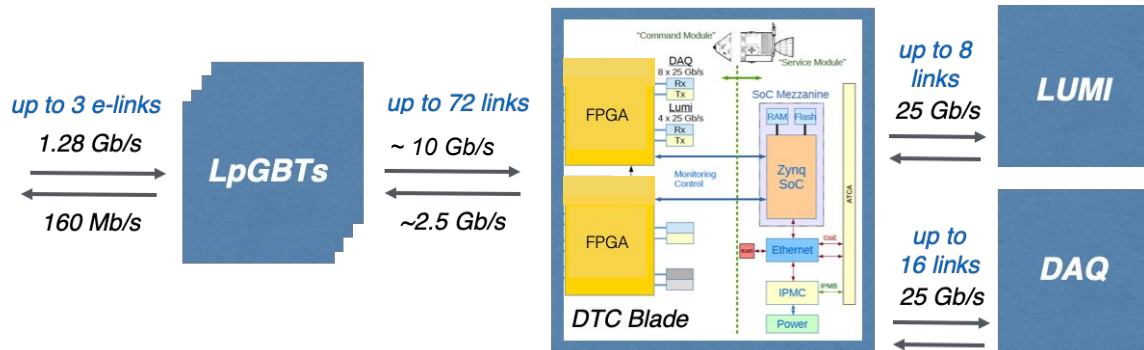
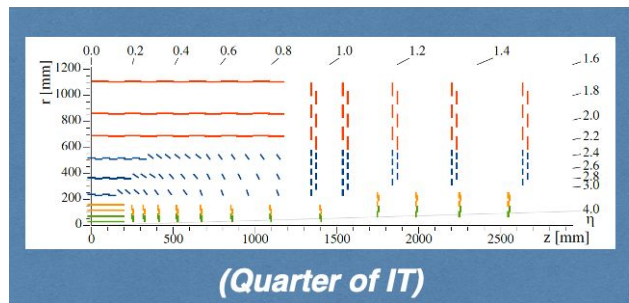
- ATCA interface
- 400 W power (12V to Command Module)
- UltraScale+ Zynq (Enclustra) SoC
- OpenIPMC & Wisconsin ESM Ethernet Switch

## Command Module (Cornell):

- Support for large FPGAs: VU13P (32 MGT Quads)
- Transceivers per FPGA: 56 copper inter-fpga links, 56 FireFly (4x12 channel, 2x4 channel)
- [TCDS2](#) [1] Support with FPGA 1 as endpoint



# Example Application - CMS Inner Tracker (IT) Backend



The IT Data Trigger Control (DTC) system consists of 36 Apollo boards:

- Receives and packs IT data from on-detector electronics via up to 72 bidirectional IpGBT links
- Communicates with the IT on-detector electronics
- In the downstream direction, processes two data streams: to DAQ & LUMI


The aggregated compressed data bandwidth is estimated to be ~100–300 Gb/s per DTC

**Total IT data bandwidth ~ 6.5 Tb/s**



# Apollo ATCA Blade Test Stands

@ CERN, BU, Cornell

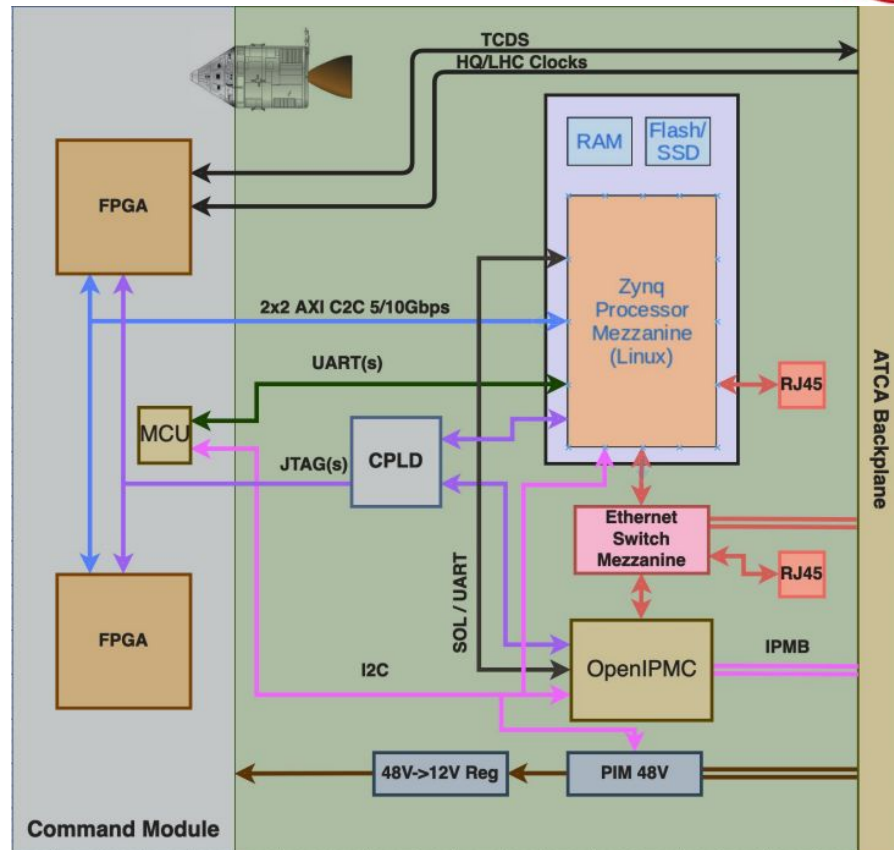


Cornell

# Service Module in Detail - Boston U

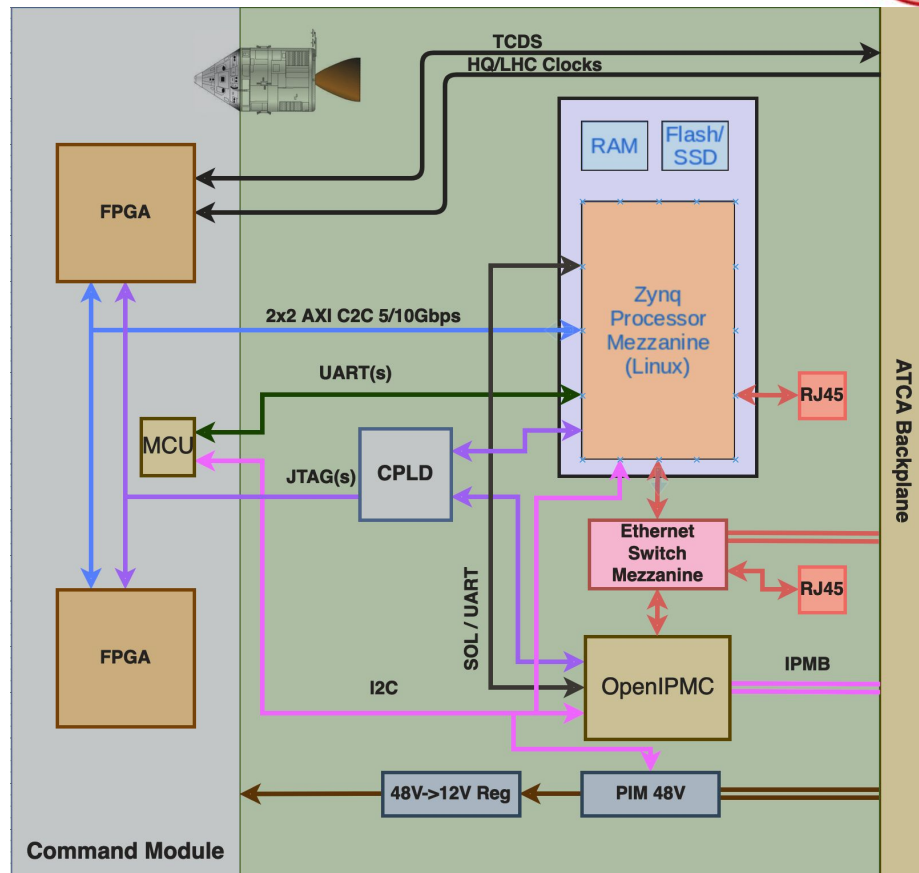


- ATCA Power entry module: 12V DC-DC converter
- Zynq System-on-Chip:
  - Enclustra ME-XU8-7EV-21-D12E-R2-1
  - OS Linux (Alma 8 -> 9) (Petalinux kernel)
  - Storage:  $\mu$ SD, EMMC or SSD
  - Option to boot and run from network file system
  - Expert UART interface for control and status
- High-speed connectors to CM
  - AXI bus via Chip2Chip connection, 5/10G link
  - LHC clocks (320MHz -> 40MHz)
  - TCDS 10.24G link
- Low-speed connector to CM
  - JTAG - for FPGA programming and debugging via XVC
  - 2 Serial Ports for information exchange between ZYNQ and CM MCU
  - I2C bus for sensor readout and control
  - General Purpose IOs
- Open IPMC module
  - Reading and reporting sensors
  - Power up/down and rebooting ZYNQ
- Ethernet Switch
  - Two gigabit connections via backplane
  - Connections to ZYNQ and IPMC modules



# Zynq Firmware: Apollo SM


- Gitlab Development Environment
  - CI Build pipelines used in all projects
  - SM Zynq: Firmware, OS file system and software tools
- Linux OS
  - Currently Alma 8
- Storage:
  - SSD support
  - First stage boot from eMMC or µSD
  - Option to boot and run from network file system
- DMA Drivers:
  - **JTAG DMA Driver:**
    - Allows to program CM FPGAs faster → **O(20sec) vs O(300sec)**
    - Implemented in the SM FW
  - **AXI C2C DMA Driver (Very Preliminary!!):**
    - Huge increase in R/W speed **20Mb/s vs 1Gb/s**
    - Will be crucial for the FE calib/config
    - Status: WIP on the SM FW implementation





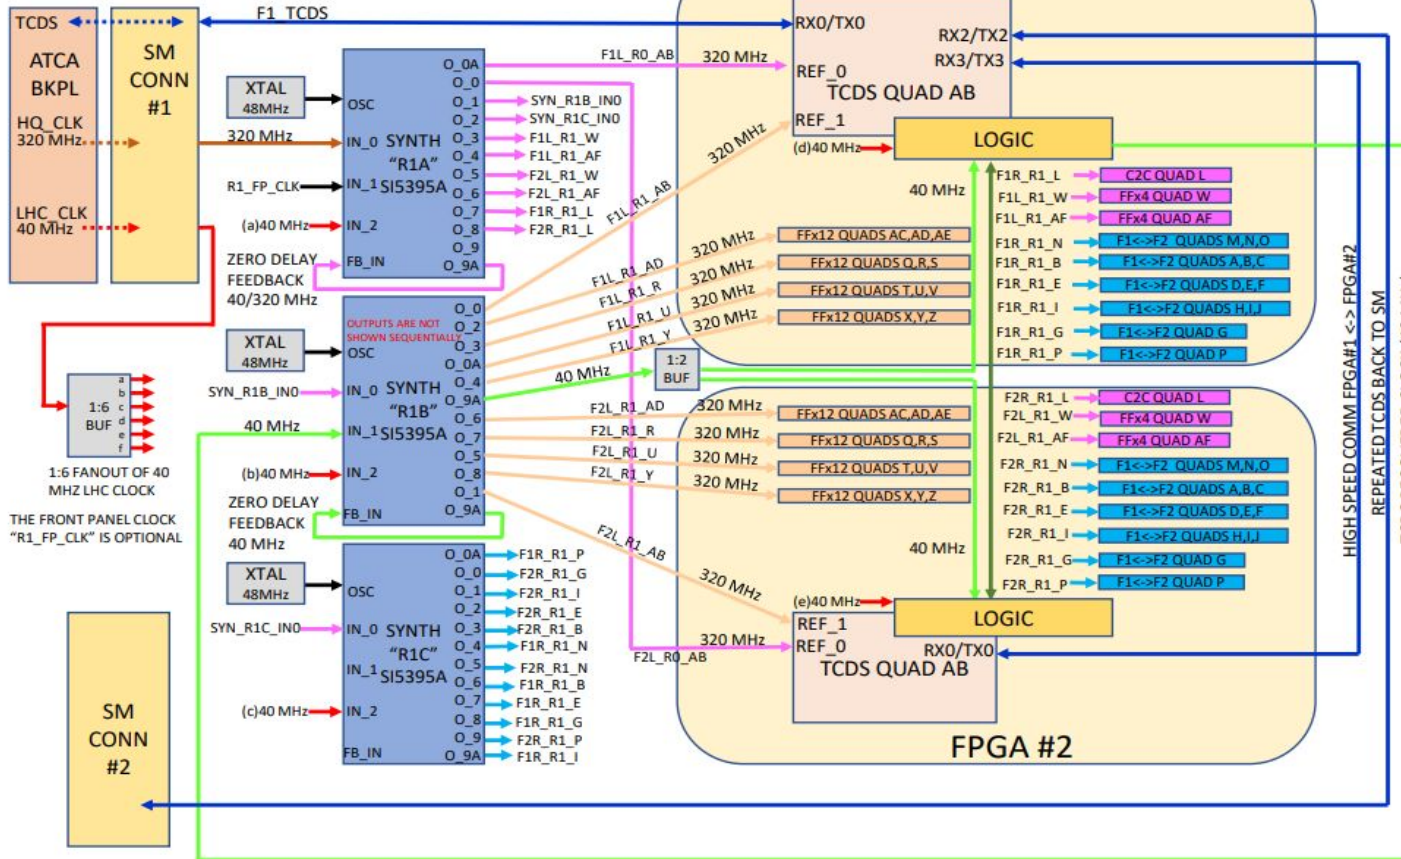
# CMS Command Module - Cornell

- 2 VU13P FPGAs:
  - -2 speed for CMS TF & BRIL
  - -1 speed for CMS IT-DTC
- 4x12ch + 2x4ch FireFly sites per FPGA
  - Dual voltage (3.3v or 3.75v) for 12ch transmit
- Optical Layout:
  - 112 optical links at up to 25G
  - 54 copper FPGA-FPGA links at up to 25 Gb/s
  - 2 copper FPGA-FPGA links reserved for TCDS TTC/TTS relay
- PCB
  - 18 layers, EM-890 halogen free
  - Conventional build - No HDI technology (no sequential lamination nor blind/buried vias)
- FW for CMS
  - [EMP](#)[1] Framework
  - FW Infrastructure & SW Tools



# Apollo CMv3: External Clock Sources

## ATCA Clock and TCDS Clock/Data



- GBIT TCDS CLOCK/DATA COMBINED
- LHC 40 MHz ATCA CLOCK
  - This is the 40 MHz clock passed from the ATCA backplane through the SM. It is fanned out through a 1:6 buffer. The frequency of this clock changes when the LHC is ramping.
- LHC 320 MHz ATCA CLOCK
  - This is the 320 MHz clock passed from the ATCA backplane through the SM. The frequency of this clock changes when the LHC is ramping.
- 320 MHz REF CLOCK
  - If the top synthesizer is using the clock on "IN\_0", then these 320 MHz clocks all have zero phase offset relative to the incoming LHC REFERENCE CLOCK signal.
- 40 MHz TCDS RECOVERED CLOCK
  - This clock is recovered from the incoming TCDS signal. The TCDS LOGIC synchronizes this clock to the bunch crossing. It also adjusts the phase to compensate for distribution delay changes. It will always maintain a fixed phase relative to the bunch crossing. The frequency also varies during filling. This clock is made available to the logic in the FPGAs for synchronizing operations.
- 320 MHz TCDS RECOVERED CLOCK
  - These clocks drive the detector-facing FireFly devices, as well as the quad that sends the outgoing TCDS signal back to the SM. They track the TCDS RECOVERED CLOCK.
- FPGA TO FPGA R1 CLOCK
  - These clocks drive the R1 reference for the FPGA quads that connect to the other FPGA. The frequency follows the 320 MHz REF CLOCK.
- TTC/TTS DATA/CONTROL
  - These signals possibly contain clocks/data/control extracted from the incoming TCDS signal (TTC) or destined for the outgoing TCDS signal (TTS). They are used within each FPGA, and can also pass from one FPGA to the other.
- OTHER CLOCKS
  - These include the front panel inputs and outputs of various crystal oscillators. These can be used for testing or for adding flexibility to the synthesizer outputs.
- GTY QUADS
  - FFX12 QUADS 12-lane FireFlies. For the IT-DTC, these will be detector facing.
  - FFX4 QUADS 4-lane FireFlies.
- F1<=>F2 QUADS Connections between the two FPGAs.
- TCDS QUAD Dedicated for TCDS function.

# Apollo: Test Summary



Summary of the conducted tests. Details of tests in later slides.

- **AXI C2C Tests:**
  - Data integrity: BER
  - Speed tests SM <-> CM
- **Link Tests:**
  - FireFly sites tested at 25G and 10G to BER<10<sup>-16</sup>
  - FPGA-to-FPGA links tested at 25G to BER<10<sup>-16</sup>
  - C2C links between Zynq and FPGAs tested at 10G and 5G
- **Clocking Test:**
  - TCDS2 links tested at 10G
  - Clock chips tested with on-board oscillators and external sources
- **IT-DTC LPGBT RD53 Milestone:**
  - Tested the Apollo <-> LPGBT <-> Front End link that will be used in CMS IT
- **Thermal Test:**
  - Power supplies tested
  - FPGA cooling verified
  - FireFly cooling verified
- **MCU code tested**
  - Power supply configuration & Power sequencing and monitoring
  - FireFly and Clocking configuration and monitoring
  - Environment information to the IPMC and the Zynq
- **Service Module functionality tests**

# Apollo SM - CM BER / speed test (Rev2)



- Rate tests from Zynq Linux OS to endpoints in the SM and CM were carried out
- Currently AXI C2C links do not employ DMA, which is under development
- Objective: Test SM -> C2C bit error rate and write speed
- Status:
  - Test utility capable of a variety of tests from SM to SM or CM with varying block sizes
  - New CM firmware with BRAMs to test block writes
  - Tests on SM211 and SM230 showed no errors with 343 GB (96 hr) and 47 GB (14 hr) respectively
  - Current max write speed: 24 Mb/s with block size > 4096 words
- Follow up:
  - Run tests on Rev3 boards when ready
  - Develop and test AXI DMA driver for SM to CM C2C links that allows for a transfer speed closer to the physical link speed of 5 Gb/s

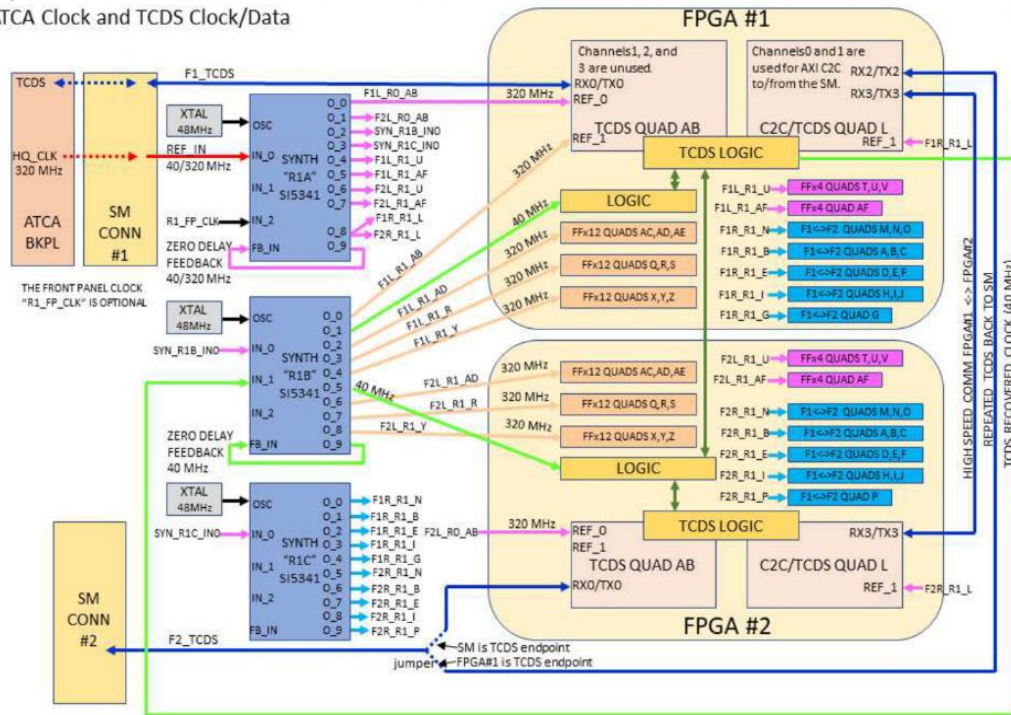


# Clocking Testing of CM Rev2

- Testing CMS Clock and Trigger distribution System ( [TCDS2](#) ) [1] within the CMS back end system firmware framework [EMP](#) [2]
- DTH [1] -> Backplane -> CM FPGA1
- CM FPGA1 Relay -> FPGA2
- TCDS2 cms-tc2ds-firmware V0.1.1:
  - Lightweight Endpoint only
  - Tested shown working for both FGPA1
  - Use recovered clock directly
  - TCDS2 Relay working
- TCDS2 cms-tc2ds-firmware V0.2.0rc1:
  - Full Endpoint
  - Tested receiving stream on FGPA1 and FGPA2: recovery of clock 40 successful
  - Tested both internal and redistributed 320 MHz Ref clock & LHC (40 MHz) clock
  - FGPA1 -> FGPA2 Relay under development

## Apollo CM Dual A2577: External Clock Sources

ATCA Clock and TCDS Clock/Data



Rev 39, 12 Feb 2023

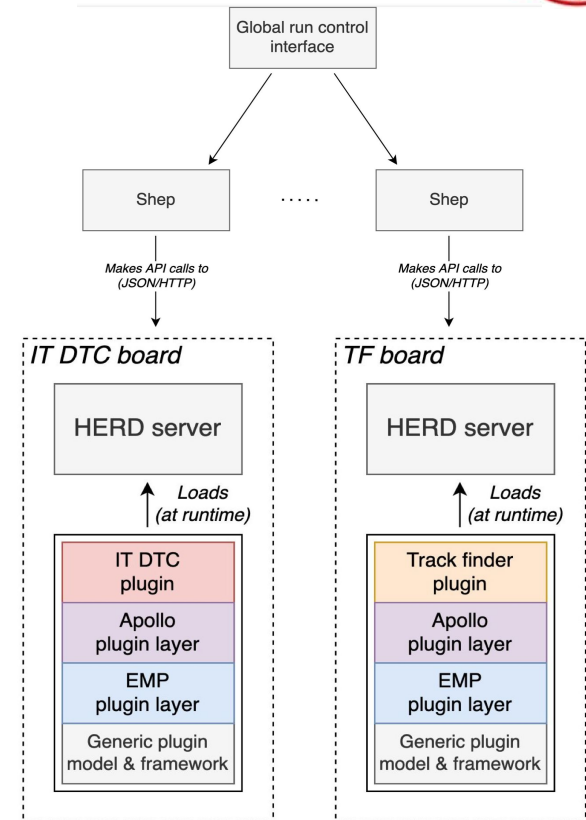
- [\[1\] The Phase-2 Upgrade of the CMS DAQ & HLT TDR](#)
- [\[2\] the Phase-2 Upgrade of the CMS Level-1 Trigger](#)



# Apollo CMS EMP + Shep/Herd Plugin



- **EMP + Shep/Herd is the CMS Online Software Framework**
- **Explicit Commands Implemented:**
  - Powering On/Off FPGAs
  - Programming FPGAs
  - Algo/Link tests
  - Direct Apollo address tables register reads through BUTool
- **Abstracted Objects:**
  - Service Module - For monitoring status tables from BUTool
  - Command Module FPGAs - For programming/link tests
  - Microcontroller - For reading registers via BUTool
  - Copper Links + Fireflies - Link test endpoints and monitoring
- **Monitoring:**
  - BUTool status tables
  - FF temperatures + optical power
- FPGA FW builds, time built, clock frequencies, temperatures



# Apollo CMS EMP + Shep/Herd Software Link Tests



- Testing CSP links using EMP + Shep/Herd Software Framework
- [Spreadsheet detailing all statuses](#)
- Copper links in single FPGA loopback and with cross FPGA mappings tested @ BU & CERN
  - Tested both using PRBS and CSP with TCDS2 (Requiring correct timing alignment)
  - Tested in EMP Butler CLI and Apollo Herd
- 4-Channel XCVR FFs in single FPGA loopback and with cross FPGA mappings tested @ CERN
  - Tested both PRBS and CSP with TCDS2
  - Tested in EMP Butler and Apollo Herd
- 12-Channel 25G FFs in single FPGA loopback and reverse loopback tested @ CERN
  - Tested both PRBS and CSP with TCDS2
  - Tested in EMP Butler and Apollo Herd

# Key Test Results from Rev3 SM

Service Module was received on Friday Jul 19:

SM only:

- Initial tests conducted on 3 units
  - Resistance and power measurements
  - Ethernet switch and IPMC
  - Zynq SoC
  - SoC - OS
  - TCDS2 eyescan
  - Tests of the new over/under-voltage protection

SM + Rev2 CM:

- Conducted tests:
  - Programming FPGA1/2
  - Registers R/W over C2C AXI
  - R/W over IPBUS to FPGA1/2
- To be done in consort with the CM :
  - TCDS2 configuration
  - Optical and copper link tests



SM Rev3 on a test stand @ BU




# Apollo SM Rev3: Test Summary

- Service power check
- Plug in IPMC module
- Check payload power by enabling it with jumper
- Program CPLD
- Power up, check that IPMC has enabled payload power
- Plug in SoM module
- Power up, check all power voltages enabled by SoM
- Plug in pre-programmed SD card and SSD, boot
- Check Ethernet connection
- Test SM sensor readout
- Test overvoltage/undervoltage protection circuitry
- Plug in Command module
- Tests with Command module
  - Programming FPGAs via JTAG
  - Test Chip2chip links
  - Test connection to MCU
  - Test CM sensor readout
- Plug into ATCA crate
  - Test Ethernet access
  - Test TCDS




# Apollo SM Rev3: TCDS Link IBERT Tests

- From SM to DTH, no errors



- From DTH to SM, no errors



- From DTC to Apollo 230 and back
- Plot quality depends on backplane slot
- Worst case shown

- No errors detected
- SM Rev 3 to/from DTH, PRBS-31, DFE





# Pre-production Schedule for the Apollo Boards

- Strategy is to align the production for all systems so the procurement and production costs can be reduced
- Pilot production ~30 boards (primarily TF expected Q1 2025)
- Production quantities in table expected Q3 2026
- First round of Apollo pre-production is needed for the Track Finder system: Same SM & Similar CM boards (speed grade -2 instead of -1)
- Plan:
  - Once the current SM Rev3 boards are fully tested, proceed with full pre-production numbers ASAP
  - Will produce all PCB boards for CM, but only populate **3 boards** (2 with FPGAs, and 1 with other parts only). The three boards will be tested with the new 12 Ch parts (~Oct/Nov '24), awaiting successful testing and final go-ahead of SamTec 12 ch 25G parts before more boards will be populated.

| System | Quantity          |
|--------|-------------------|
| TF     | 162 + spares (10) |
| IT-DTC | 36 + spares (5)   |
| BRIL   | 16 + spares (4)   |



# Summary

- HW:
  - Apollo Rev2 is being actively used for development & testing in several locations: BU, Cornell, CERN
  - SM Rev3 has been produced and is being tested at Boston U
  - CM Rev3 is close to going to manufacture
  - SM Board functionality validation is complete for Rev2 and ongoing for Rev3
- FW:
  - SM Zynq mature and well tested
  - CM TCDS2 Clocking scheme testing is well advanced and confidence in final design for Rev3 is high
  - CMS Application framework EMP is becoming mature and well tested by IT-DTC and currently work to build TF & BRIL under EMP are ongoing
  - Initial testing of IT-DTC DAQ Chain to commence at CERN ~ Nov/Dec '24
- SW:
  - CLI Tools are very mature and well tested
  - EMP Board level plugin almost complete and under test
  - Application specific EMP plugins to follow



Thank You.



# Backup



# Challenges

- Many and Various!
- Thrown in at the deep end with all previous engineers having left the project.
- Hardware issues:
  - CM: Copper traces on some boards giving noisy links and unhealthy eye diagrams
- Spicy mixture of development tools & frameworks:
  - Gitlab vs Github
  - SM: Hog framework for Zynq FW Builds
  - CM: CMS EMP Framework for application builds
  - Vivado and all it's curiosities. Currently working with a mixture of 2022.2 and 2023.2
    - Two bugs identified causing build issues
      - TCL Script behavior (2023.2)
      - Large project crashes (2022.2)





# Apollo SM Rev3

## Overvoltage/Undervoltage Protection Tests

- Undervoltage sensor tested when DC-DC is off
- Overvoltage tested by adjusting threshold to a lower value
  - By soldering an extra resistor in the divider
- Both sensors are sent to IPMC
- IPMC needs a software module to check OV/UV sensors and react appropriately
  - In progress now ...

# Apollo History



Apollo Hardware now has 3 major revisions (Rev1, Rev2 & Rev3). In the updates from Rev1 to Rev2 and Rev2 to Rev3, the following main changes were implemented:

- Rev1 (5 boards) → Rev2 (11 boards) :
  - Service Module: Replaced 7-series with US+ Zynq (Enclustra). Tested OpenIPMC, Terragreen material, Full CERN TCDS2 compatibility
  - Command Module: Halogen free material, replaced KU15P and VU7P with two VU13P FPGAs, added support for higher voltage on 25G FireFly transmitters
- Rev2 (11 boards) → Rev3 (1 board) :
  - Service Module:
    - Implemented a larger high-reliability (automotive) hold up capacitor with longer lifetime
      - Holdup capacitor allows the boards to “ride” through short power dips created by other boards in the chassis blowing fuses
    - Protection against OV/UV conditions on the 12V DC-DC converter output
  - Command Module:
    - FF configuration: 4x12-channel transmitter/receiver pairs and 2x4-channel transceiver per FPGA
      - Increases the number of FireFly channels from 52 to 56 per FPGA
      - Two additional 3.8V regulators and switches & controls for the two added 12-channel transmit sites
    - All 5 clock chips are now SI5395A (previously 4x SI5395A and 1x SI5341A)
    - 2 additional general-purpose 25 Gb/s GTY links between the two FPGAs (54 total)
    - Narrower FPGA heatsink and Wider FireFly heatsink:
      - New 25Gx12 FireFlies have a long latch that couldn't close with the old FPGA heatsink
      - FPGA heatsink 5% reduction, FireFly heatsink 28% increase
- Rev2 boards have been running reliably at TIF, BU & Cornell for almost 2 years
- First SM Rev3 board is under test in BU since July 2024
- In this presentation:
  - All feature specifications are for Rev3
  - Most reported test results are for Rev2 with preliminary results for Rev3 SM + Rev2 CM.

# Hardware: Apollo SM Rev2 → Changes Rev3



Apollo SM Rev3

## Service Module:

- Implemented a larger high-reliability (automotive) hold up capacitor with longer lifetime
- Protection against OV/UV conditions on the 12V DC-DC converter output
- More robust TCDS2 routing





# Hardware: Apollo SM/CM Rev2 → Changes Rev3

## Service Module:

- Implemented a larger high-reliability (automotive) hold up capacitor with longer lifetime
- Protection against OV/UV conditions on the 12V DC-DC converter output
- More robust TCDS2 routing

## Command Module:

- FF configuration: 4x12-channel transmitter/receiver pairs and 2x4-channel transceiver per FPGA
  - Increases the number of FireFly channels from 52 to 56 per FPGA
  - Two additional 3.8V regulators and switches & controls for the two added 12-channel transmit sites
- All 5 clock chips are now SI5395A (previously 4x SI5395A and 1x SI5341A)
- 2 additional general-purpose 25 Gb/s GTY links between the two FPGAs (54 total)
- Narrower FPGA heatsink and Wider FireFly heatsink:
  - New 25Gx12 FireFlies have a long latch that couldn't close with the old FPGA heatsink
  - FPGA heatsink 5% reduction, FireFly heatsink 28% increase



# Hardware: Apollo SM/CM Rev2 → Changes Rev3

## Service Module:

- Implemented a larger high-reliability (automotive) hold up capacitor with longer lifetime
- Protection against OV/UV conditions on the 12V DC-DC converter output
- More robust TCDS2 routing

## Command Module:

- FF configuration: 4x12-channel transmitter/receiver pairs and 2x4-channel transceiver per FPGA
  - Increases the number of FireFly channels from 52 to 56 per FPGA
  - Two additional 3.8V regulators and switches & controls for the two added 12-channel transmit sites
- All 5 clock chips are now SI5395A (previously 4x SI5395A and 1x SI5341A)
- 2 additional general-purpose 25 Gb/s GTY links between the two FPGAs (54 total)
- Narrower FPGA heatsink and Wider FireFly heatsink:
  - New 25Gx12 FireFlies have a long latch that couldn't close with the old FPGA heatsink
  - FPGA heatsink 5% reduction, FireFly heatsink 28% increase

- TCDS2 changes (See next slide for block diagram) → In order to be as flexible as possible and able to support all of CMS TCDS2 stream endpoint implementations several small changes were required:
  - a. Moved the forwarded TCDS signal connected to FPGA#2 from channel RX3/TX3 to channel RX0/TX0 in the TCDS quad.
  - b. Connected a 320 MHz clock from synth R1B to the REF\_1 input of the TCDS quad on FPGA#2.
  - c. Added a 1:2 clock buffer to fanout the 40 MHz clock from synth R1B to feed the logic in the two FPGAs.
  - d. Added a 1:6 clock buffer to fanout the 40 MHz clock from the backplane to various synths and the the FPGA logic.



# Thermal Test Details

- Firmware implementation
  - Heater unit developed by authors of <https://doi.org/10.1016/j.micropro.2013.12.001>
  - Small LUT oscillators that can be enabled dynamically
  - All available MGTs enabled at 25 Gb/s
- To avoid large thermal gradient inside FPGA
  - Spread out the heater units
  - Always enable the same amount of LUTs in each cluster
    - 1 heater unit = 8k LUTs/cluster
- Temperature measurement:
  - Enabled sysmon in each SLR (4 in total)
  - Additional temperature readout from an embedded diode (DXN/DXP)
- Power measurement:
  - 12 V total current to CM measured by onboard resistor
  - Current read out from each DC-to-DC converter

