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

Imperial College London

Giacomo Fedi

on behalf of the CMS collaboration and Serenity Consortium

TWEPP 2022 - Bergen - Norway 19-23.09.2022

22 September 22

Overall Picture (2028)

Trigger Data – 40MHz
L1A Data – 750kHz

Front-End
5G-10G

Trigger Data
Back-End
16G-25G

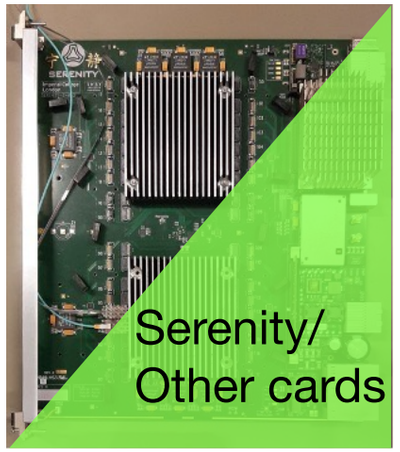

L1A Data
DAQ
25G

DAQ
L1A Data

600 ATCA Cards,
each with typically
200 optical links

L1A Distribution via TCDS

DAQ - DTH Card

DAQ

DAQ
HLT – High Level
Trigger

Planned adoption of Serenity board in CMS Phase-2 Upgrade:
Tracker, HGCAL, MTD (timing), RPC (muon), L1 Trigger, and possibly BRIL

Serenity 1.2 (latest prototype)

FPGA Daughter Card
VU13P, VU9P, VU7P, KU15P

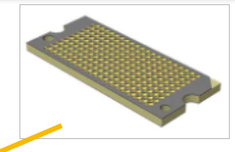

Samtec interposer

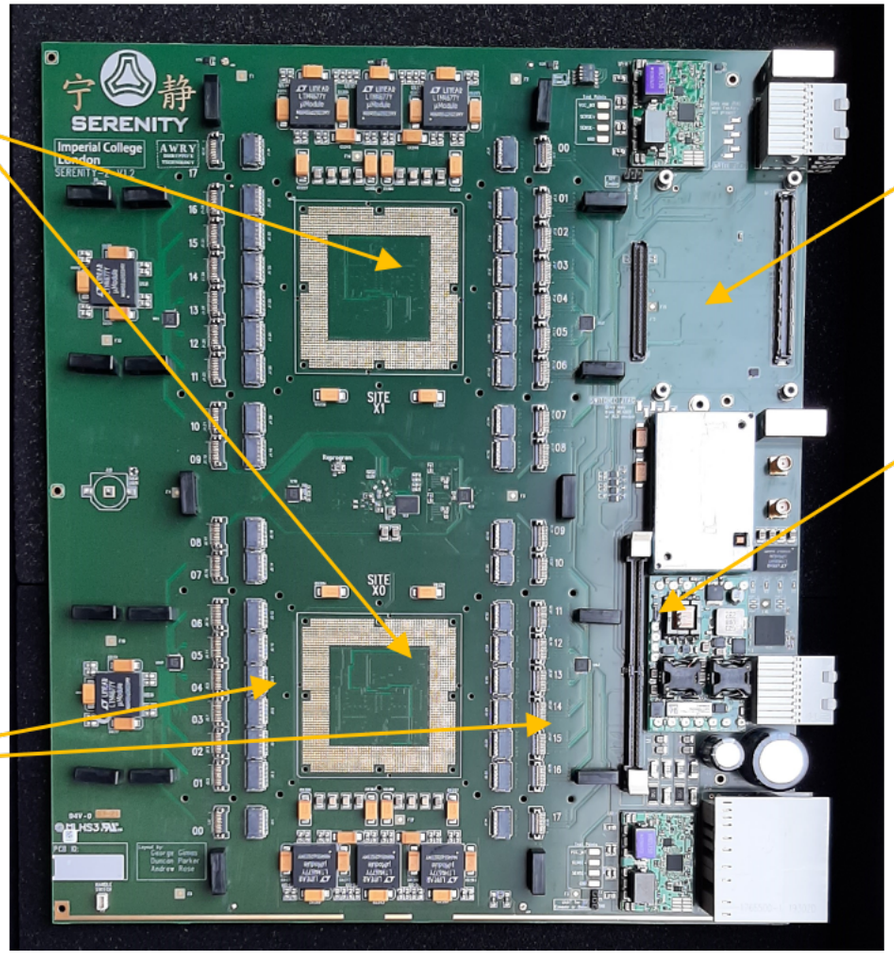

COM Express Mini Module Type10
Intel® 11th Gen Core Processors
8 or 16GB LPDDR4
NVMe SSD, PCIe Gen-4
USB 4

Optics
Back-End & Front-End

Open IPMC Module
New development
based on modern
micro-controller

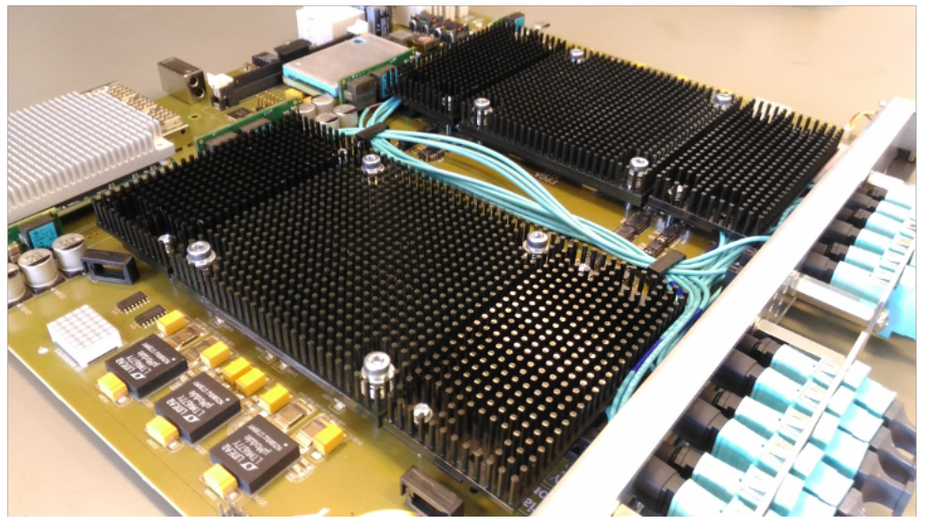

Serenity board [TWEPP-18]

- Two main prototypes developed: v1.1 and v1.2
- 12 boards produced, in used for subsystem development and for optoelectronic testing
- HW modification since LTMs (DC/DC converters) require very low ESR capacitors (or compensation) → a number of capacitors were changed

Firmware/Software frameworks

EMP Framework - Application management

- Configurable implementation of infrastructure FW & SW that can be re-used for a wide range of applications
- Clean **separation of algorithm** from well-tested infrastructure
- Infrastructure already present:
  - MGT links
  - SLINK interface
  - TCDS2 interface
  - Rx/Tx Buffers
  - IPBus environment (Eth or PCIe)
  - Early DAQ interface

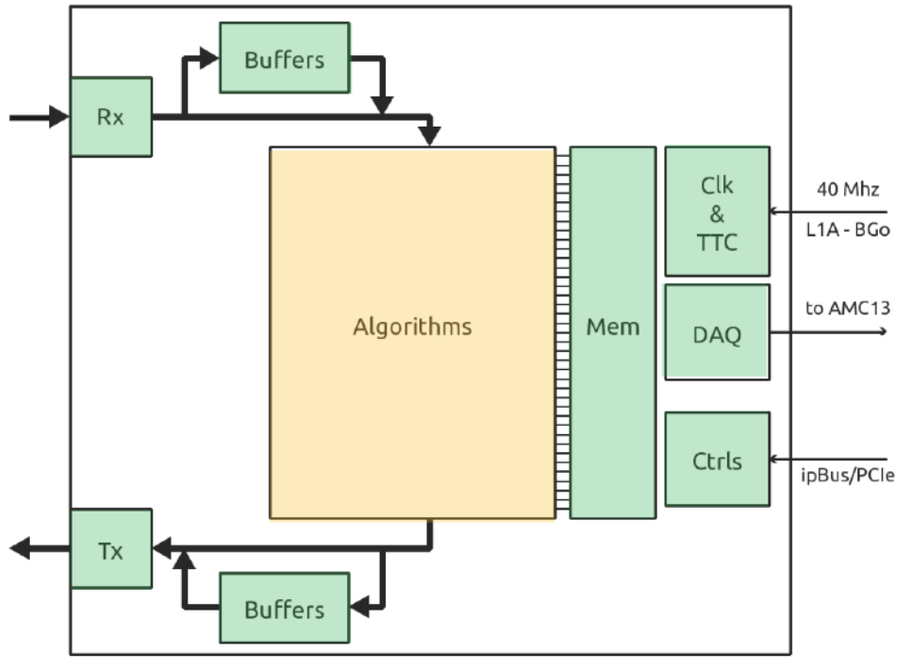

Smash - Board management

- Software/Firmware tool that permits an easy and flexible configuration of the board
- Based on pluggable modules
  - Signals can be re-routed on the fly (I2C, IPBus, GPIO, JTAG, etc.)
- Board physical config described in SMASH script
- Powerful tool
  - E.g. VU7P firmware FPGA loading time: ~8s

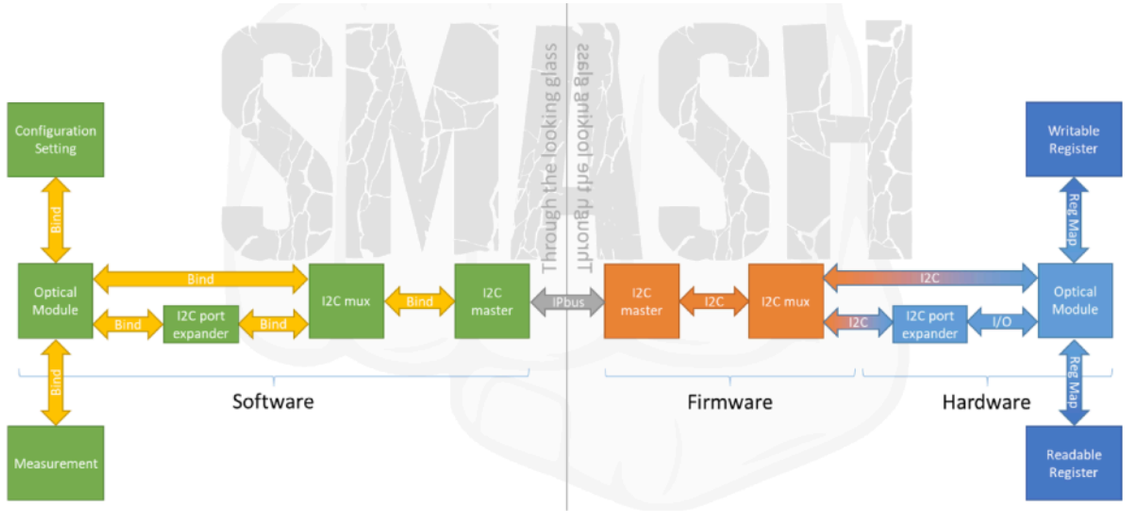

[A.Rose]

OpenIPMC

- **OpenIPMC** is an architecture-independent, free and **open-source** Intelligent Platform Management Controller (IPMC) software for ATCA boards [project]
- IPMC mezzanine [TWEPP-21] communicates with the shelf manager and deals with the very low level controls of the board (e.g. power on/off, temperature monitoring, etc.)
- Based on widely-supported FreeRTOS operating system and Linux-based open toolchain
- In Serenity, OpenIPMC is used in a LAPP-pin compatible DIMM called **OpenIPMC-HW** mezzanine
- OpenIPMC-HW mezzanine is an **open-source hardware**
  - Based on STMicroelectronics STM32H745 microcontroller
  - Cheap < $100

- **Serenity implementation**
  - FPGA temp protection
  - Temperature sensor monitoring
- **Available features**
  - PICMG-compliant IPMI functions
  - Telnet/Serial console
  - DHCP
  - Project modularisation (board customisation)
  - Remote FW upload/update (HPM.1)
  - Xilinx XVC protocol (remote FW uploading)
- Currently testing on a multi-board environment
- **Under development**
  - FLASH file system for configuration data

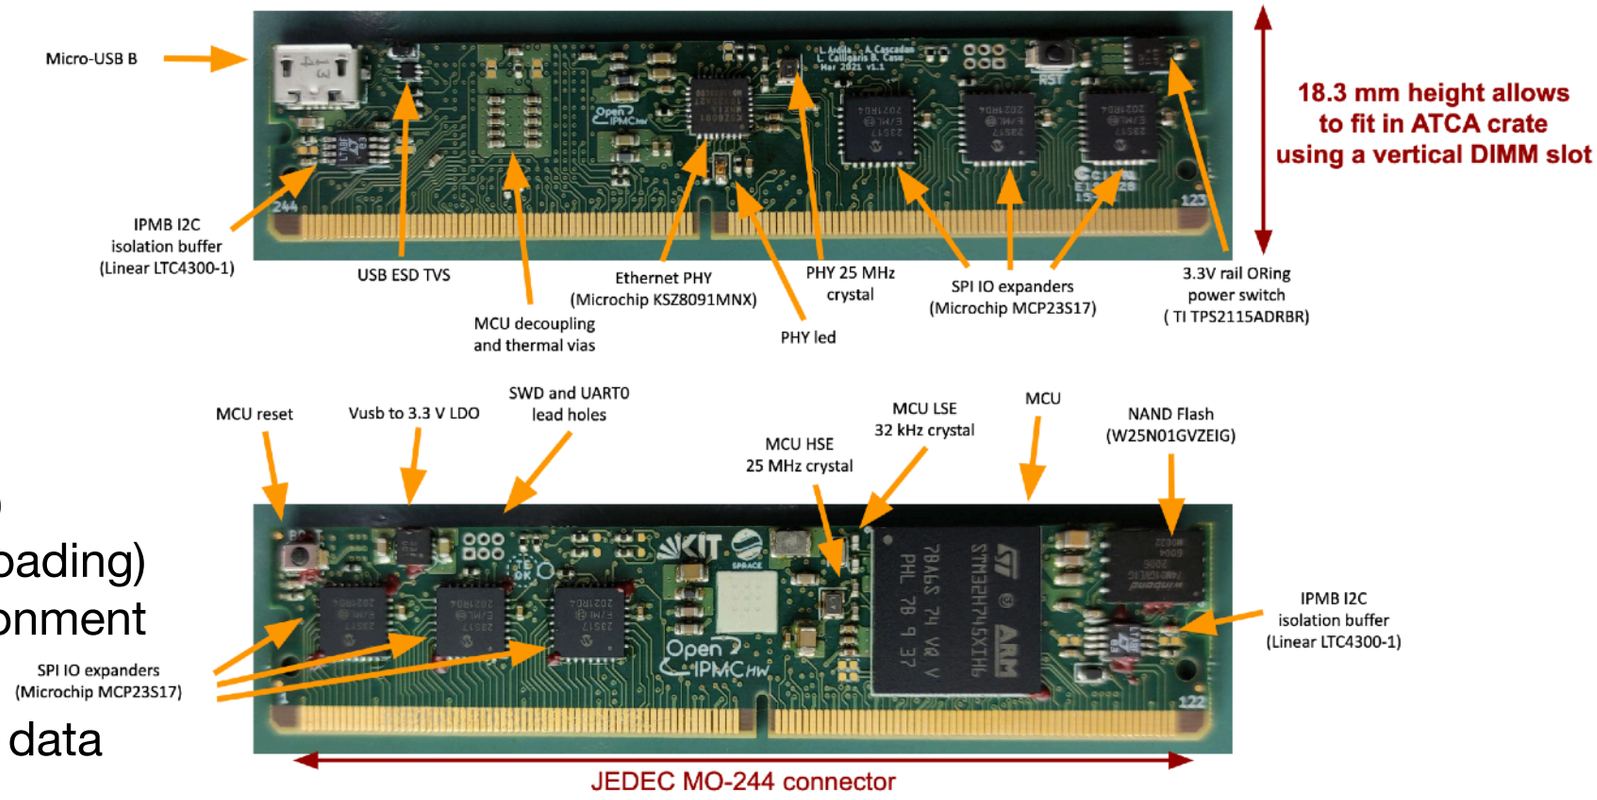

Building 186 Tracker test stand

CERN Building 186

- **Tracker integration** ongoing in CERN building 186 (TIF)
- Test stand for most of Serenity system integration (so far)
  - Board-to-board link testing, DTH integration (TCDS/SLINK)
  - Thermal/Power studies under full load
  - Infrastructure firmware/software testing
  - DTC FE firmware testing -> full project integration

Test stand

- ATCA board arrangement changes depending on test requirements
- **Available boards**
  - 2 Apollo ATCA boards
  - 2 Serenity v1.1 with 2xKU15P per board
  - 1 Serenity v1.2 with a KU15P and a VU13P
  - 1 DTH p1v2
  - 1 Emerson F125 switch

Rack A16 in building 186
(Couple of months ago)

Thermal measurements

Serenity thermal measurements

Motivations

- For LHC Phase2 most likely use VU13P FPGAs
- **Power consumption can reach 200W**
  - Need **10-year lifetime** for FPGAs and optics (<100C FPGA, <50C FF Optics)
  - Possible **cooling issues** (especially with 2 FPGA per board)
- Phase 2 cavern racks will be limited to 10kW max power
  - CMS planning **~4kW for ATCA-shelf electronics**

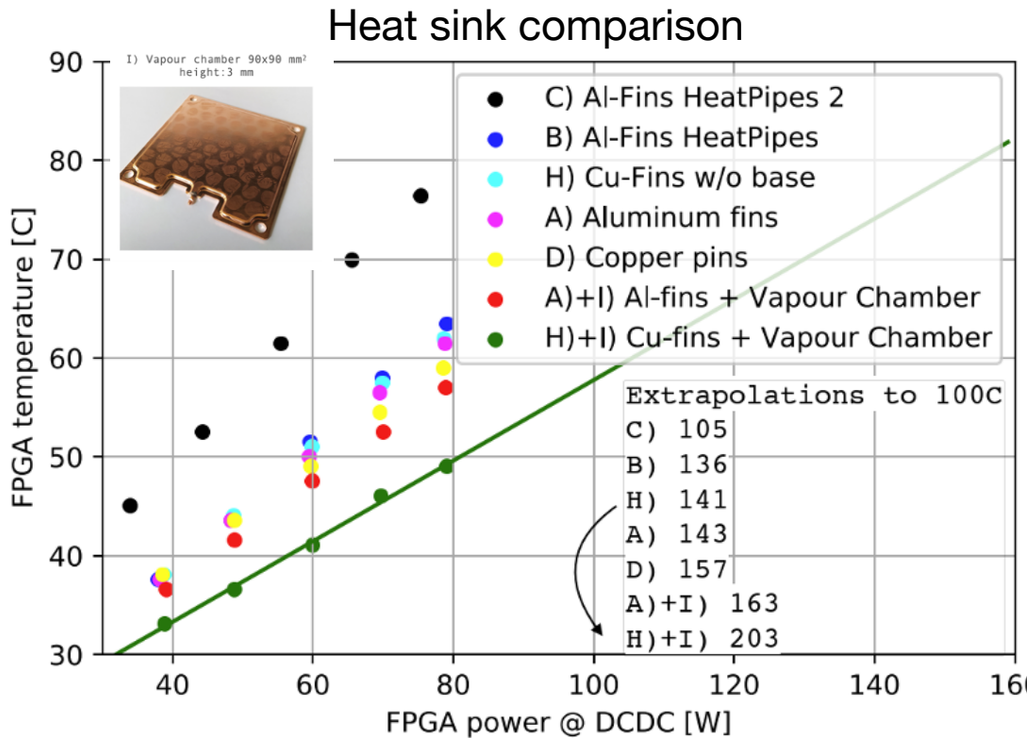

Board cooling tests

- Studied using a Serenity board (2xKU15P and a single VU13P) in a Schroff-LHC ATCA shelf
- Extrapolations shows we can deal with 200W with a single VU13P board but we should limit to 120W/FPGA in case of 2 FPGA/board
  - Promising solution under study: **vapour chamber heat sinks** seem promising
- In **CMS the rack power limit** might limit the FPGA power consumption (e.g. **<150W per FPGA** in case of a shelf housing 10 boards)
- **Learned lesson**: removing the **interposer** improves the cooling (heat sink height and PCB thermal resistance gain)

Rack thermal measurements

Rack cooling

- Measured the **overall rack cooling performance** in the racks at building 186 at CERN
- **Shelf fan speed** kept to 10/15 to optimise power consumption 500W/shelf
- Current chilling system **cannot provide enough water flow** to cool down the rack air loop properly when the shelves @ maximum power (10kW)
- **Investigations are ongoing** to improve the chilling system in our system and check if a similar issue is present in the CMS cavern

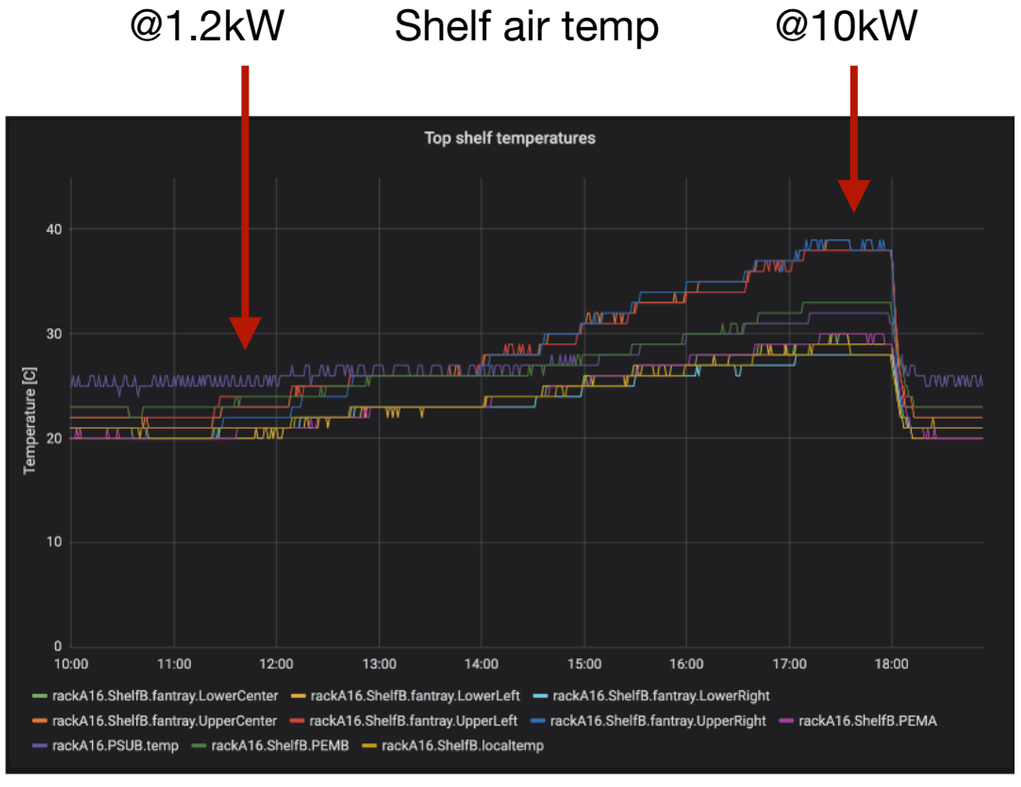

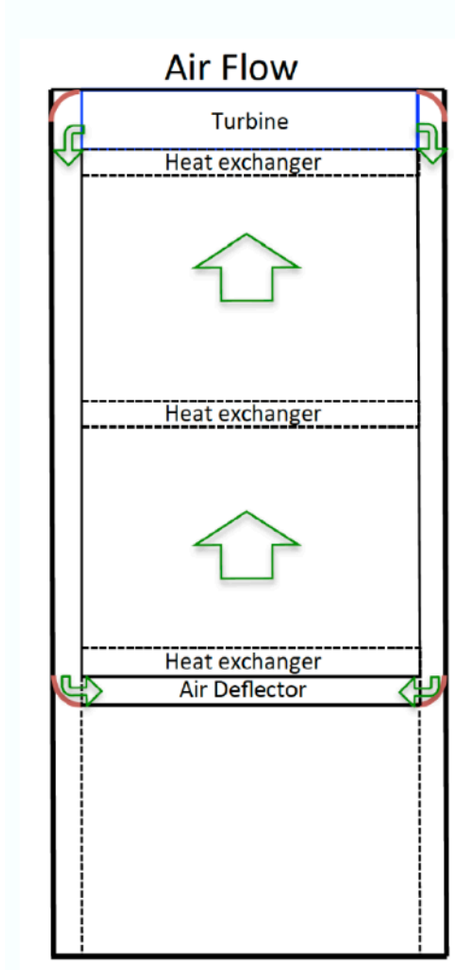

Optics testing

12 channel @25 Gb/s FireFly

Motivation

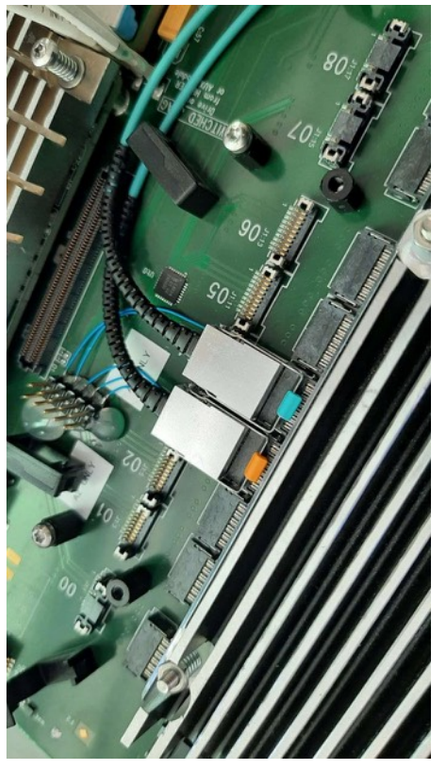

- Serenity board is designed for 4Ch and 12Ch Firefly transceivers
- Provide optical links at **low BER** compared to commercial links (i.e. 10^{-12} vs 10^{-5})
- 4Ch parts use a bi-dir socket operating up to 28 Gb/s
- **12Ch socket is uni-dir**, which is more suited to our applications, but **existing parts** only operate **up to 16 Gb/s**
- 12Ch connector also used for Versatile Link+ (LpGBT)
- Many CMS sub-detectors would **prefer a 12Ch part, but operating at 25 Gb/s**
- Samtec have developed this part over the couple of years
- We have received alpha and beta parts for evaluation

12(Rx/Tx) @25 Gb/s FireFly beta testing in Serenity

- 12 Rx-Tx pairs of 12 channel 25 Gb/s parts were received
- Results from **8 pairs presented here** with remaining 4 under test by CERN optics group
- Total of **96 channels tested**

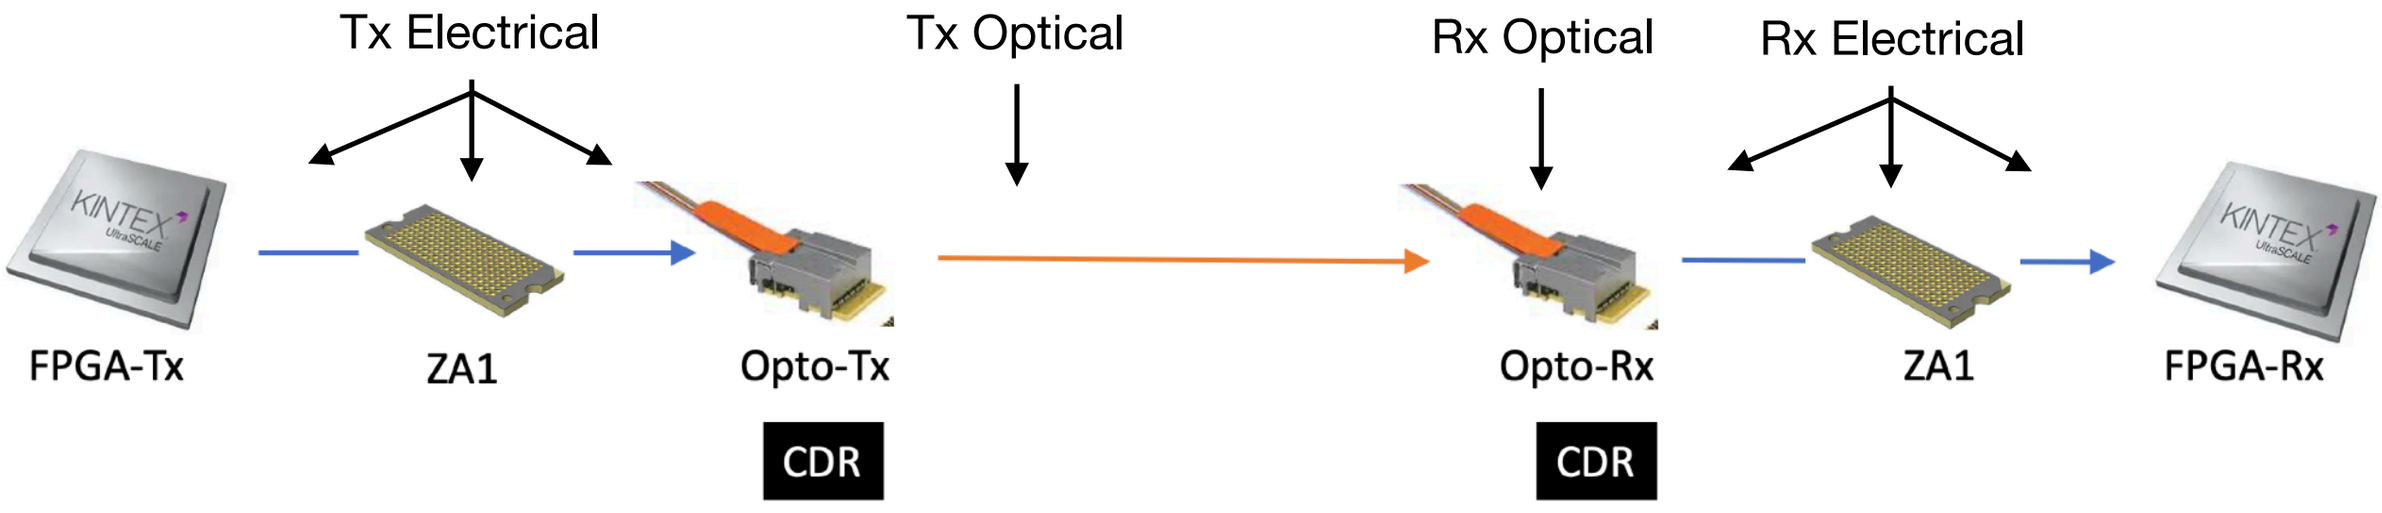

CDR = Clock Data Recovery

Tx Electrical

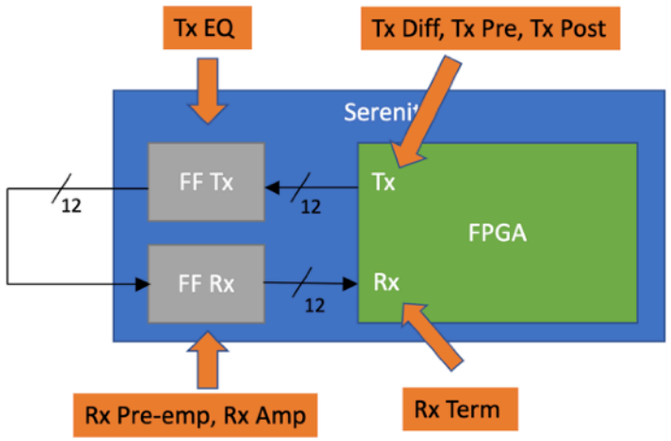

- Tests with Alpha parts suggested issues in the TX electrical domain (FPGA → FF)
  - Can not easily access the Tx electric signal
  - IBERT **parameter-scan script** to investigate the electrical performance
- Parameter scans show a **good parameter stability** (yellow area)

Parameter scan

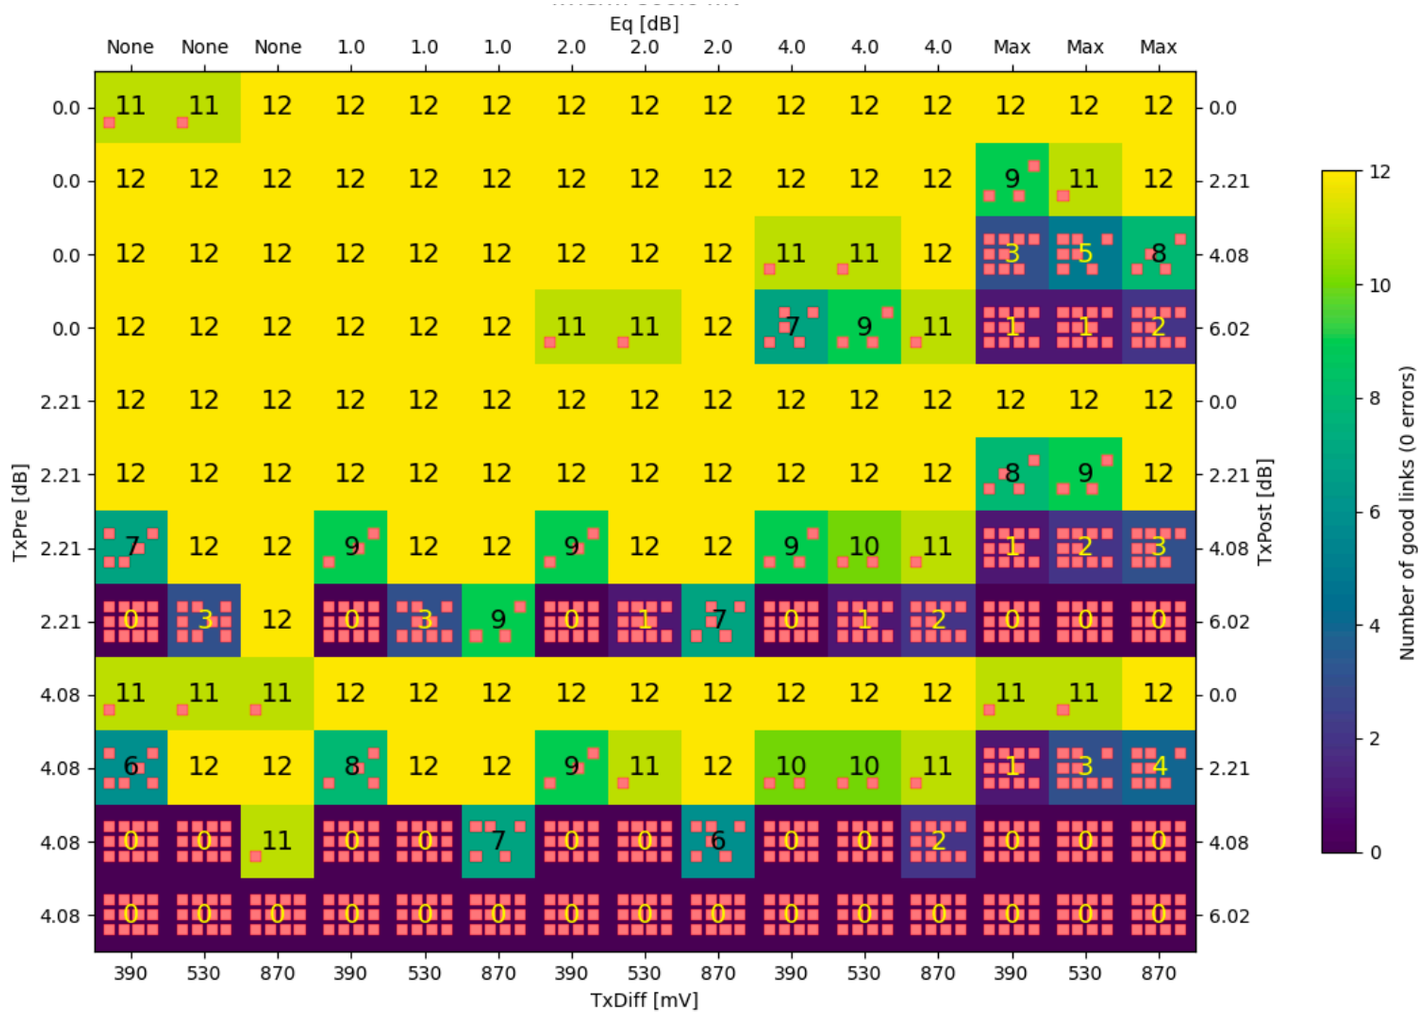

Legend

- Two parameters per axis (four in total)
- Each bin corresponds to different settings
- Number/colour of the bin represents the number of links without errors
- Each pink square in a bin represents which link is giving errors

Tx Optical

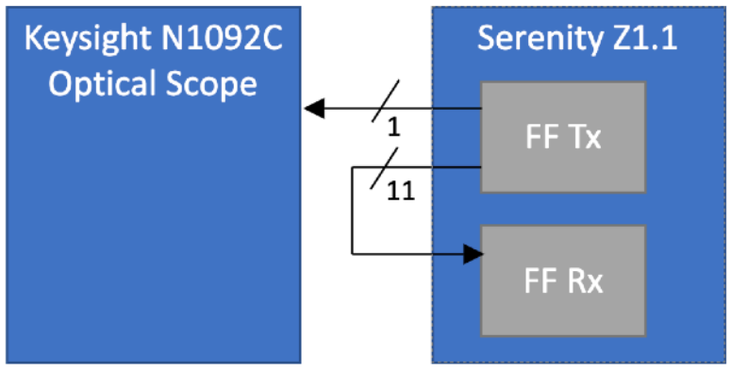

Eye diagrams of a Tx FF 12Ch @ 25 Gbps

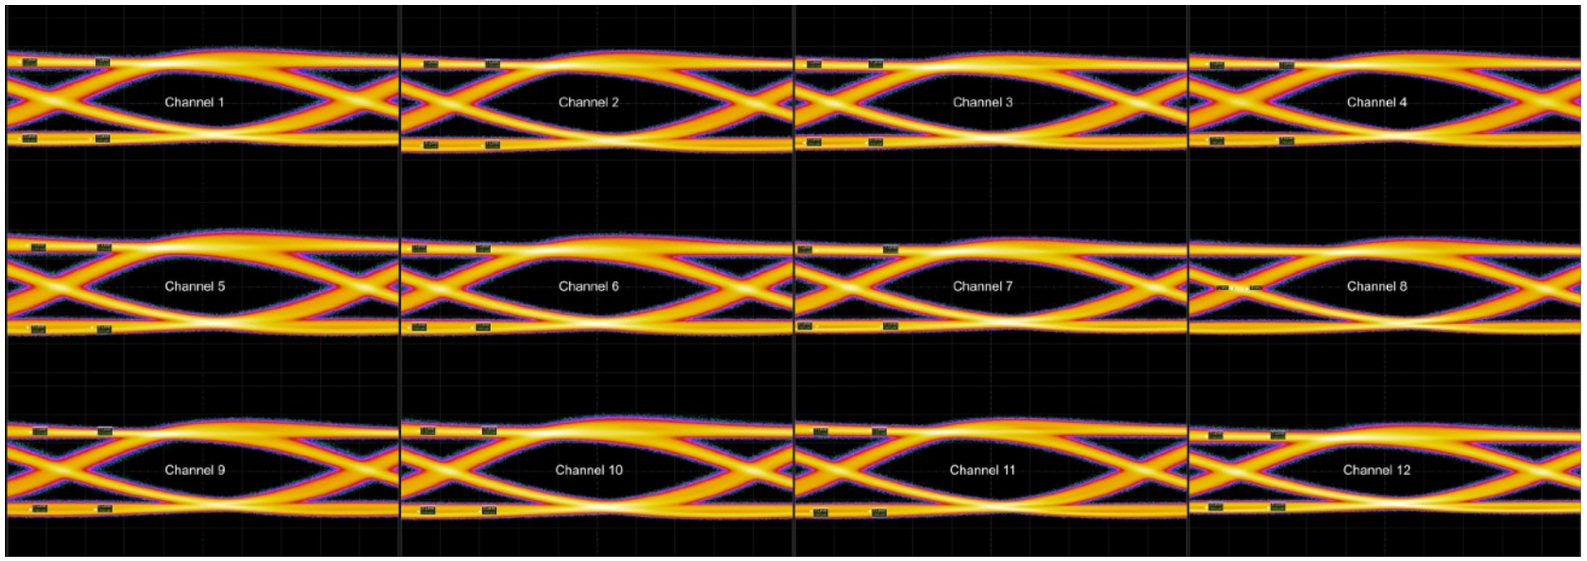

Eye diagram, OMA, and Extinction Ratio

- Optical Modulation Amplitude (OMA) and Extinction Ratio **meet specification**
- Optical eye **diagrams are clean** (both PRBS7 & PRBS31)

Issues

- One Tx channel failed (corrupted optical eye)
- Failure not seen by Samtec
- Cause remains unknown, but could be different **power sequence** or insufficient anti-static precautions
- Keenly awaiting results from Samtec reliability tests

Extinction Ratio

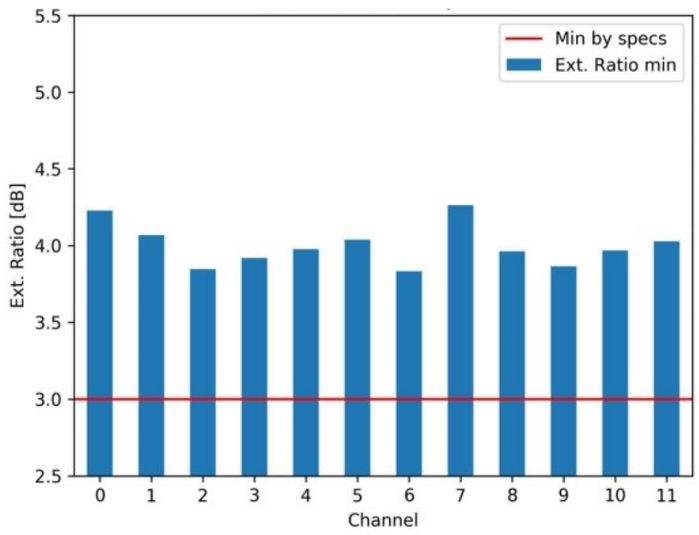

Optical Modulation Amplitude

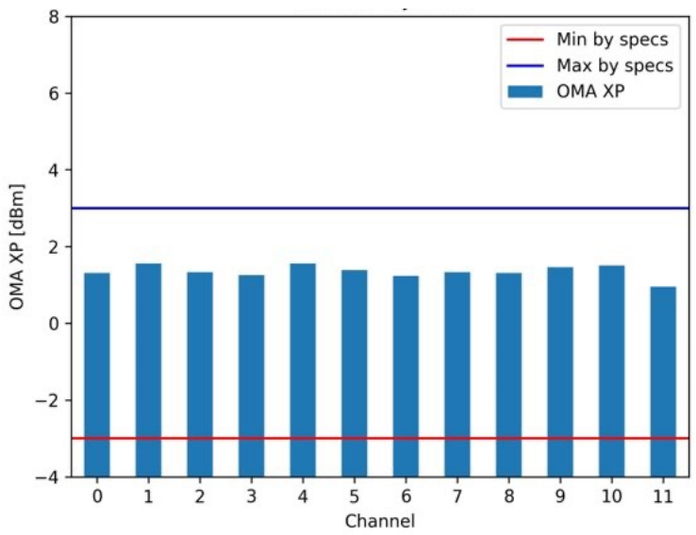

Rx Optical

- Measure **Rx sensitivity** to attenuation of optical signal using an **optical attenuator**
- All links were connected in a fibre loopback (except the link under test) to include cross-talk
- BER vs attenuation consistent within all the working channels
- **BER vs Rx OMA** (corrected for Tx OMA and attenuator offsets) is **within spec** (<-6dBm) at BER<10^{-12}

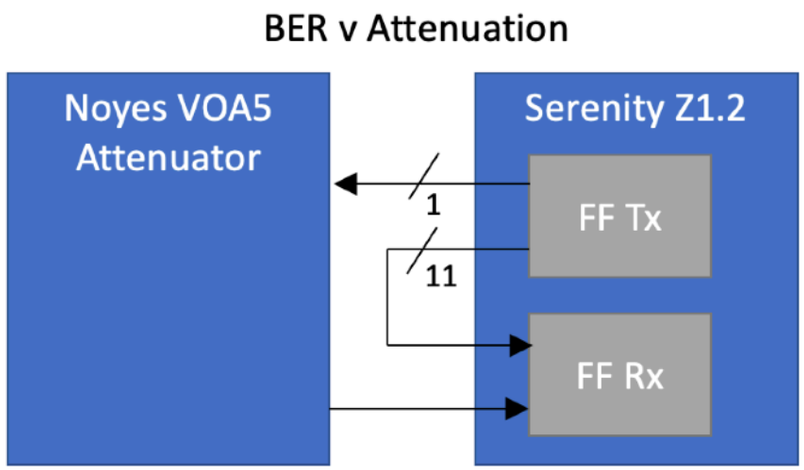

Calibration step

BER vs attenuation

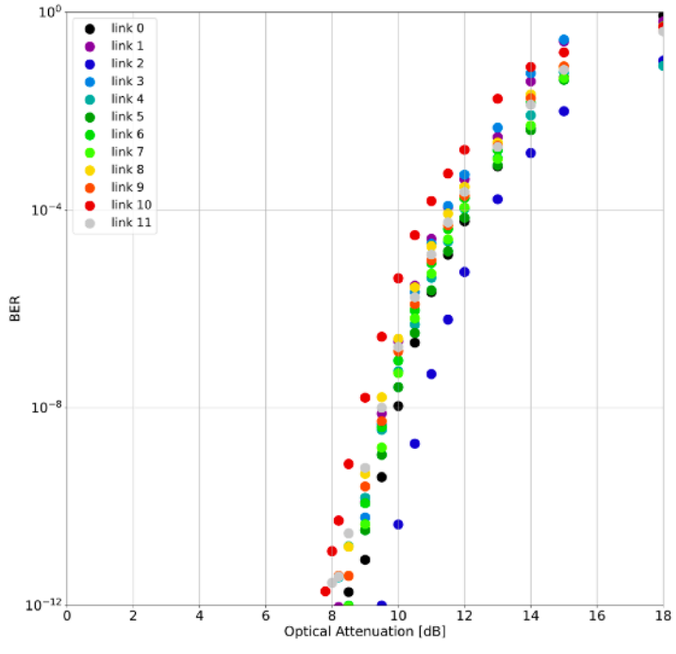

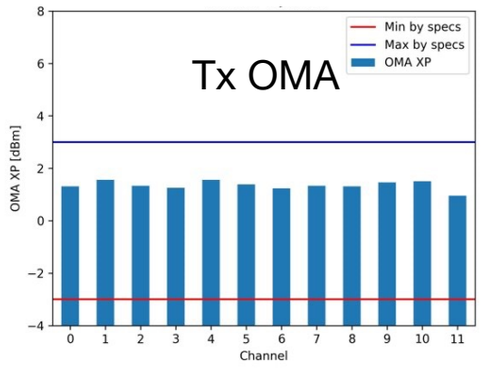

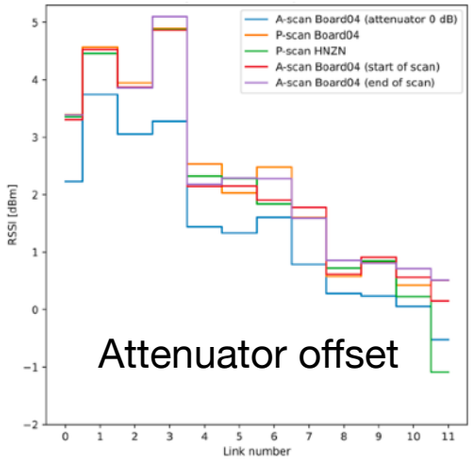

BER vs Rx OMA

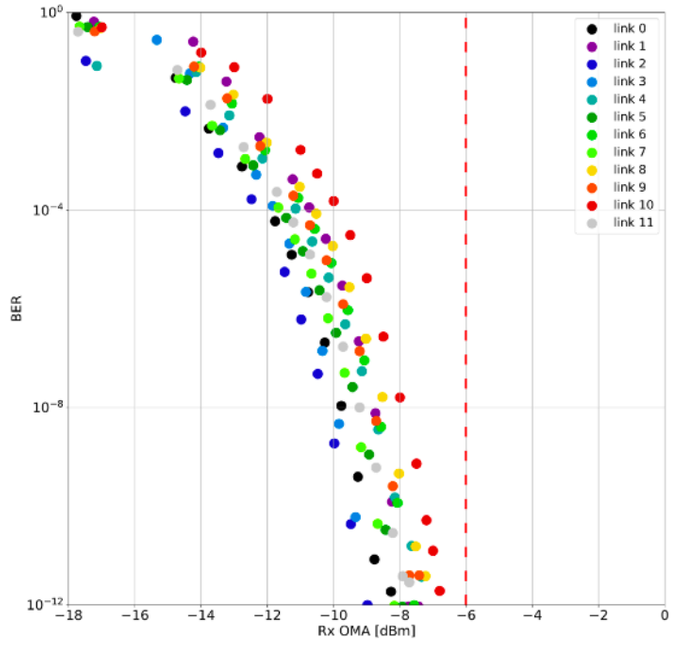

Rx Electrical

Bathtub plots

- BER vs sampling time-shift
- **BER extrapolated via function fit**
- BER extrapolated to 10^{-12} **good for all the channels** (>30% opening)

Issue

- Just one channel in a device couldn't perform well because of a CDR problem (not seen by an independent group)

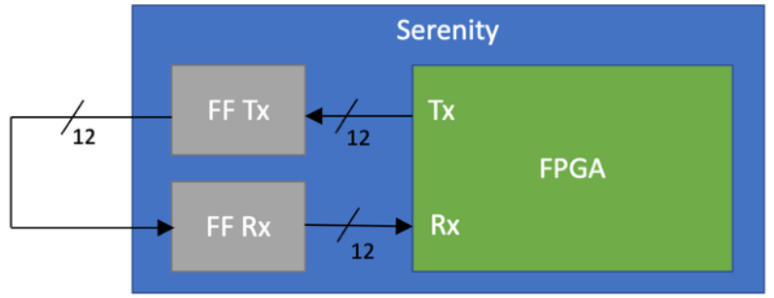

BER extrapolation

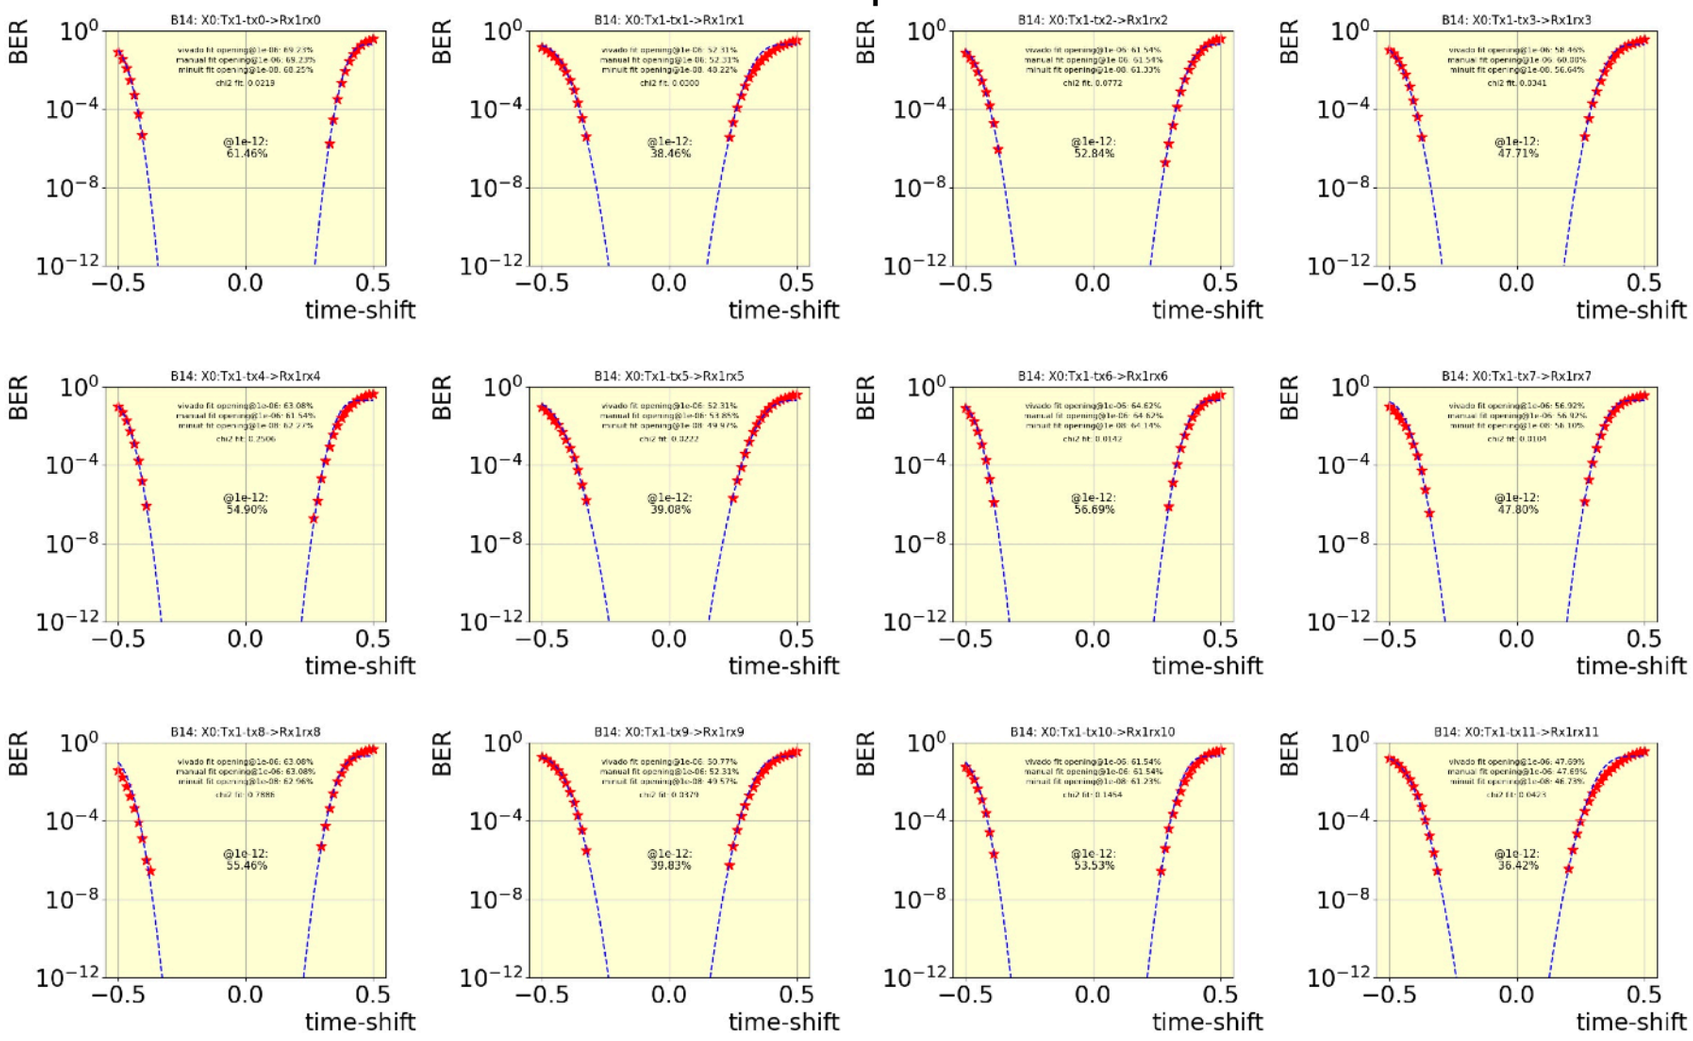

Tracker vertical slice integration

Tracker vertical slice integration

Tracker Phase 2 Goal

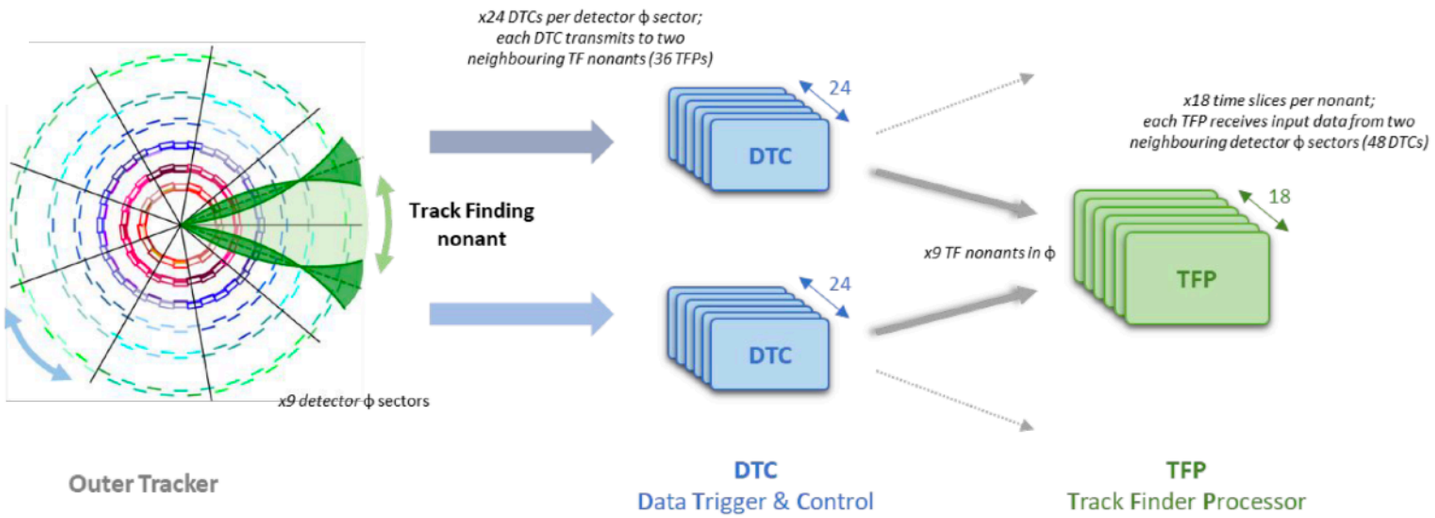

Integration made with **current prototype hardware**, thus before moving to pre-production

- Module pre-production begins in early 2023
- BE board pre-production begins in late 2023
- Aim to to assemble larger scale slice tests in 2024

Current progress

2S Front End module testing

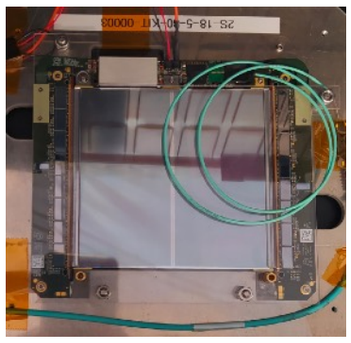

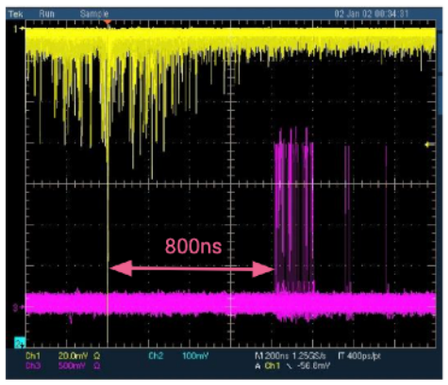

FE - BE(DTC) on test beam

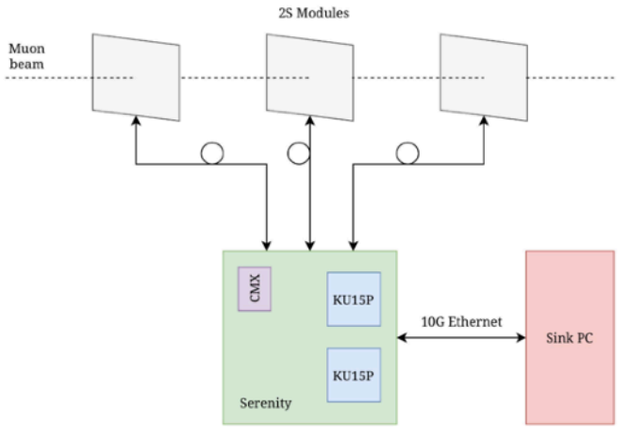

DTC - TFP(Apollo board)

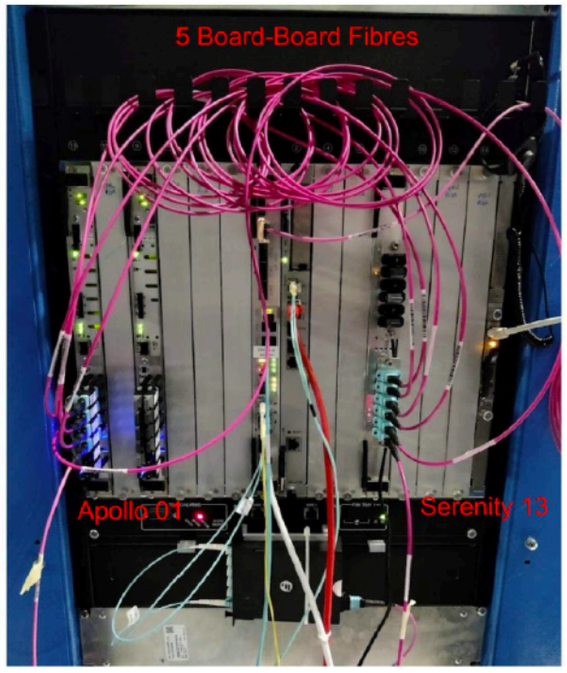

TFP(Apollo) - L1T

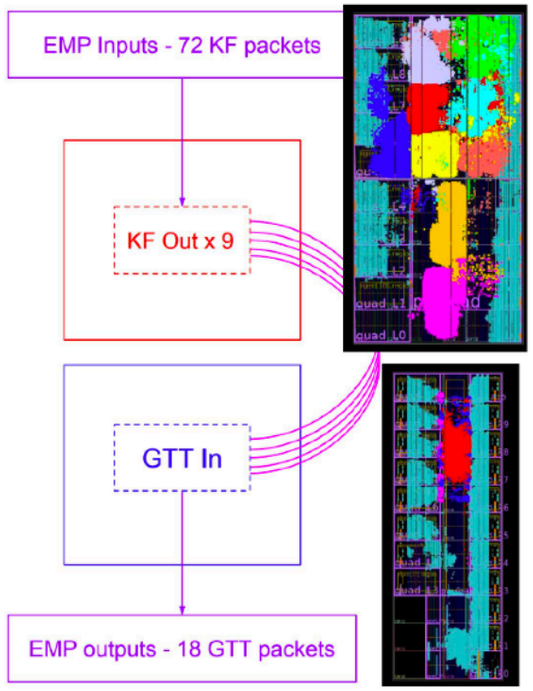

Vertical slice: test beam

Test beam @M2 muon beam @CERN

- First slice test for tracker (11/2021) made in conjunction with MuonE experiment
- **4 2S Tracker front end modules** interfaced with VTRx+
- 1 Back end serenity with interface and router FW
- 40MHz readout of stubs over lpGBT → **30 TB data sent via 10G eth to a PC**
- Stable for many hours

Early DTC firmware running on Serenity

Learning how to **synchronise FE data under realistic conditions**

- Necessary for track finding
- Understanding impact on system latency

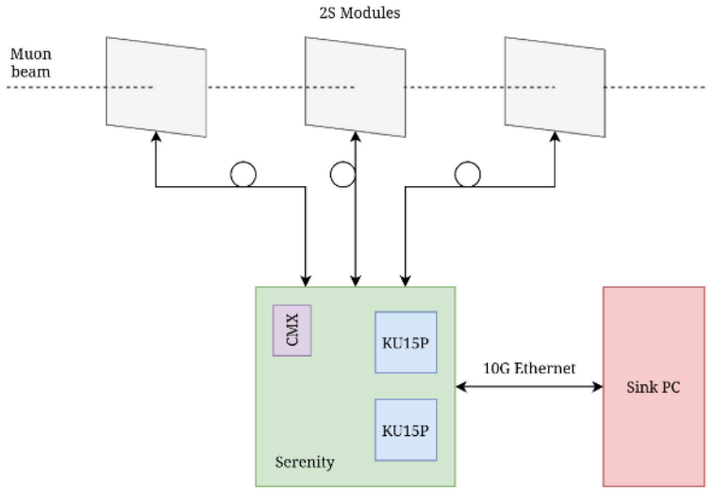

Test stand@M2 CERN

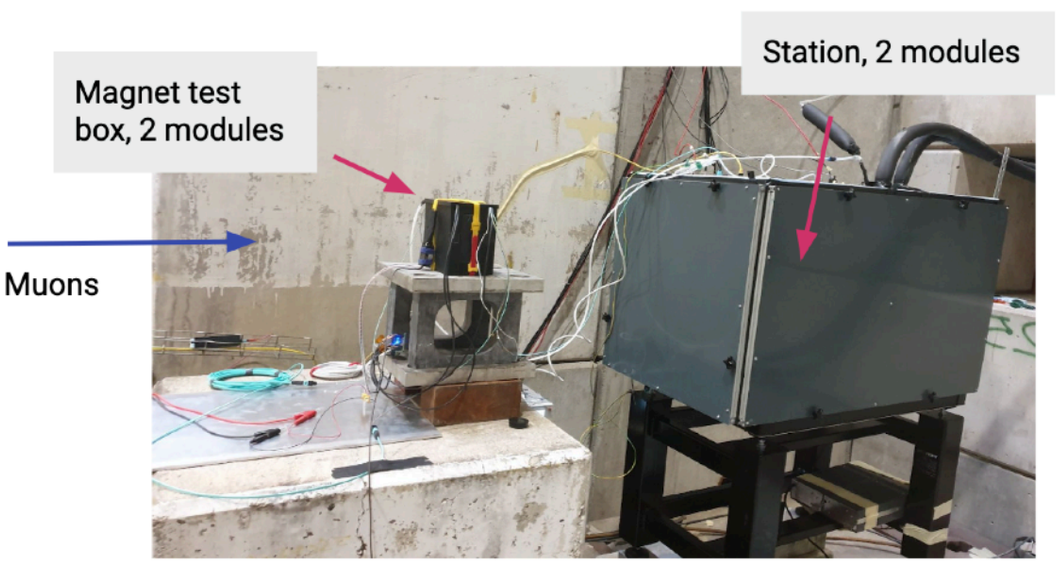

Correlation within two modules

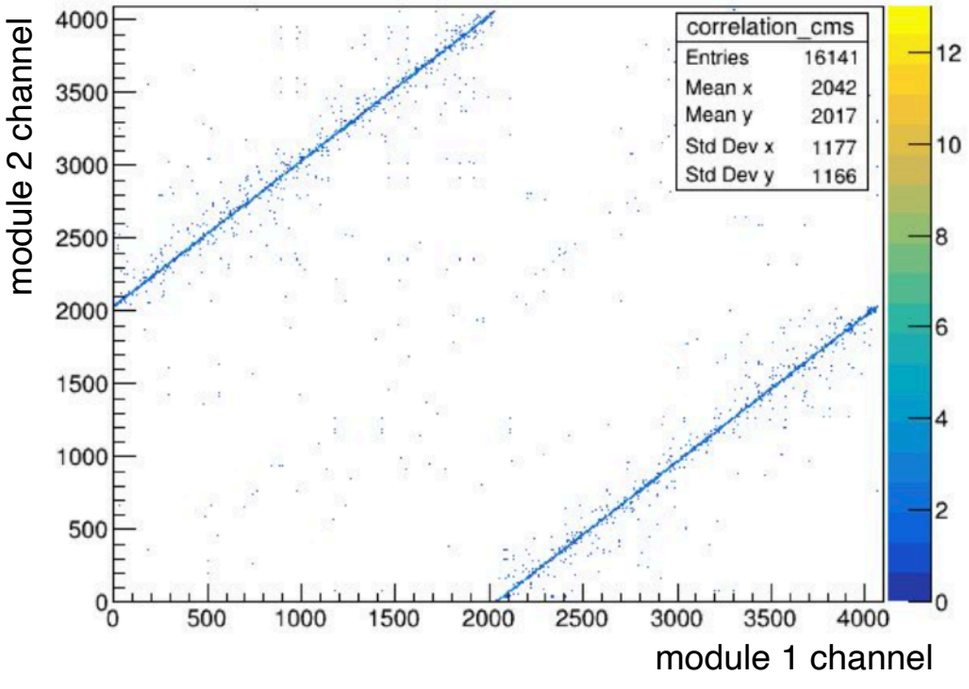

Next Serenity

Pre-production Serenity

- Design of a two-flavour board: **S (single FPGA)** and **D (double FPGA)**
  - Benefit from a modular design
  - **Zynq SoM**, targeting Kria devices
  - **No interposer** → less flexibility, but cheaper, simpler and easier to cool down
  - TCDS2 backplane signal correctly managed
  - Increased number of layers (16 → 18)
- Electrical design has been finished and reviewed
- Expected to have **first boards mid-2023** (10 boards)
  - High uncertainty due to component shortage

Serenity S

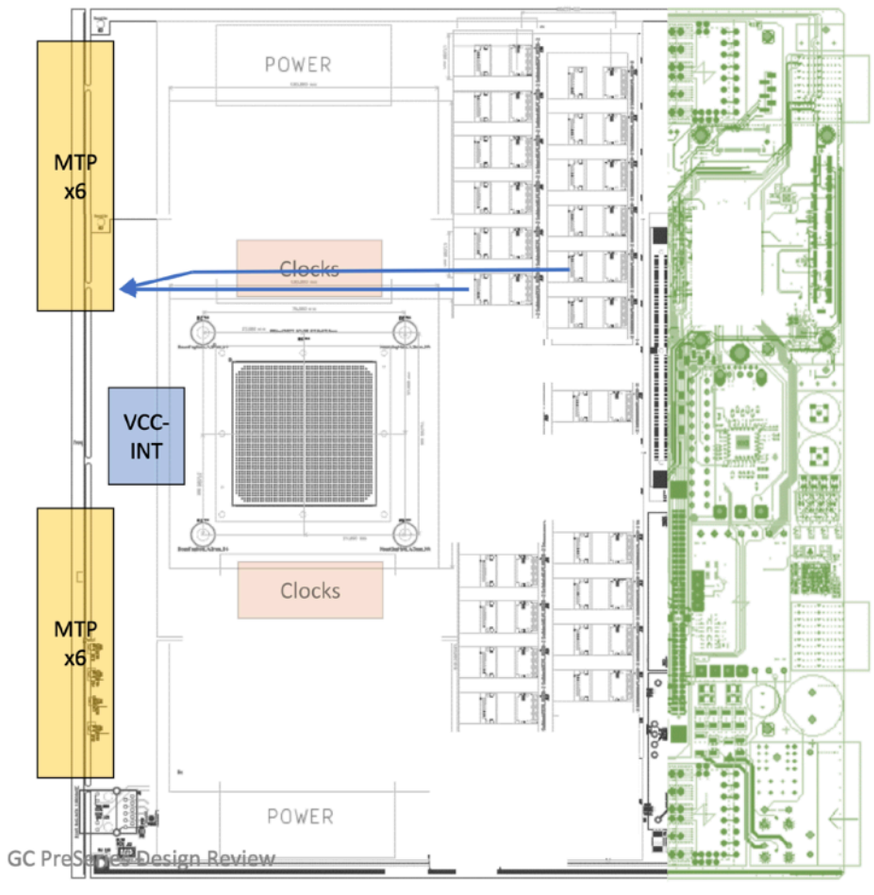

Serenity D

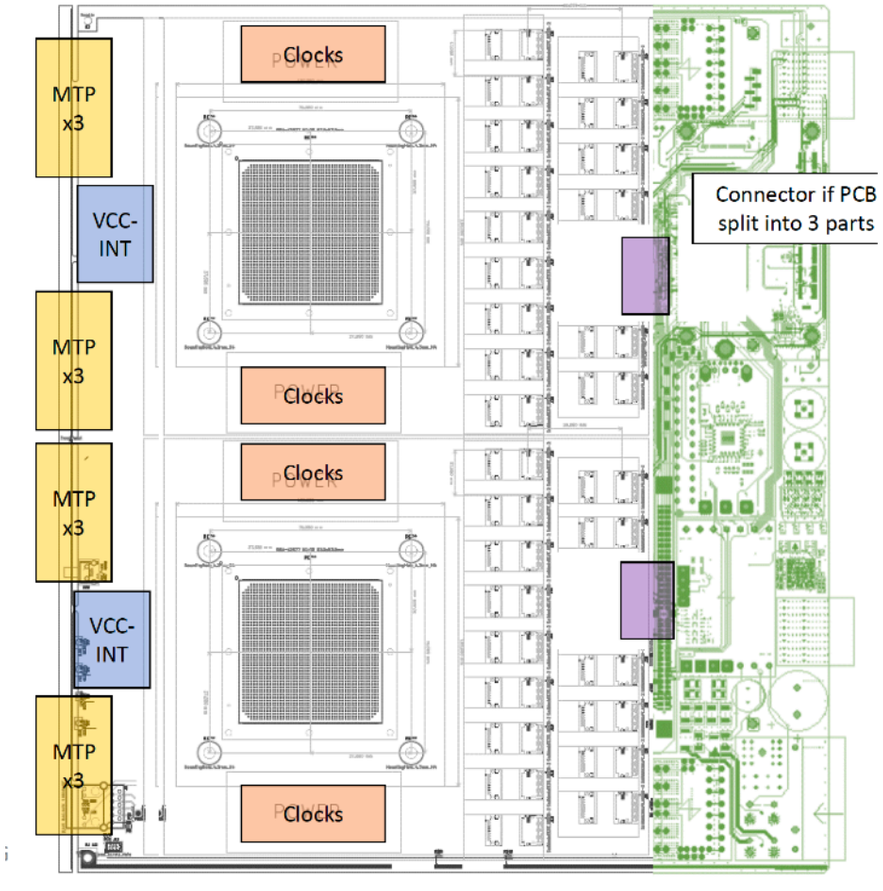

Conclusions

Optics measurements

- We have tested 8 Firefly Beta parts 12Ch 25 Gb/s
  - **Stable operation across many parameter combinations**, apart from a link with a CDR issue
  - Tx parts respect the OMA and Extinction Ratio specs
  - Rx parts respect the attenuation specs
  - 95 out of 96 links work fine (issues might be due to our testing procedure)
- Next:
  - Test of inter-optics compatibility (QSFP-DD, FF, Amphenol Leap, Finisar BOA)

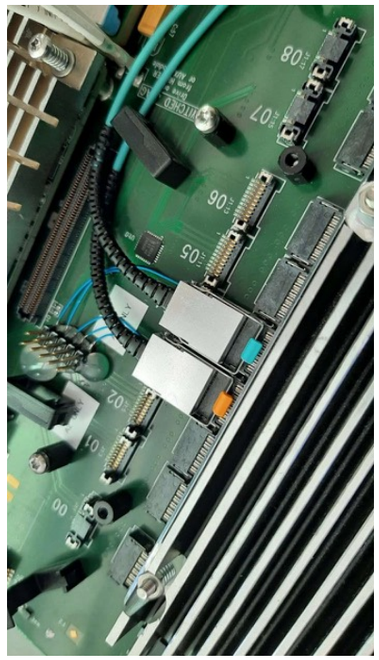

Vertical slice integration

- First beam test: all aspects of the **chain worked successfully**
- All the components of the tracker TDAQ chain have been tested
- Next beam test expected for this month aiming at testing 6 modules

Serenity ATCA board

- During the last years we gathered a **wide experience** on the board
- Some lessons we learned:
  - Interposer reduces the cooling performance and complicates the board assembly
  - COM Express is easy to use and set up, but it’s relatively expensive
  - TCDS2 signal routing is delicate and needs proper handling
  - A power sequence is needed, especially for the Samtec 12Ch 25Gb/s parts

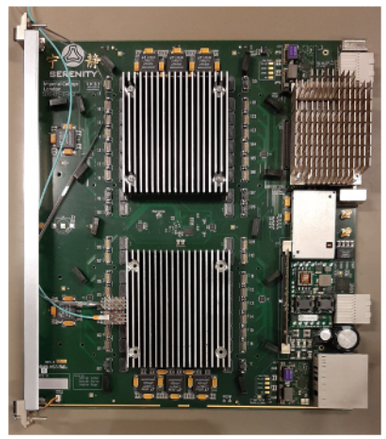

Backup

Serenity thermal measurements

Caveat: curve not directly comparable as the heat sinks have different geometries

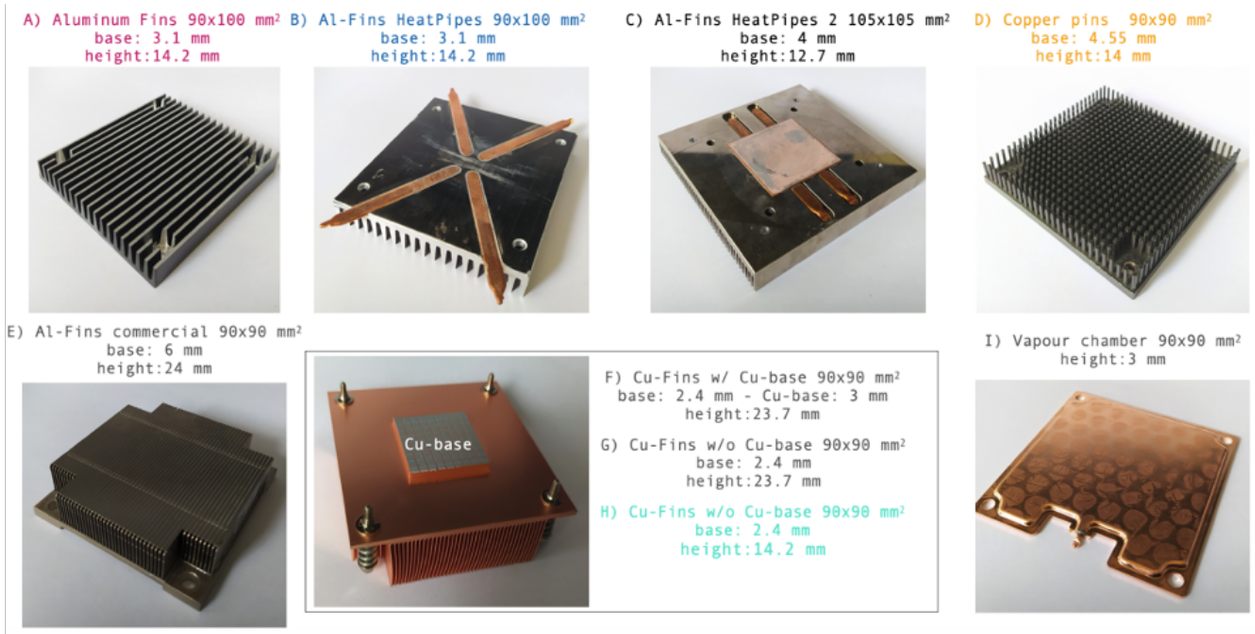

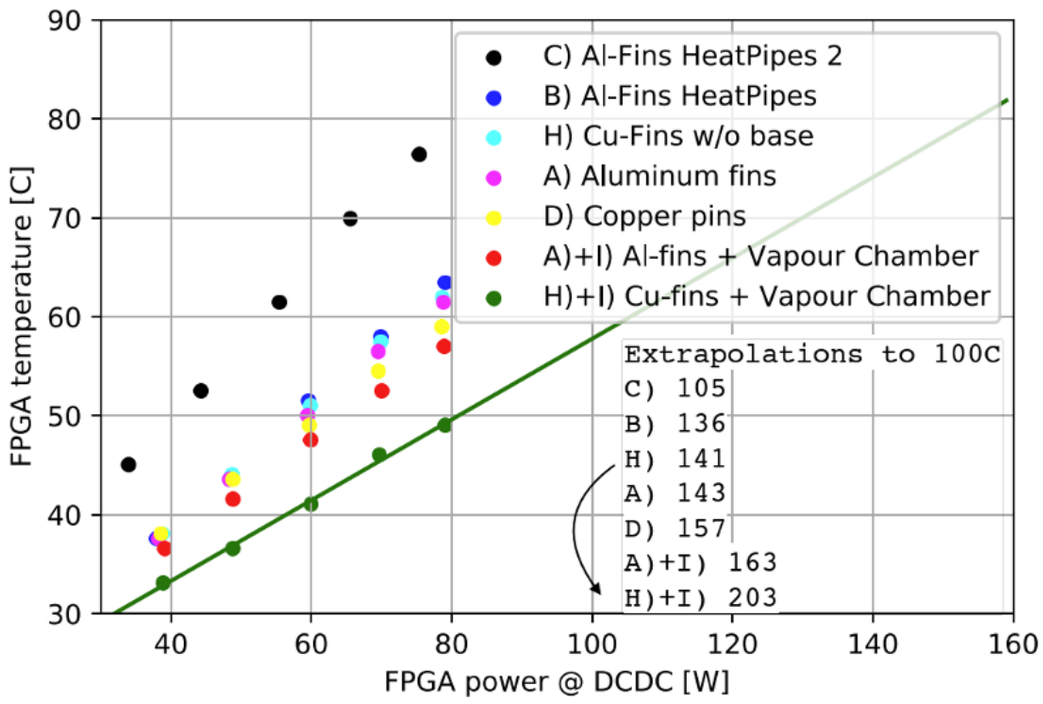

Parameter scans

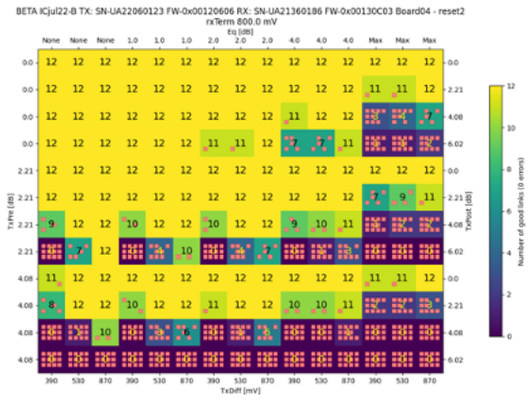

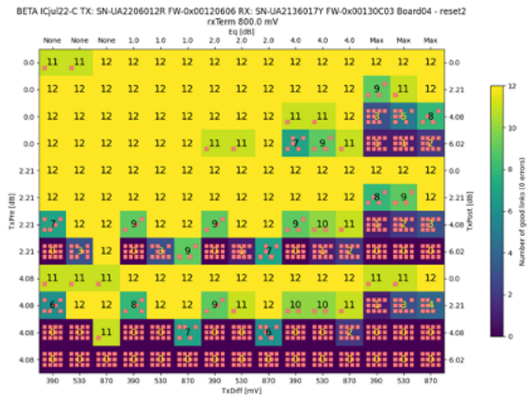

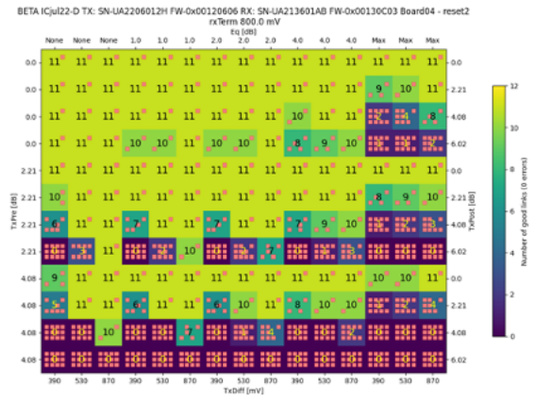

CDR Rx Issue

Broken LC-LC Adapter

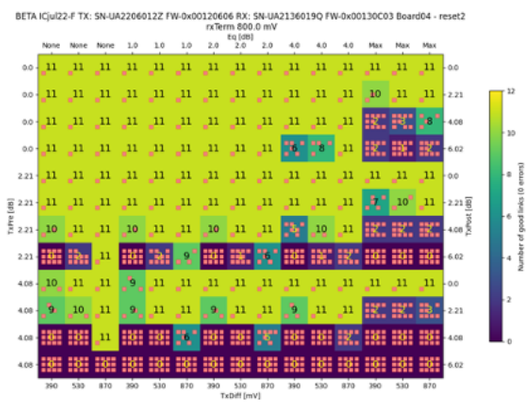

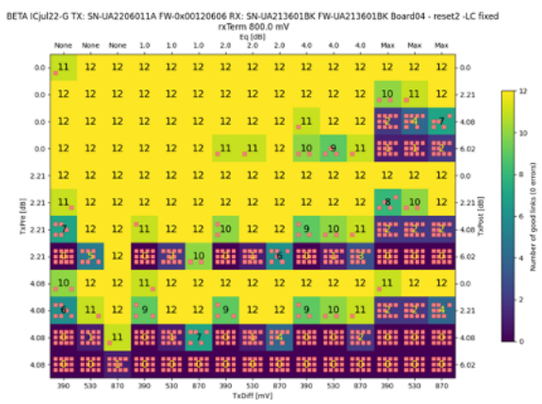

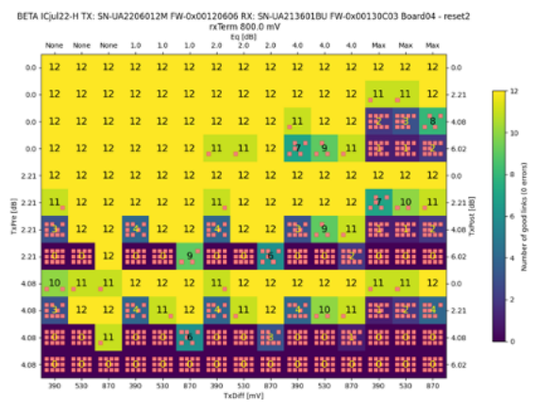

Bathtub

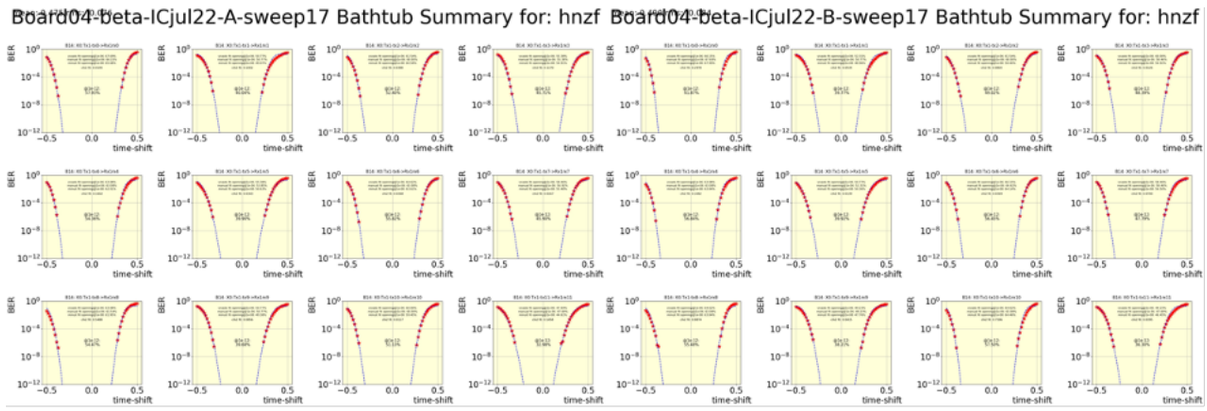

Broken LC-LC Adapter

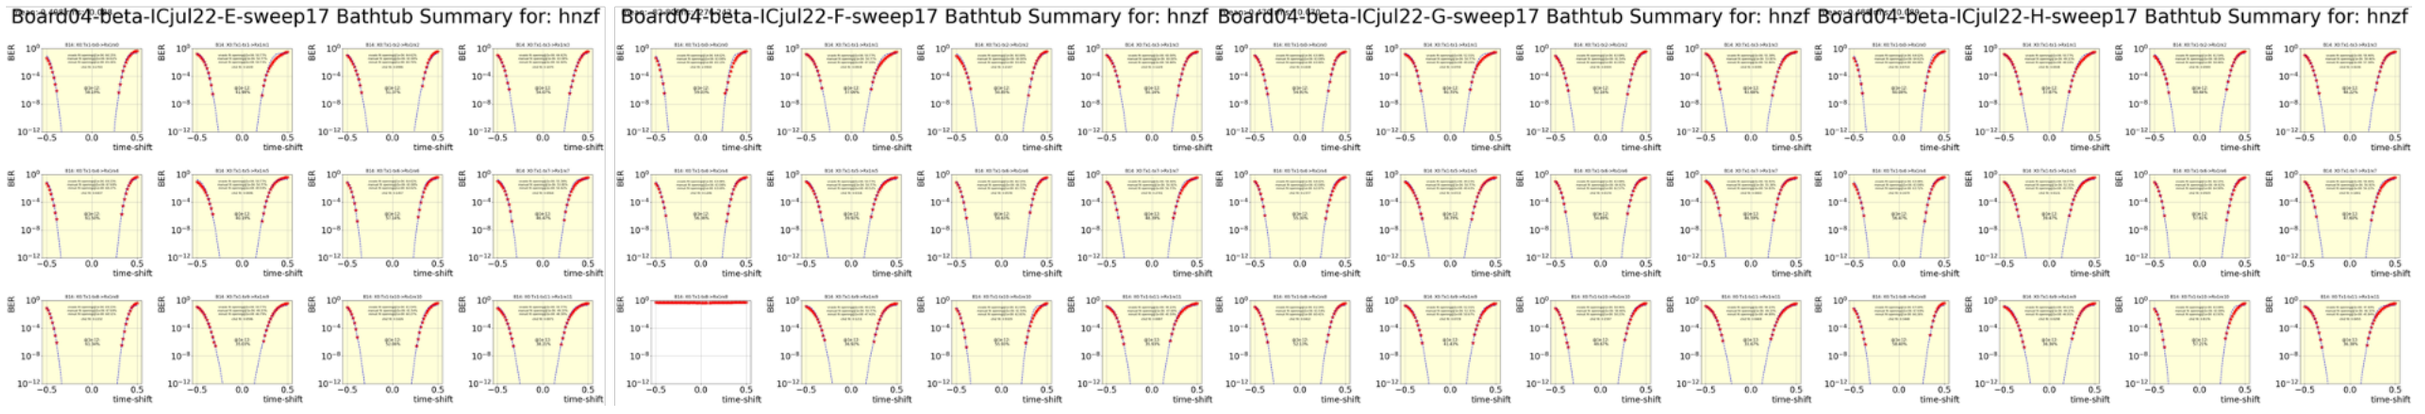

Optical Modulation Amplitude

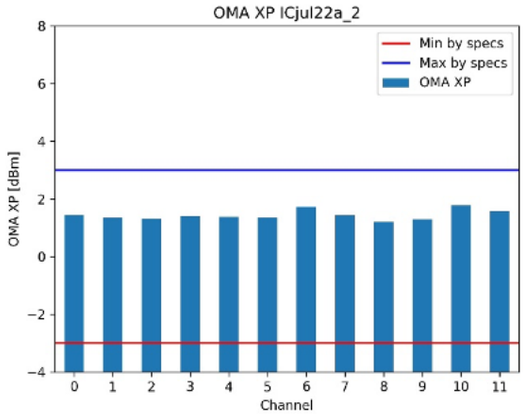

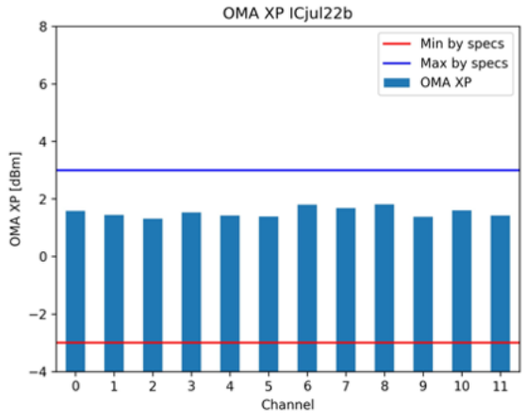

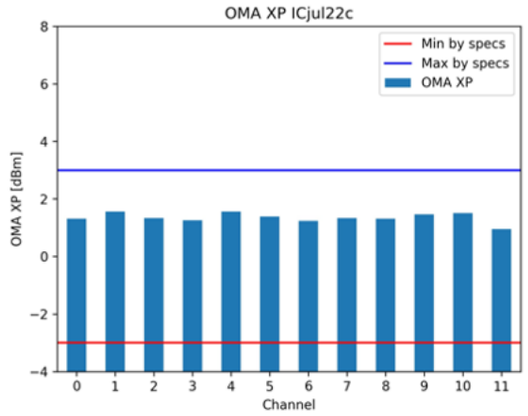

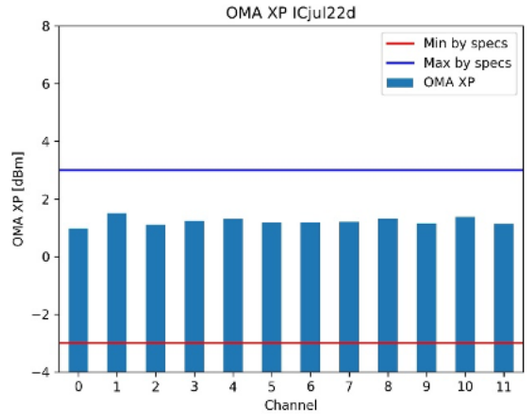

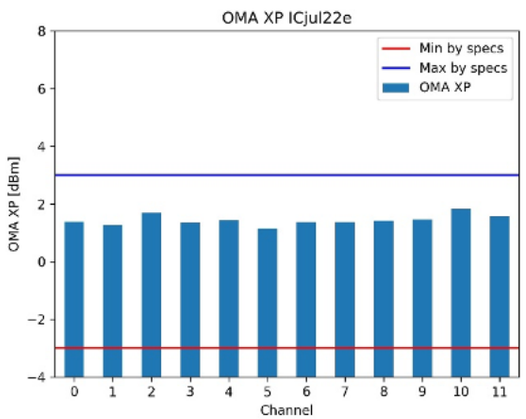

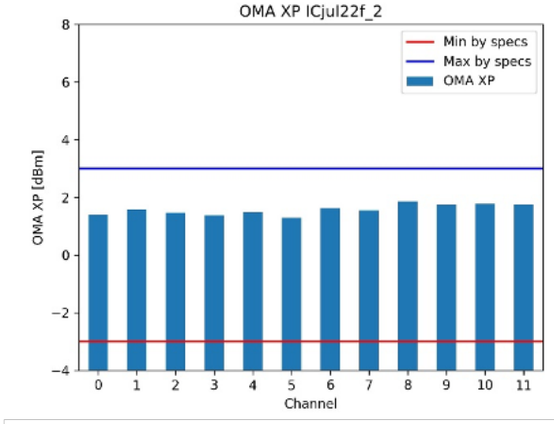

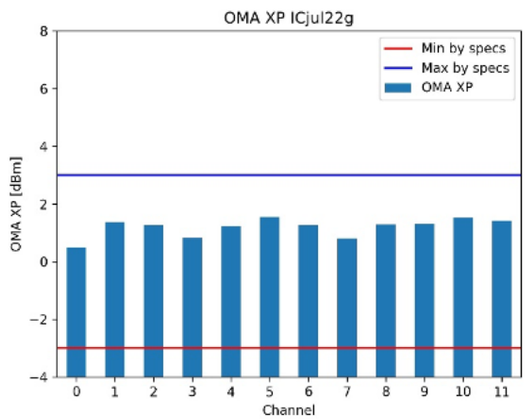

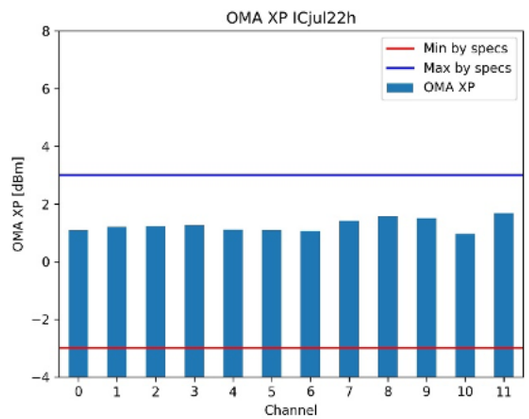

Extinction Ratio

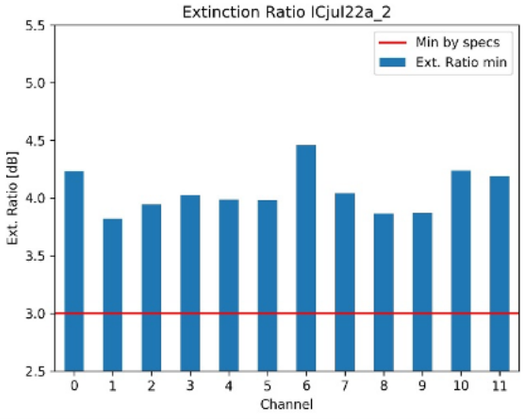

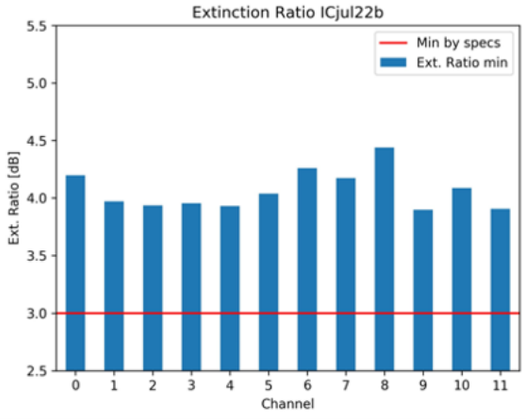

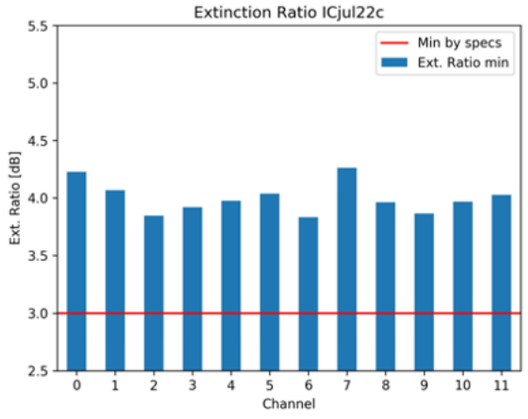

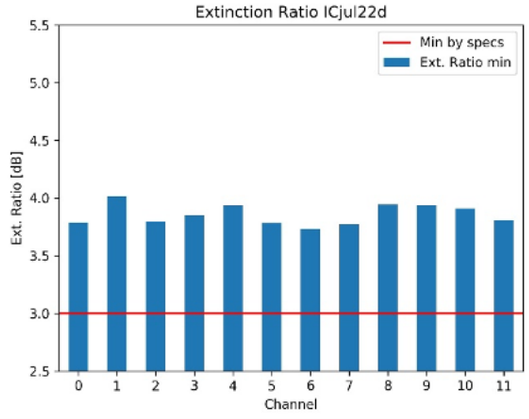

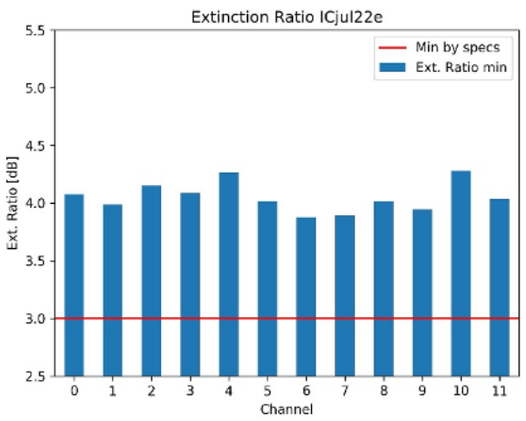

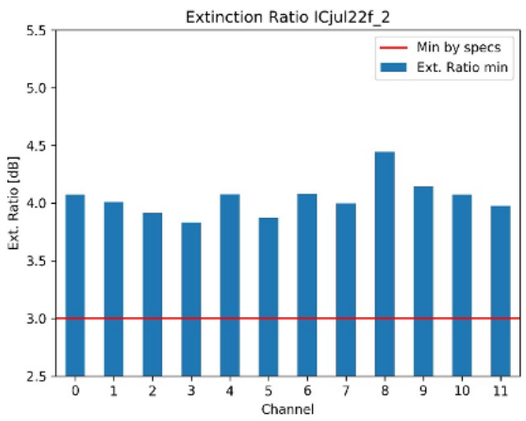

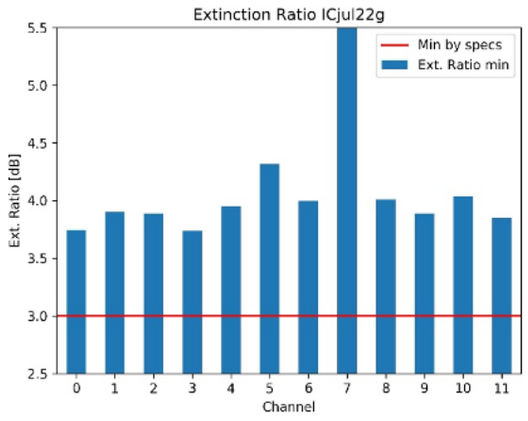

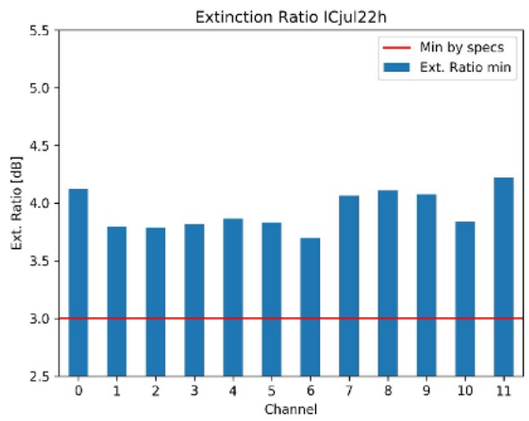

BER vs Attenuation (no corrections)

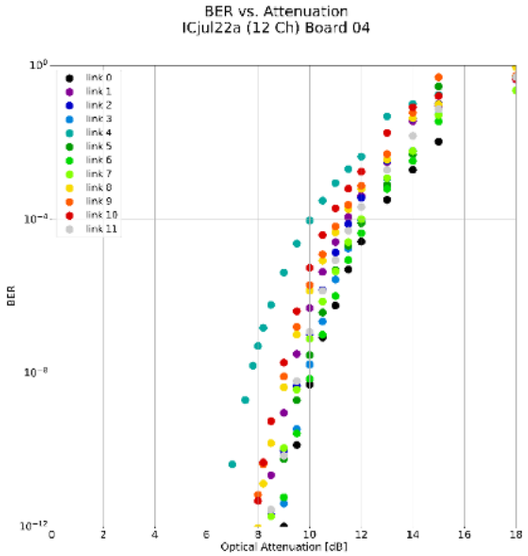

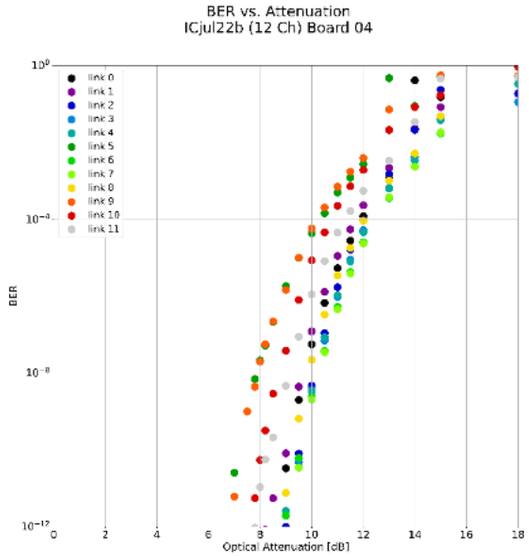

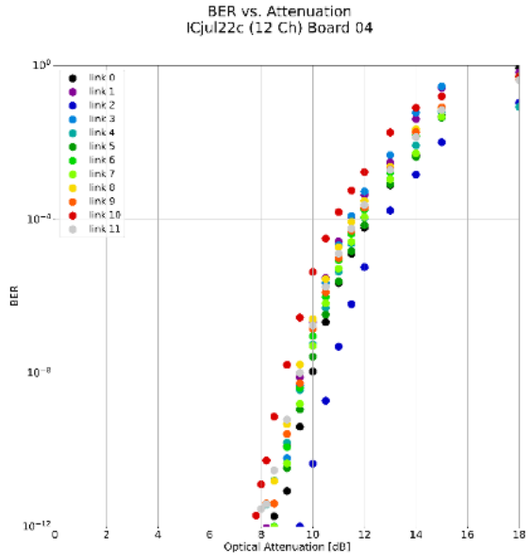

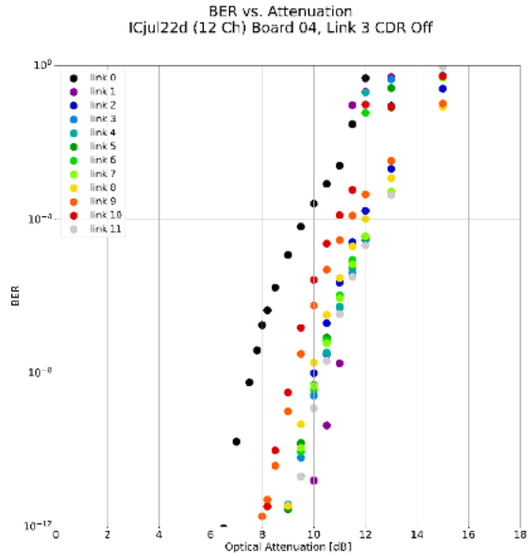

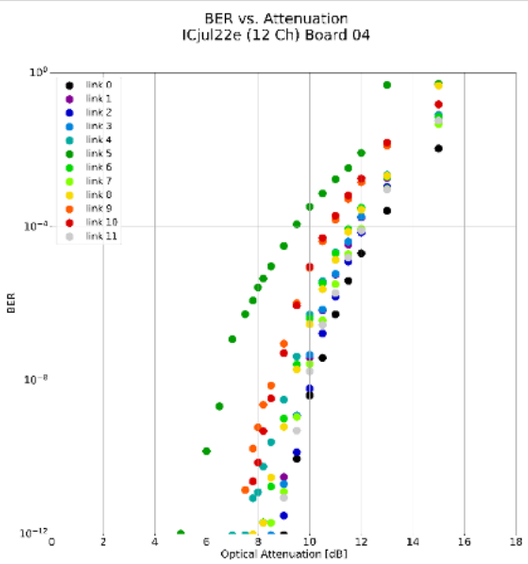

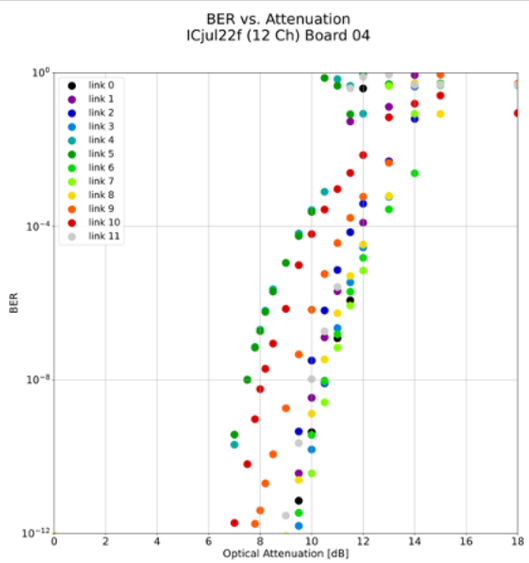

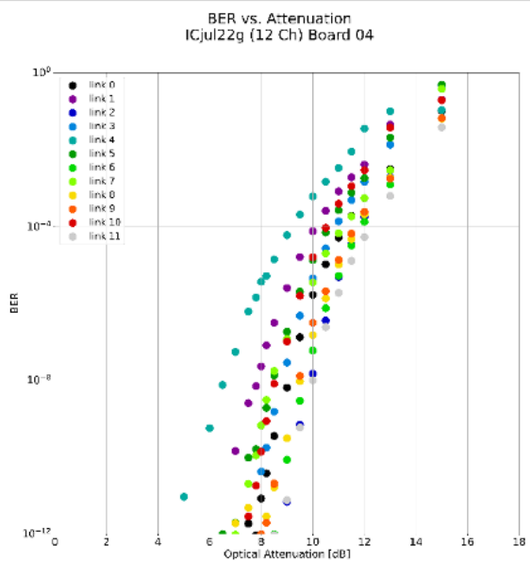

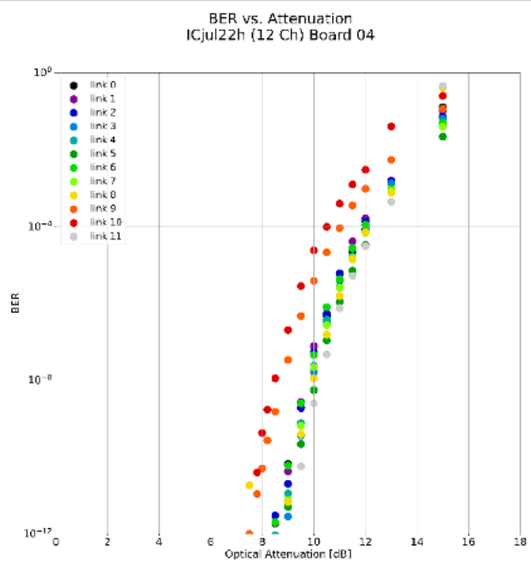

BER vs Rx OMA (corrected)