PANDA Compute Node Upgrade: Status and Plans

Thomas Geßler[1] Wolfgang Kühn[1] Jens Sören Lange[1]
Yutie Liang[1] Zhen'An Liu[2] Jingzhou Zhao[2]

[1]II. Physikalisches Institut, Justus-Liebig-Universität Gießen

[2]Trigger Lab, Institute of High Energy Physics, Chinese Academy of Sciences

PANDA DAQ FEE Workshop

May 29, 2018
GSI

May 2015: Compute Node Carrier Board (CNCB) v3.3

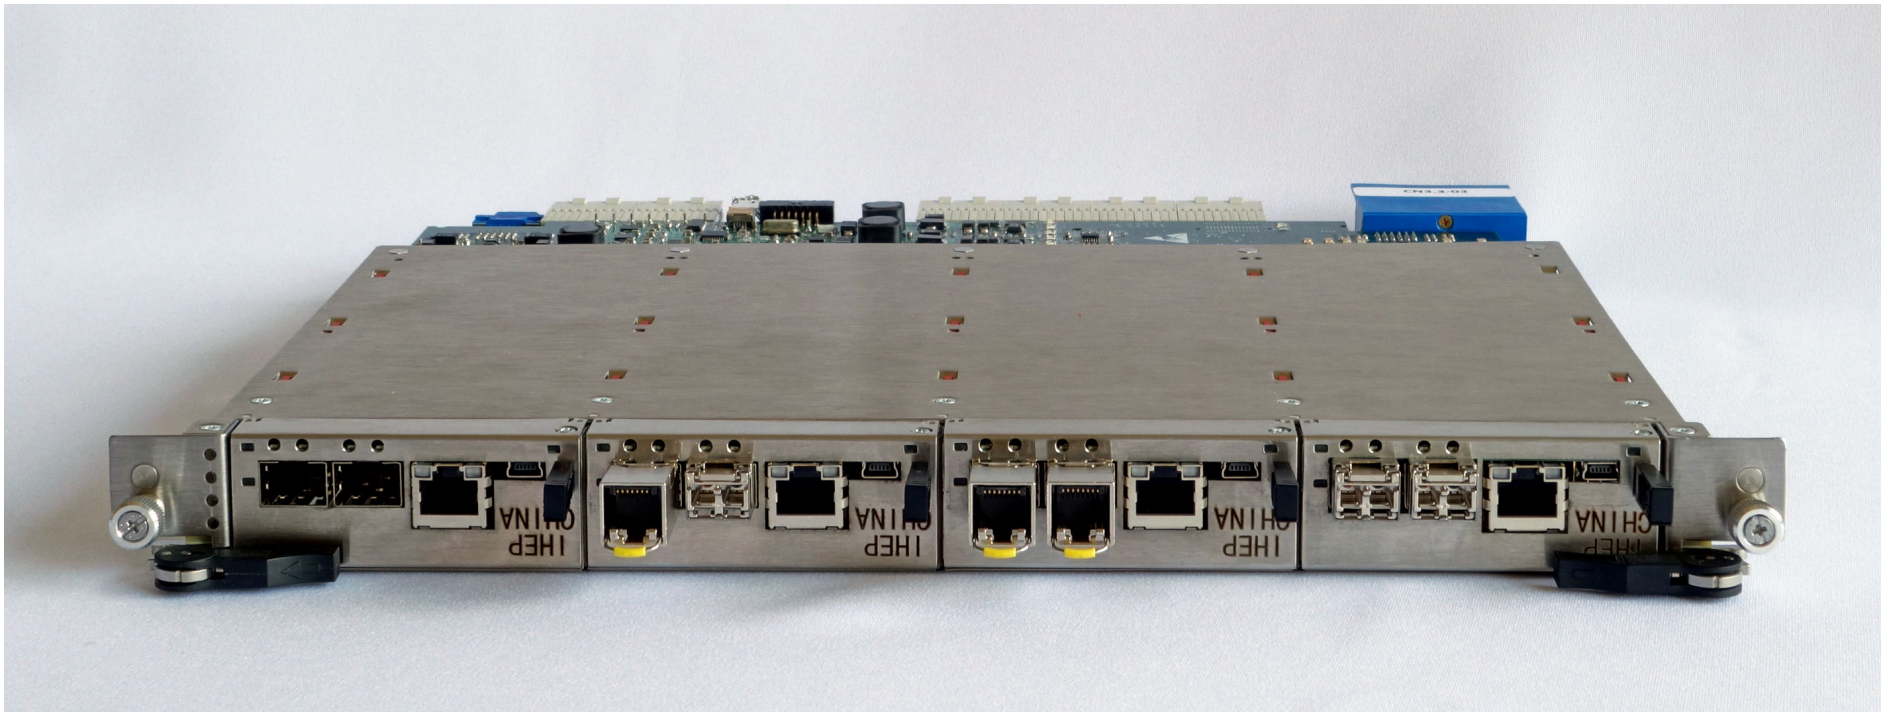

- Four full-width **AMC slots**
- Virtex-4 FX60 FPGA as **switch** to ATCA backplane

May 2015: Compute Node Carrier Board (CNCB) v3.3

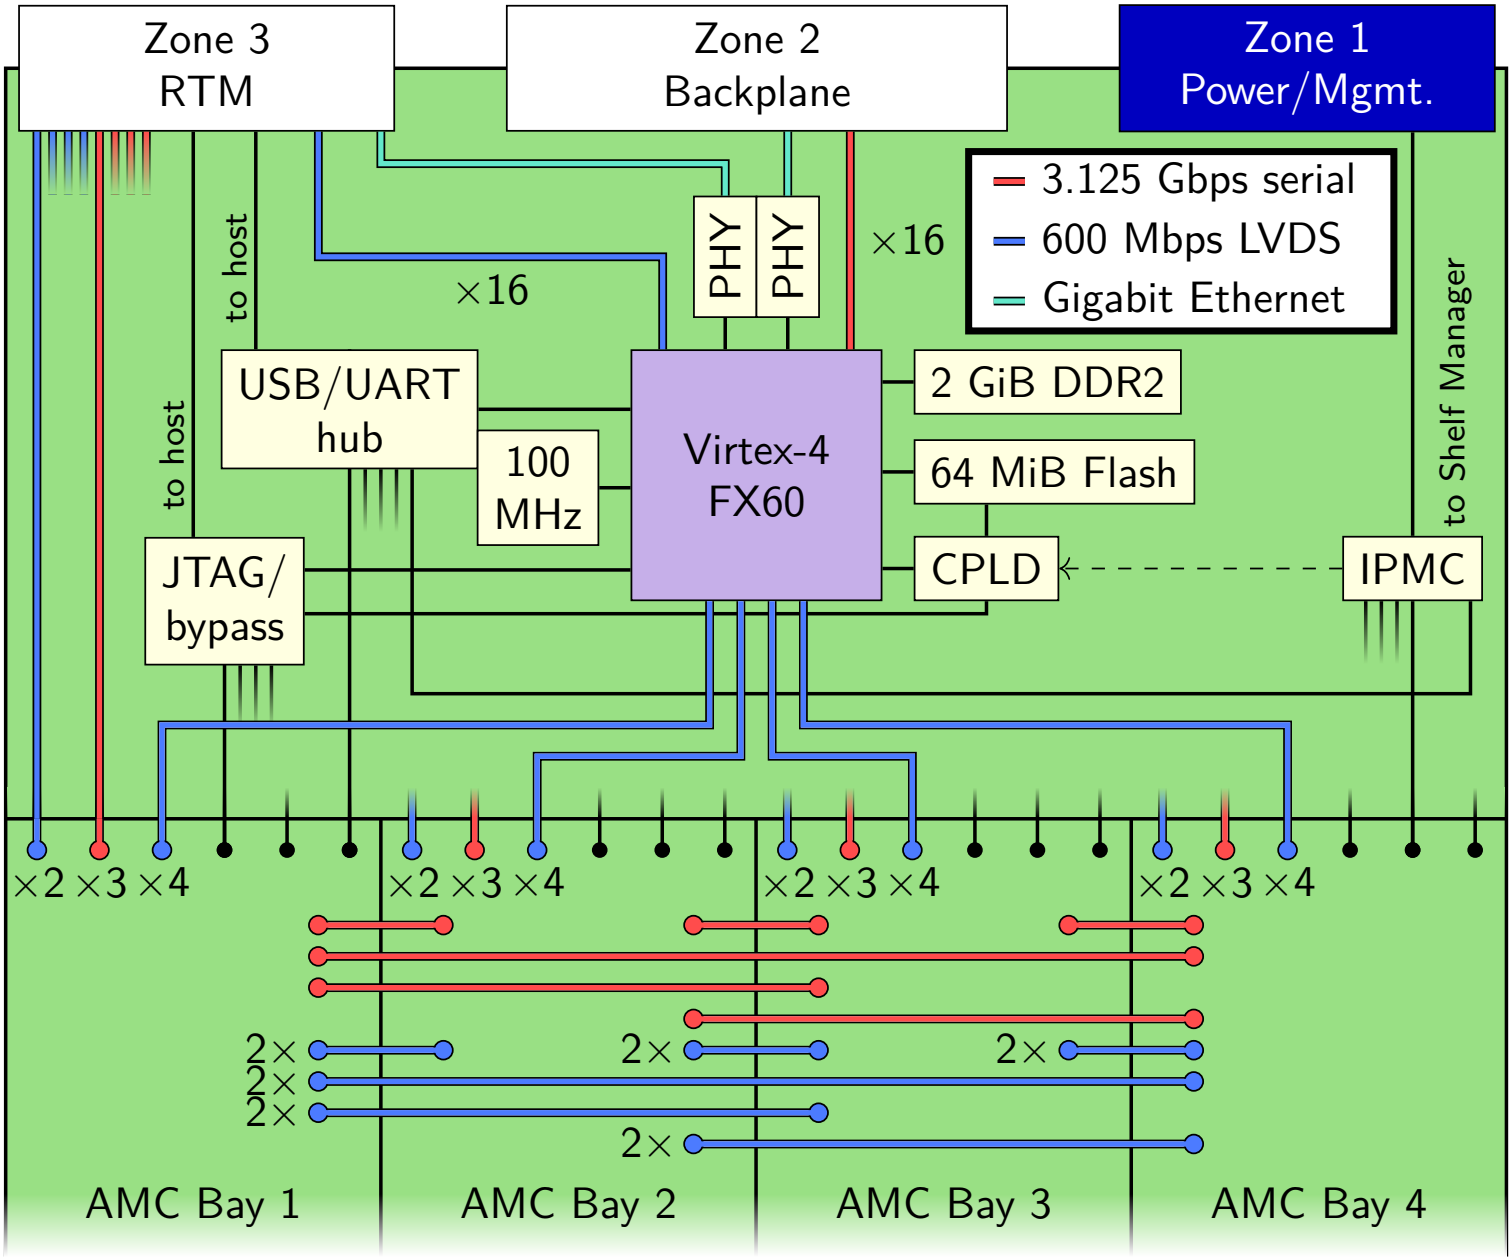

Nov 2014: xTCA-Based FPGA Processor (xFP) v4.0

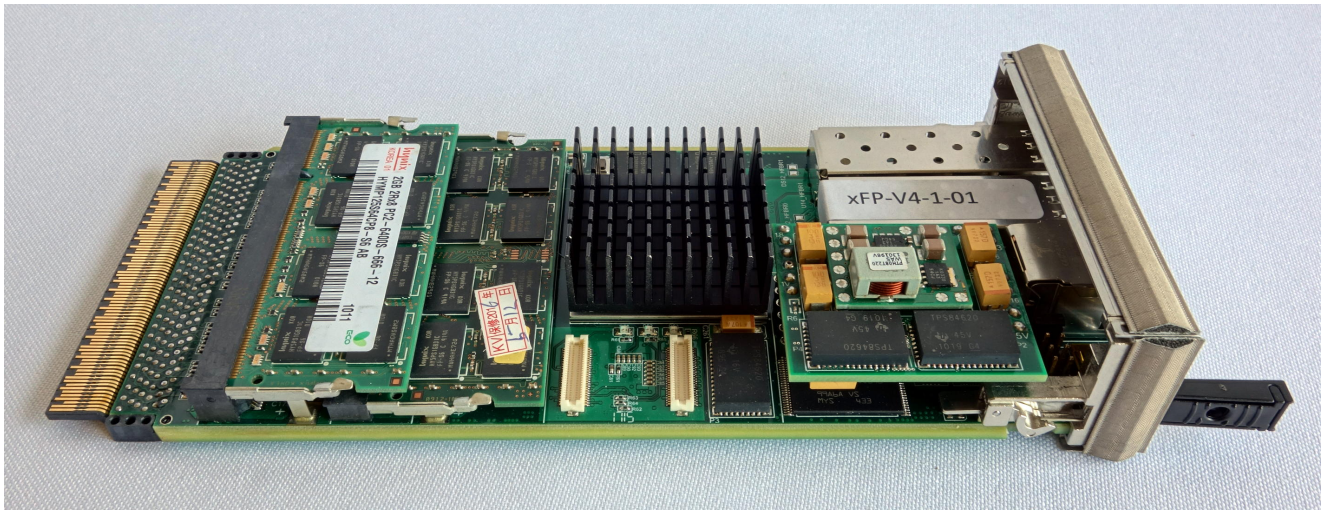

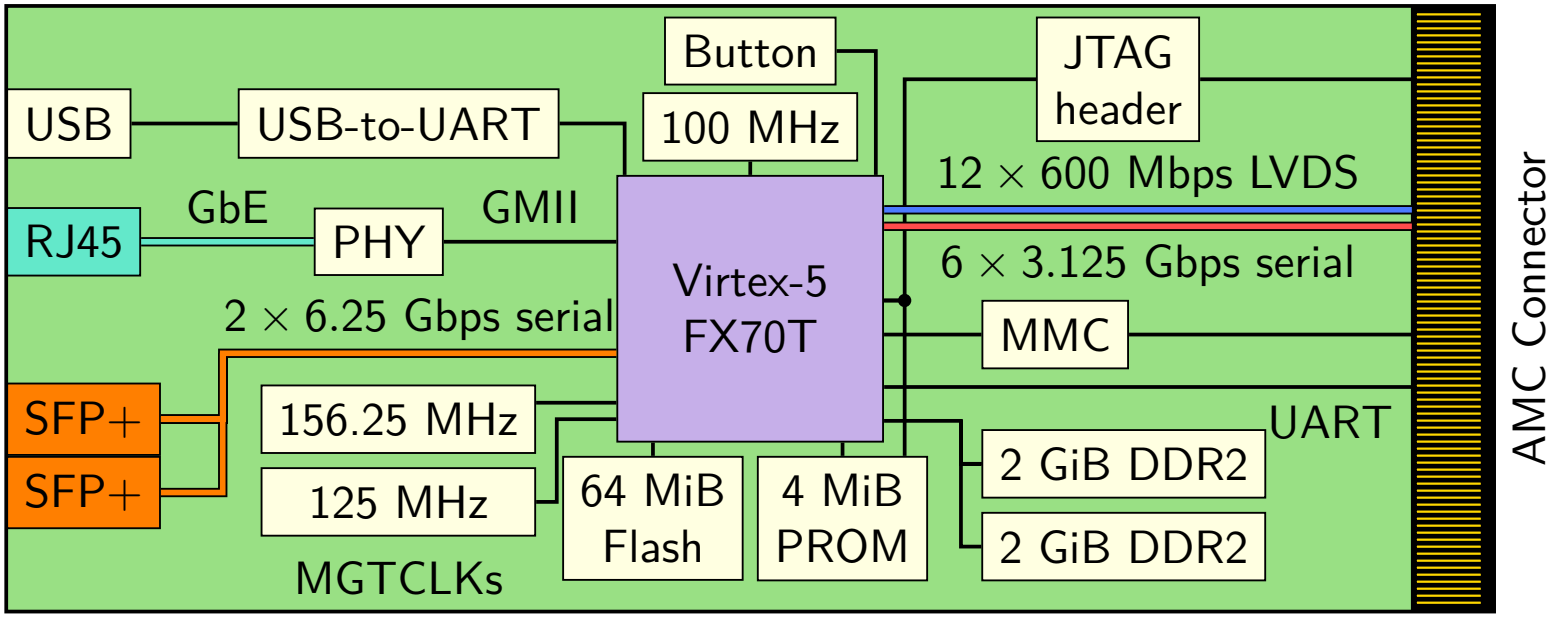

Compute Nodes in the ATCA Shelf

Two fully equipped Compute Nodes, one backplane GbE switch

Prototype Triggerless DAQ (Tests by Milan Wagner)

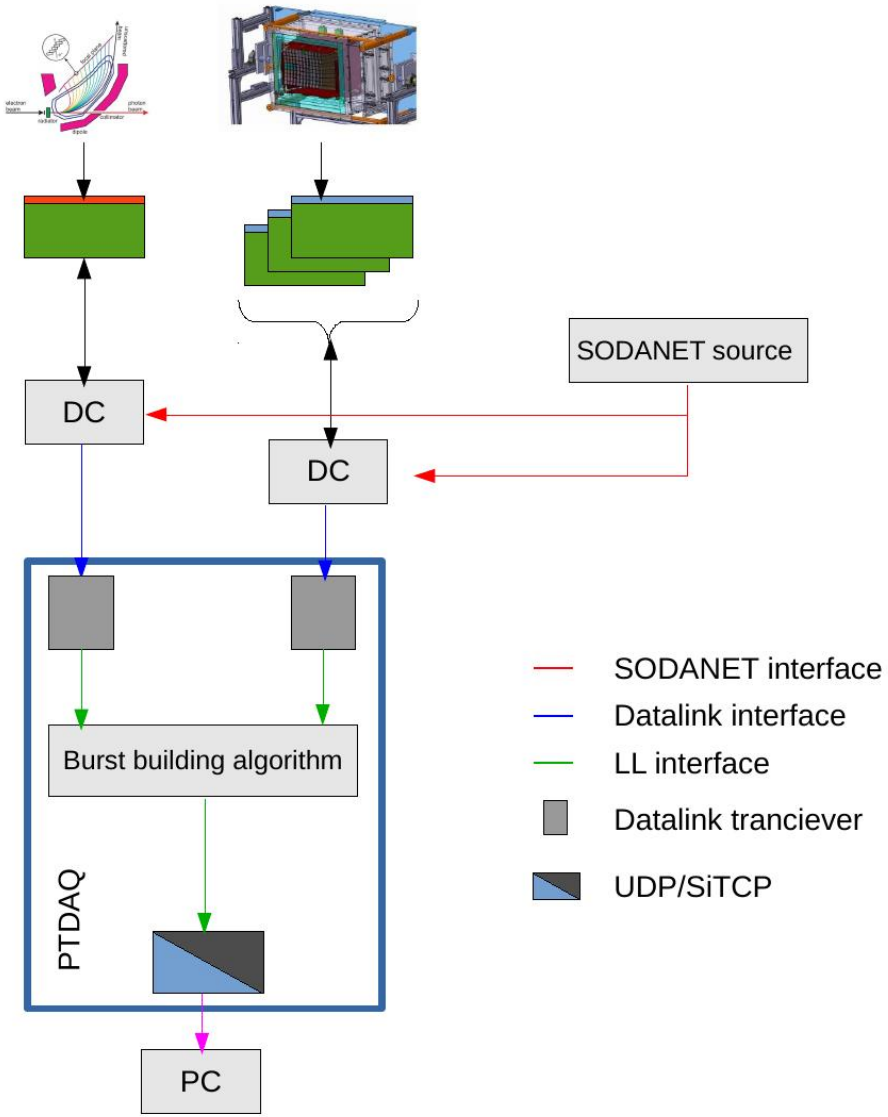

- Tested during **EMC beam-test** at MAMI in 2015
- Read out with **PTDAQ**:
  - 48 crystals of Proto120
  - 16 channels of the photon tagger
- **Successful burst-building** of 88 million events without errors

Belle II: PXD Data Reduction with Regions of Interest

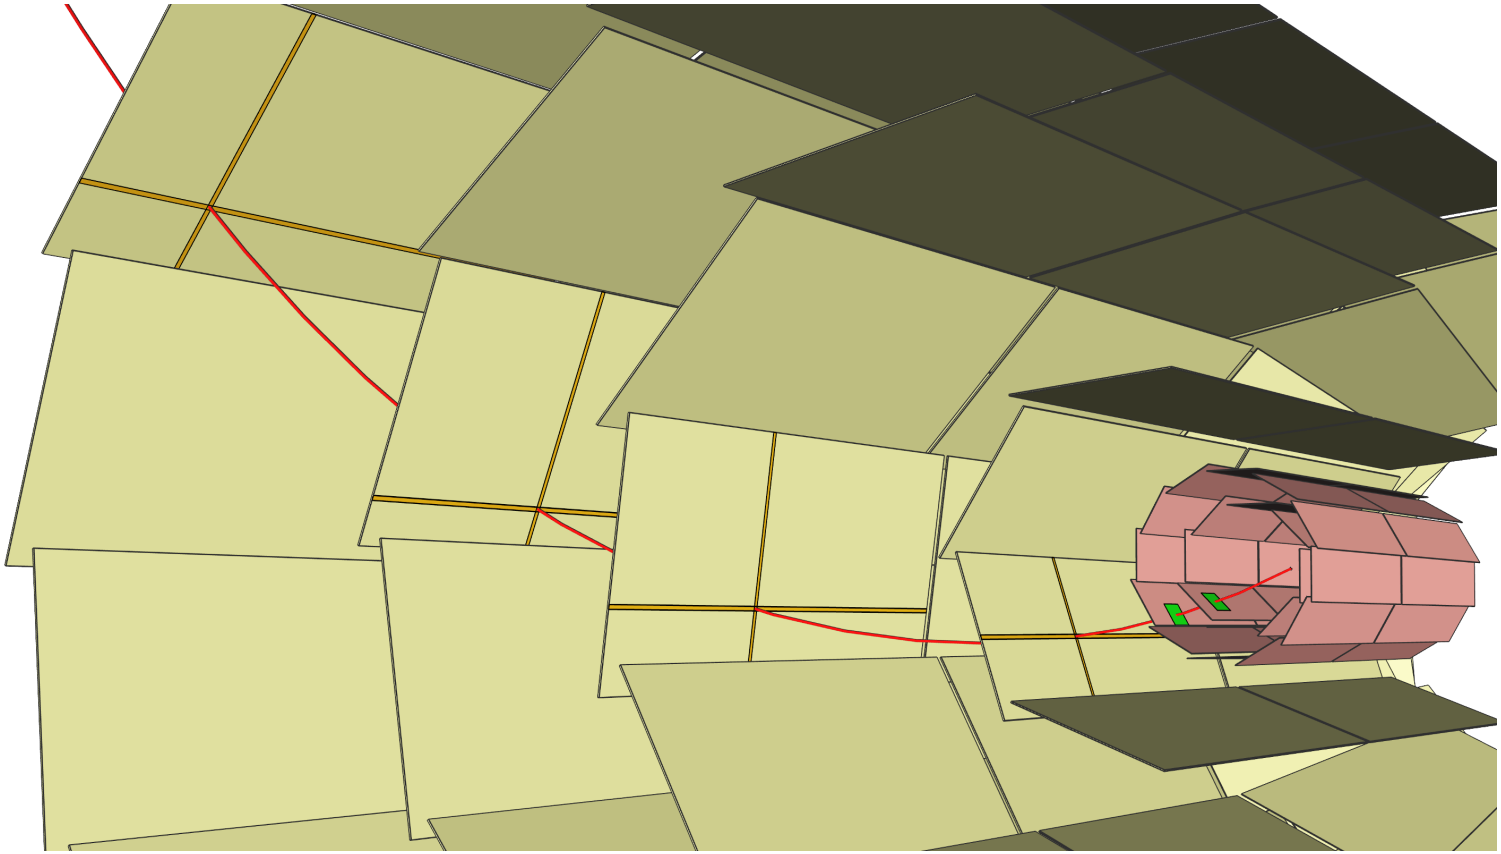

- Compute Nodes **store pixel-detector data** (20 GB/s) in RAM
- **Regions of interest** are received from high-level trigger (PC farm) and applied online
- Remaining pixels are **sent to event-building system**

Belle II: Compute Node Data Flow

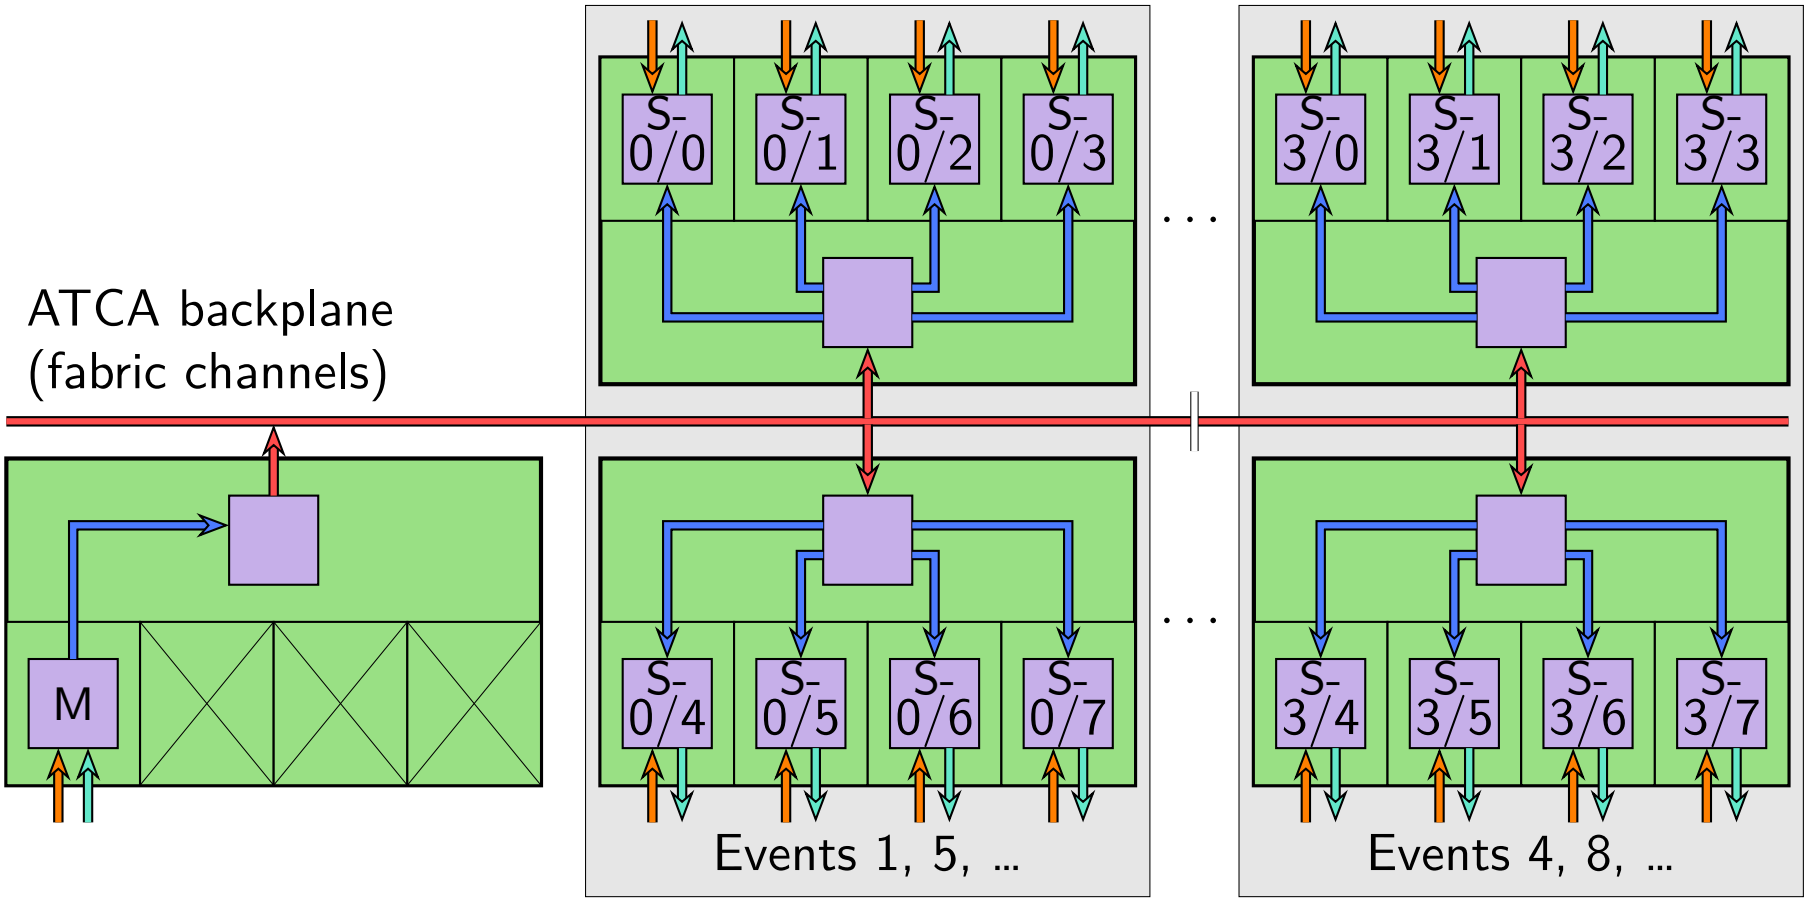

- **Selector (S):** Receive–store–reduce–send pixel data
- **Merger (M):** Receive ROIs, distribute over ATCA backplane
- Currently running in **Belle II Phase 2** with 1 Merger, 8 Selectors

Compute Node Upgrade: Carrier Board

- **First step:** upgrade Carrier (but stay compatible with current AMC)
- **FPGA:** Change to Xilinx UltraScale architecture

	Virtex-4 FX60 (CNCB)	Virtex-5 FX70T (xFP)	Kintex UltraScale 060 (Upgrade)
Registers	50k	44k	663k
LUTs	50k × 4-input	44k × 6-input	332k × 6-input
DSP Slices	128	128	2760
BRAM	4 Mb	5 Mb	38 Mb
MGT	16 × 6.5 Gbps	16 × 6.5 Gbps	32 × 16.3 Gbps
CPU	PPC405	PPC440	-

- **No more hard-core CPU** → Slow control on MicroBlaze or light-weight option like IPbus
  - Belle II experience shows that Linux-based slow control adds a lot of complexity

Compute Node Upgrade: Carrier Board

- **RAM:** 2 GiB DDR2 SODIMM $\rightarrow$ 16 GiB DDR4 (8 chips)
- **MGTs:** 6.25 Gbps $\rightarrow$ 16.3 Gbps
  - 4 links to each AMC card (currently: 4 $\times$ 600 Mbps LVDS)
  - 14 links to ATCA backplane
  - 1 link to RTM (10G Ethernet)
- **GbE switch:** 4 AMCs, 1 switch FPGA, 1 uplink to ATCA Base Interface
- **Configuration:** Flash/CPLD (slave serial) $\rightarrow$ automatic from NOR Flash (master BPI)
- **Programmable MGT clock**
- CPLD as **JTAG hub**
- **Keep:**
  - I2C buses, sensors
  - IPMC connector

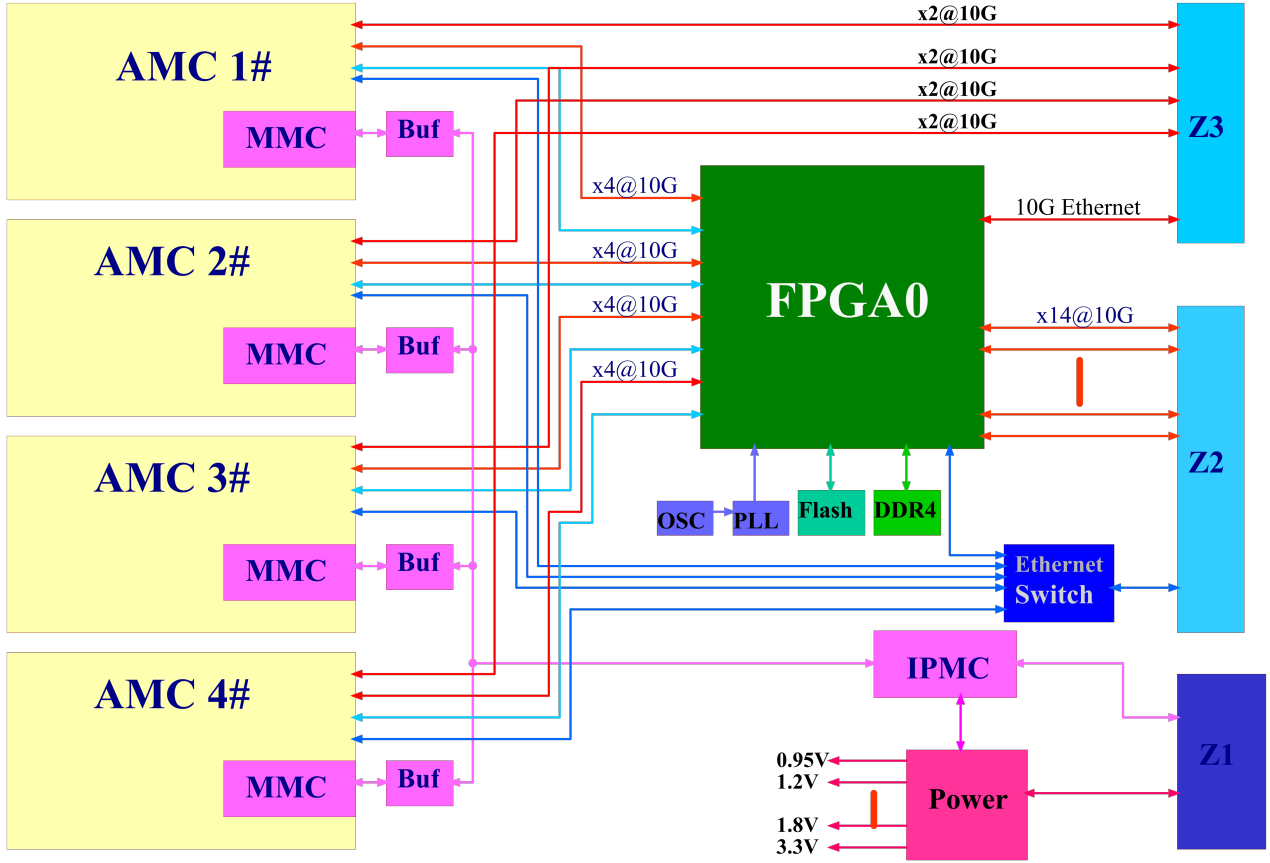
AMC 1#
MMC
Buf
AMC 2#
MMC
Buf
AMC 3#
MMC
Buf
AMC 4#
MMC
Buf
x2@10G
x2@10G
x2@10G
x2@10G
Z3
x4@10G
x4@10G
x4@10G
x4@10G
FPGA0
10G Ethernet
x14@10G
Z2
OSC
PLL
Flash
DDR4
Ethernet
Switch
IPMC
Z1
0.95V
1.2V
1.8V
3.3V
Power

CNCB v4.0: First Prototype

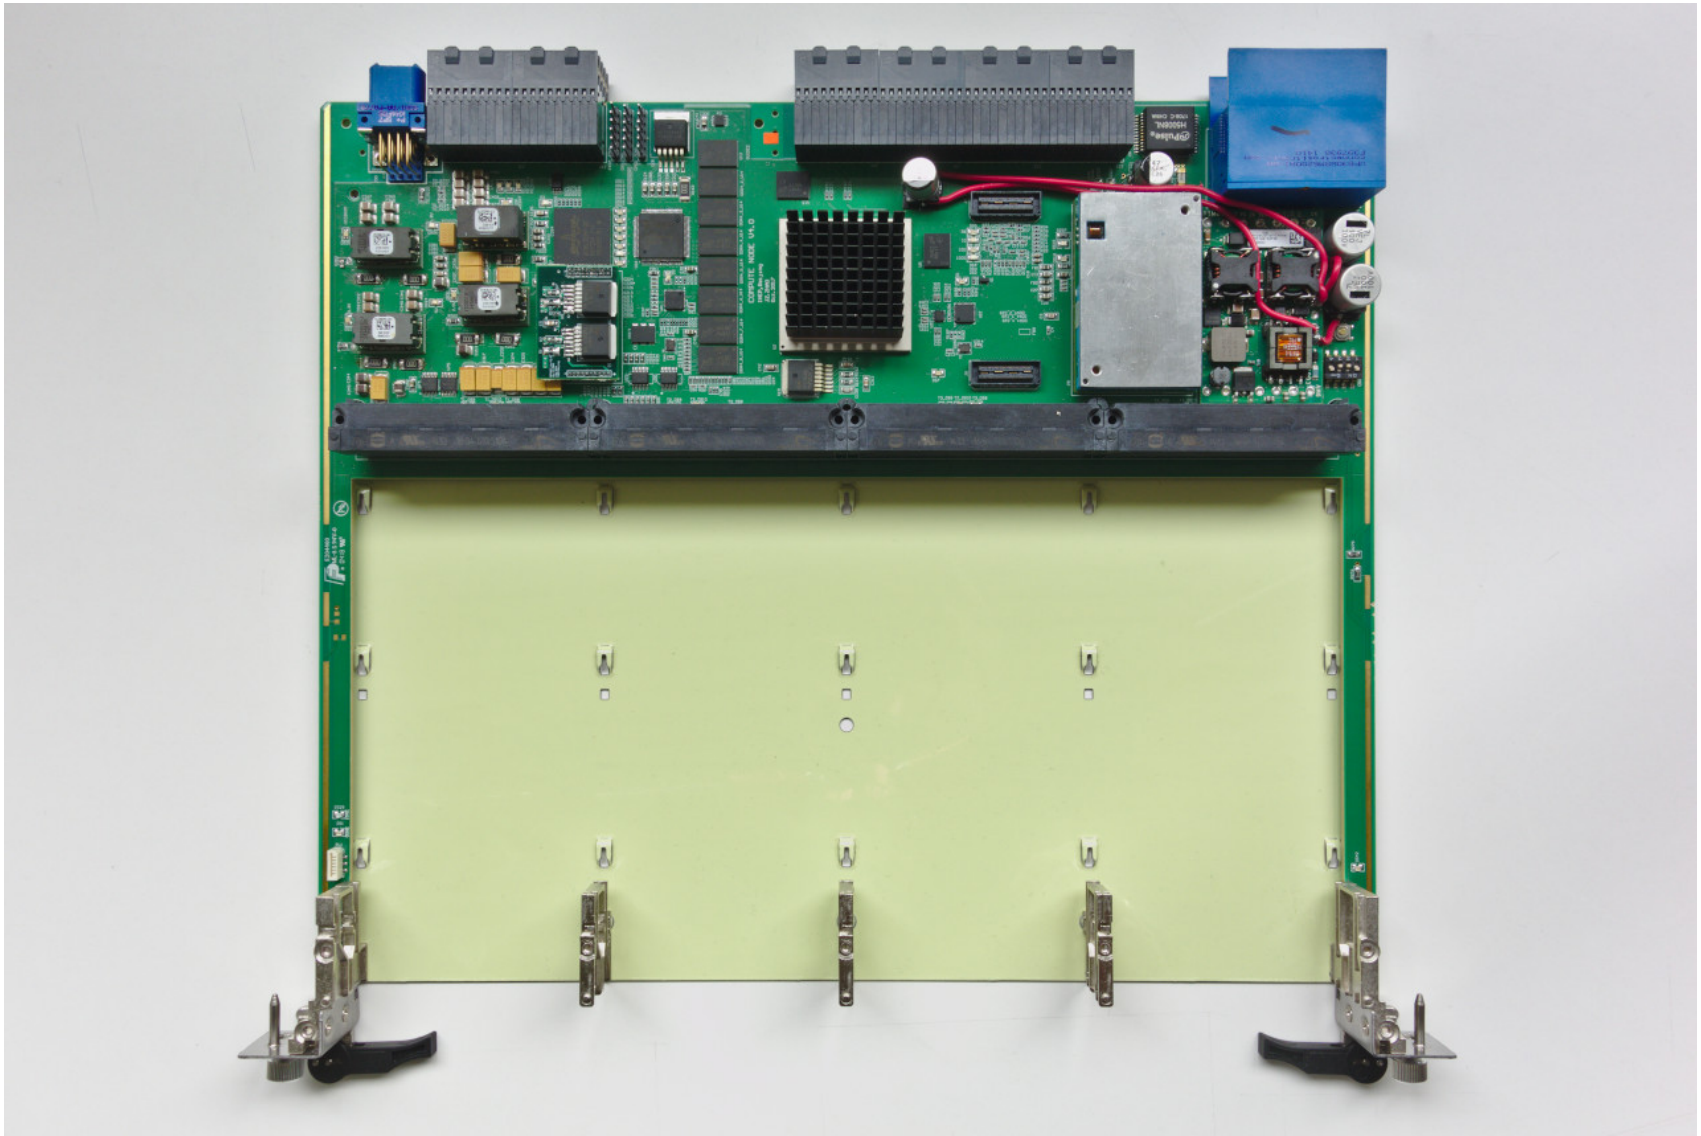

Two prototype boards were produced, then tested in Gießen and at IHEP

CNCB v4.0 Prototype Tests

	Gießen	IHEP
FPGA access	OK	OK
RAM	OK (16 GiB tested)	OK (2 GiB tested)
JTAG hub	OK (Carrier + AMCs)	OK (only Carrier tested)
PLL chip	OK	OK
Eth. switch	OK (Carrier, AMCs, backplane)	not tested
Flash chips	OK	not tested
Auto config.	OK	not tested
AMC links	OK (6.25 Gbps Aurora)	OK (12.5 Gbps loopback)
Backplane links	OK (3.125 Gbps Aurora)	OK (12.5 Gbps loopback)
IPMC interface	OK	not tested
Clock fan-out	not tested	OK
Linux on MB	OK	not tested

(IPMI functionality tested by Björn Spruck)

$\Rightarrow$ all essential functions tested successfully at at least one site

CNCB v4.0 Prototype Remaining Issues

- Main issue: **voltage drop** of several supply voltages, especially at high loads
  - 1.80 V $\rightarrow$ 1.69 V
  - 0.95 V $\rightarrow$ 0.91 V
  - Reason: DC/DC converters' remote-sense not correctly connected
  - Will be fixed in the next iteration
  - Compensated by manual adjustment $\rightarrow$ no noticeable impact
- Prototype was equipped with FPGA **speed grade** -1 instead of -2
  $\rightarrow$ MGT links tested at 12.5 Gbps instead of 16.3 Gbps
- The **MGT link to the RTM connector** (for 10G Ethernet) works only up to 3 Gbps (tested with loopback adapter at IHEP)
- DDR4 with the nominal memory clock (1000 MHz), but only **very slow R/W speeds** (~ 20 MB/s, MicroBlaze)
  $\rightarrow$ Should be tested with high R/W speeds, up to the hard limit of 16 GiB/s

Compute Node Upgrade: Rear-Transition Module (RTM)

- Two **RTM prototypes** have been produced
- Features:
  - On-board **USB-JTAG programmer** (Digilent)
  - **UART-USB interface** for 4 AMC cards + switch FPGA
  - **USB hub** for UART interfaces, IPMC
  - **SFP+ cage** for switch-FPGA 10G Ethernet
- JTAG and USB functions **successfully tested** at IHEP
- Will arrive at Gießen soon for **first tests**

Outlook

- After prototype tests are concluded, a **second design iteration** will be produced and tested
- Next step: **upgrade of AMC card**
  - Similar upgrade of FPGA and RAM
  - Many more high-speed front I/Os (e.g., Avago MiniPod)
  - In the meantime, current AMC cards can be used in the new carrier board
- Also interesting for **Belle II PXD upgrade** $\rightarrow$ **synergy**