A Readout Driver for the ATLAS LAr-Calorimeter at a High Luminosity LHC

Andy Kielburg-Jeka, Steffen Stärz

On Behalf of the ATLAS LAr Calorimeter Group

TWEPP 10 Aachen, September 20-24, 2010

Outline

- ATLAS detector and current LAr calorimeter read-out system
- upgrade to HL-LHC
- readout scheme for HL-LHC
- hardware development status
- summary and outlook

The ATLAS Detector

- 4 high granularity LAr calorimeters

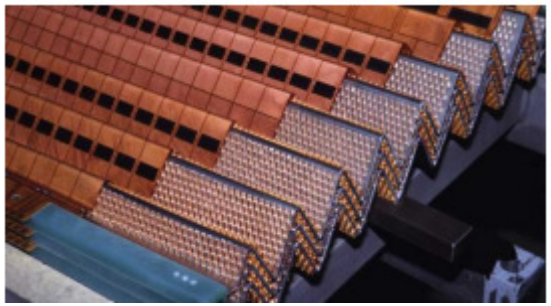

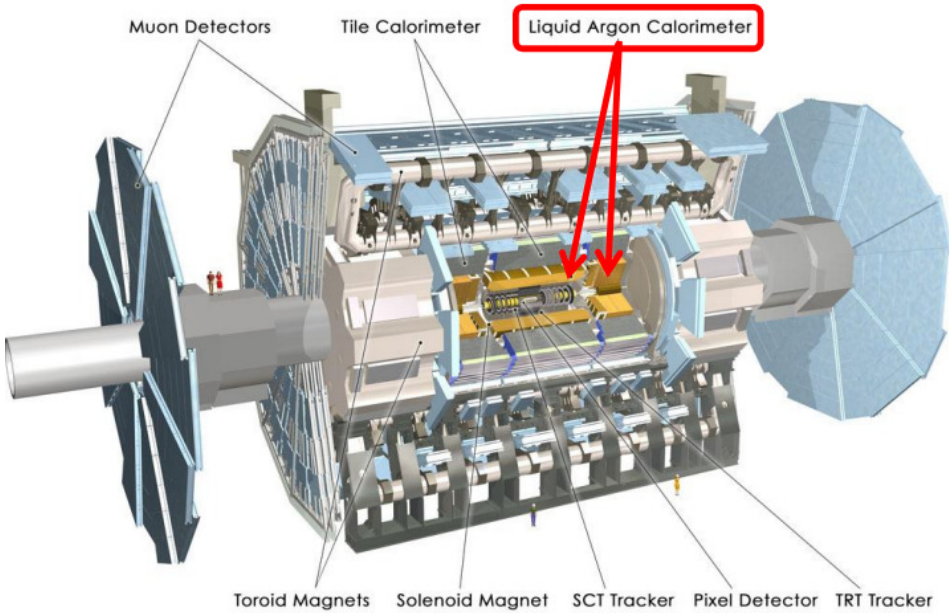

- 182486 readout channels
- 40MHz proton-proton collision rate

- front-end and trigger-sum electronics
  - mounted on detector,
  - operates in radiation environment

- back-end electronics and additional trigger logic
  - in counting room outside the experimental cavern

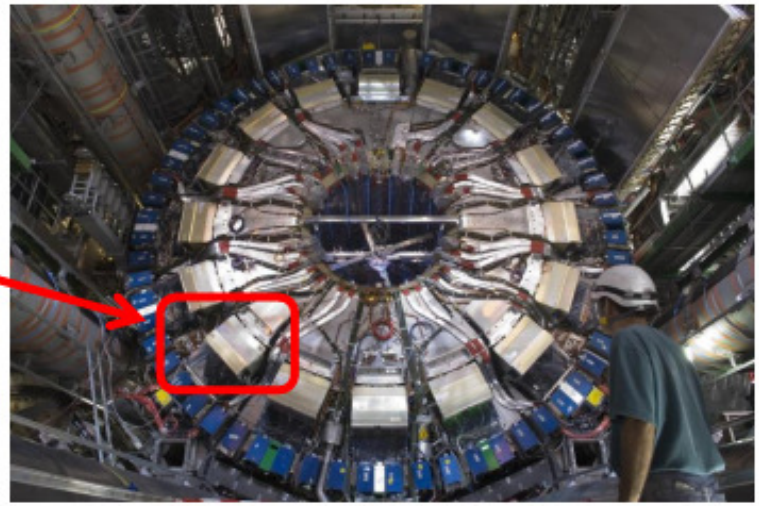

Current LAr Calorimeter Readout System

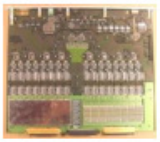

- 1524 front-end boards (FEB)
  - up to 128 channels
  - preamp, pulse shaping, buffer and sampling
- analog signal to L1 calorimeter trigger (@ 40MHz)
- readout clock is synchronous with LHC bunch-crossing (25 ns)

- connected to off-detector electronics by
  - 58 copper trigger cables
  - 1600 optical links

- 192 read-out driver boards (ROD)
  - digital filter
  - data transferred only after L1 accept @ 100kHz

- 800 optical links to DAQ PCs

- 68 read-out buffer PCs (ROB)
  - DAQ and high-level trigger buffer

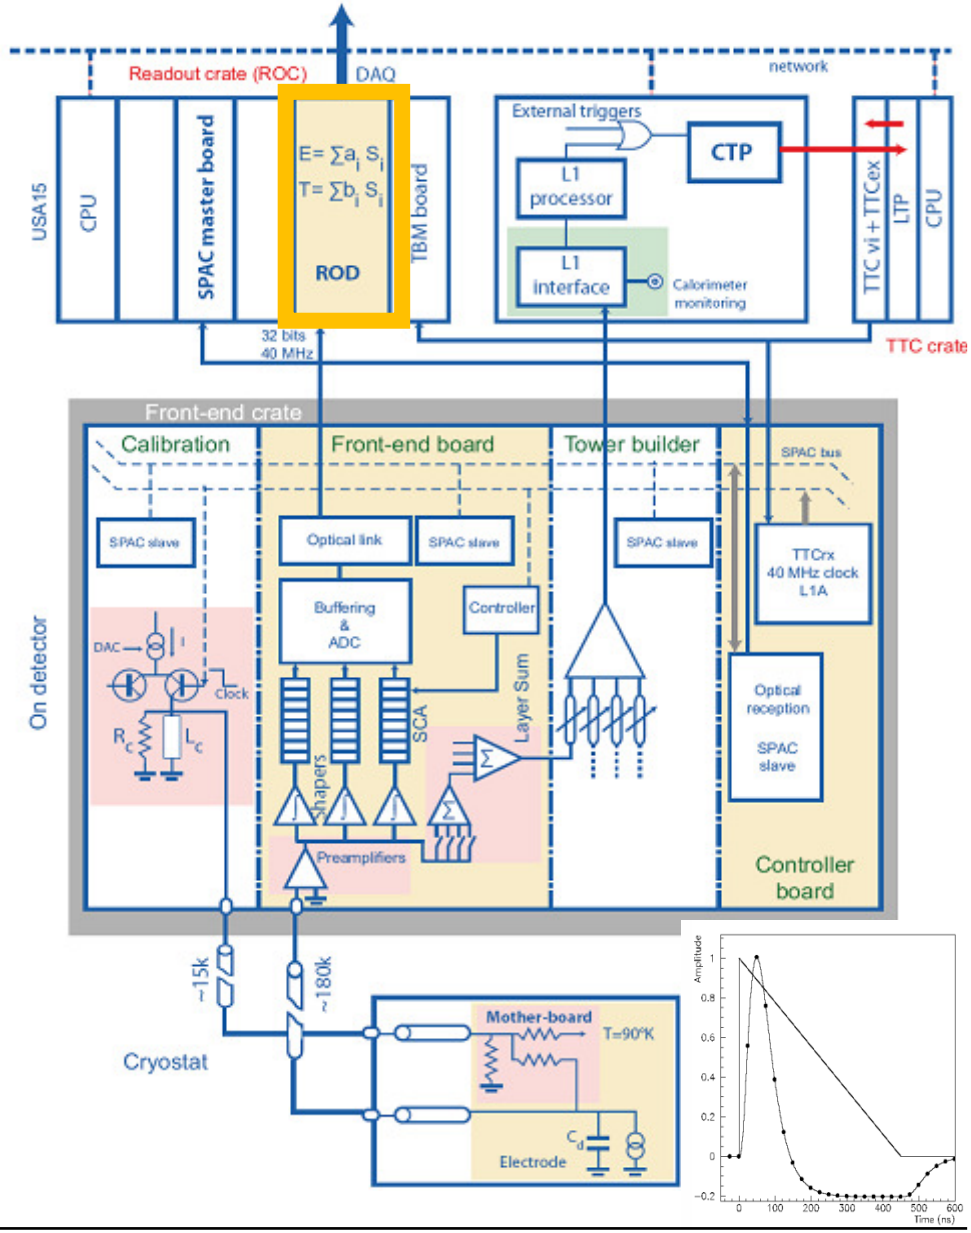

Upgrade to High Luminosity LHC

- Several reasons for upgrade:
  - radiation tolerance of front-end electronics: HL-LHC reaches up to 3000 fb^{-1}
    - qualified for ~700 fb^{-1} (+ safety margin)

    → not possible to operate hardware until end of HL-LHC
  - ageing / reliability concerns of current electronics:
    - some parts exceed 20 years age at HL-LHC start
  - limited size of on-detector data buffers:
    - current readout architecture does not allow increase in L1 latency nor bandwidth
  - if readout is upgraded → chance to increase granularity of trigger input

- HL-LHC challenges:

  ~ 5x higher radiation levels for on-detector electronics
  ~ 5x higher total ionization dose for on-detector electronics
  ~ 5x more simultaneous proton-proton interactions and pile-up background

Trigger Latency Concerns

- Current LAr latency:
  - 96 bunch-crossings (BC)
  - +10 reserve BC
- Scenario with new trigger design:
  - example: modified muon trigger
  - trigger violates the LAr latency by 6.5 BC
  - would create dead-time

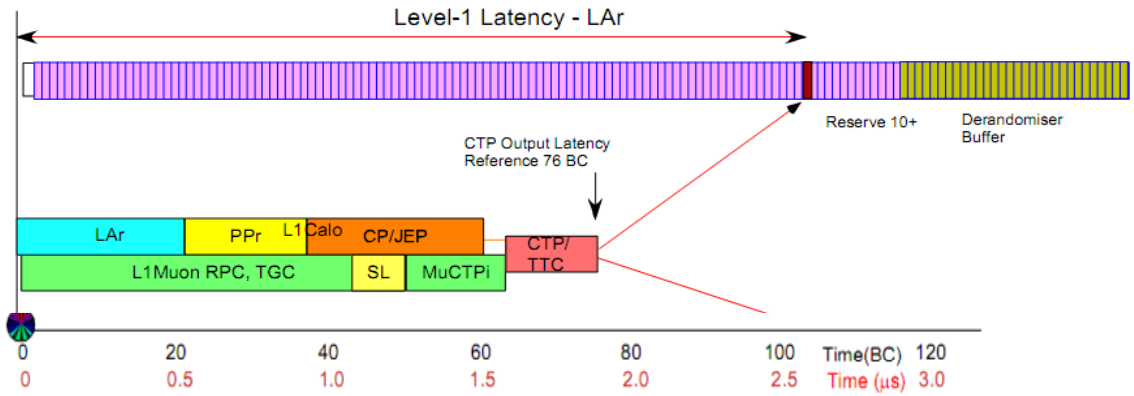

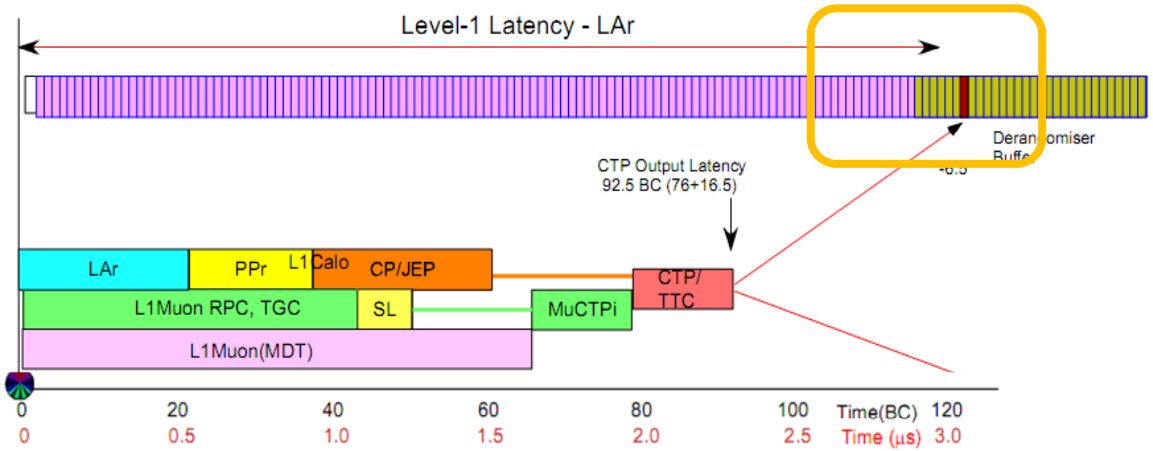

- larger buffer size needed
- digital buffer on new ROD planned to replace the analog buffer

ROD Input for High Granularity Trigger

- current performance limitations:
  - fixed analog trigger sums
- read-out and trigger upgrade at HL-LHC:
  - ROD prepares digital data for hardware triggers
  - higher granularity for better isolation of photons and electrons and pion tagging

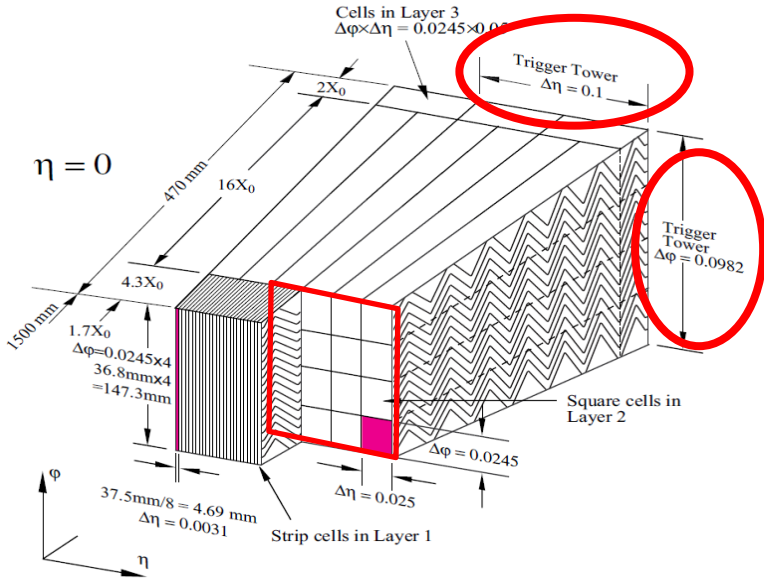

	today	upgrade scenario
hardware triggers	L1	L0 and L1
trigger input	analog	digital
trigger tower size in Δη x Δφ	0.1 x 0.1	up to 0.025 x 0.025
trigger tower depth	4 layers summed	4 layers separate
data size / 0.1 x 0.1 tower	16 bit	~200 bit
latency up to trigger input	1.2 µs (L1)	1.2 µs (L0) 6.4 µs or more (L1)

Requirements for New Hardware

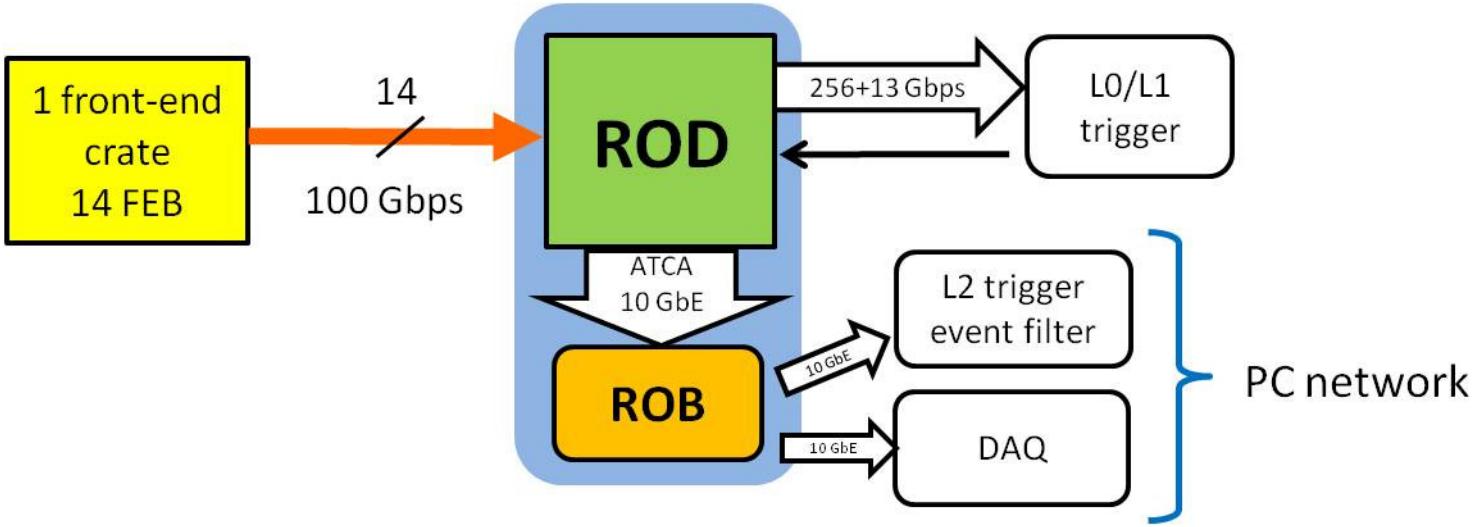

- Front-end electronics
  - preamp, shaping, digitization of data @ BC rate 40MHz (no buffering anymore)
  - high performance, radiation tolerant optical links needed for 100Gbps (see talk by A. Liu)

- Read-out driver (ROD)
  - 40MHz input rate, connected to 14 front-end boards if possible
  - digital calculation of energy, time, quality factor for each calorimeter cell
  - perform L0/L1 sum digitally based on raw data → higher granularity, flexible
  - send digitized data to L0 trigger @40 MHz
  - local data buffer, send data to read-out buffer (ROB) after L1A @ 100 kHz

Readout Scheme for HL-LHC

- each ROD receives 14 x 100Gbit/s = ~ 1.4Tbit/s
- high speed parallel fiber optical transceiver (e.g. 12 fibers @ 10Gbps)

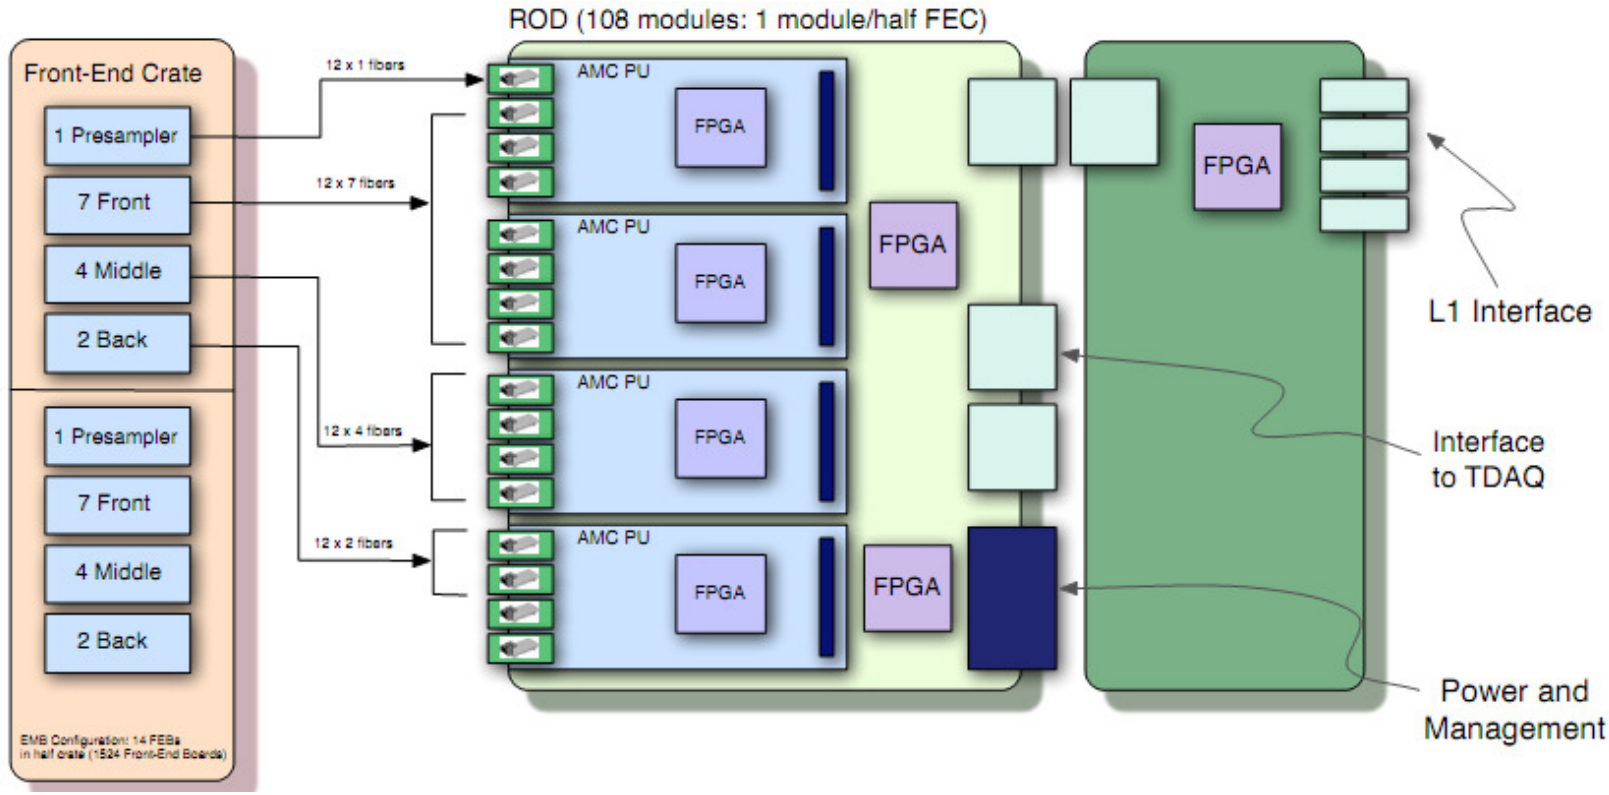

- ongoing ROD R&D:
  - use FPGAs for high performance SERDES
  - use FPGAs for parallel data processing with high bandwidth
  - follow up and explore technology evolution (FPGA, ATCA)

Sub-ROD + ROD Injector Development (BNL, Arizona)

- Sub-ROD and ROD Injector design:
  - ATCA form factor
  - 75 Gbps fiber-optic link
  - SNAP-12 parallel optical connector
  - FPGA SERDES: Xilinx Virtex 5 FX on Sub-ROD and Altera Stratix II GX on Injector
  - Control and Power logic (IPMI, power conversion, hot plug, …)
    - Sub-ROD: standard design
    - Injector: proprietary

- integration tests:
  - ATCA power and control scheme functioning
  - 6.25 Gbps per link reached → 75 Gbps total
  - used as platform for latency tests

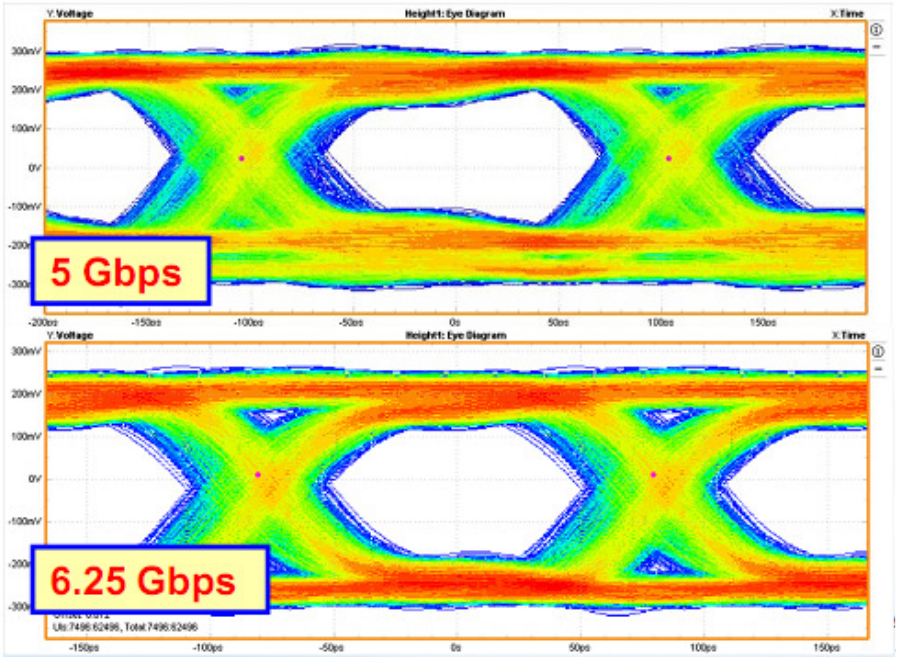

Architecture of AMC Sub-ROD Processing Unit (BNL)

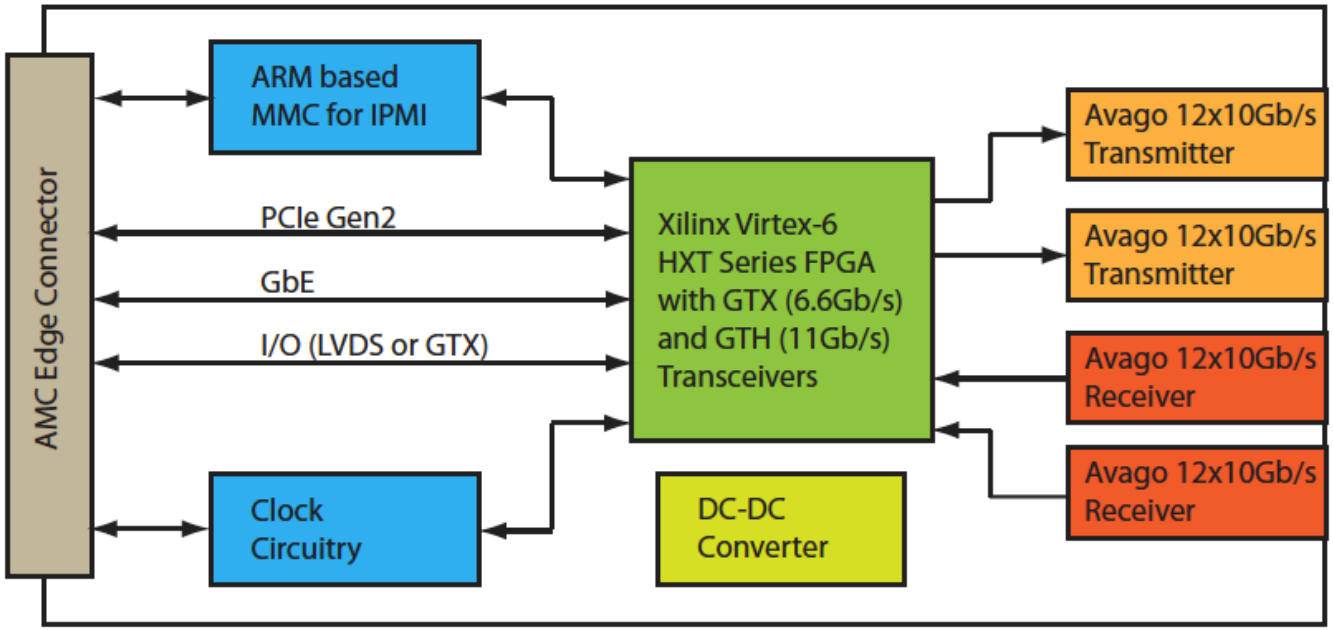

Block Diagram of AdvancedMC SubROD

- AMC format Sub-ROD Processing Unit can be tested in a low cost MicroTCA crate or on an ATCA carrier board
- MMC (Module Management Controller) is based on ARM processor for IPMI
- Precision clock circuitry and DC-DC converters for power management
- Based on high performance Xilinx Virtex-6 HXT series FPGA which has both GTX (6.6Gb/s) and GTH (11Gb/s) transceivers
- High speed (12x10Gb/s) parallel fiber optical transceiver from Avago
- Versatile interface through AMC edge connector, PCIe Gen2, GbE, LVDS and GTX

ROD Prototype in ATCA Format (LAPP, Annecy)

- parallel development of ATCA based ROD board
- 10 Gb Ethernet is tested with commercial ATCA computing blades
- first design of proprietary ATCA board is ready
- R&D goals:
  - understand IPMI and power control
  - develop control software framework
  - on board communication between FPGAs
- full ROD demonstrator will use ALTERA Stratix IV FPGA

ROD Demonstrator

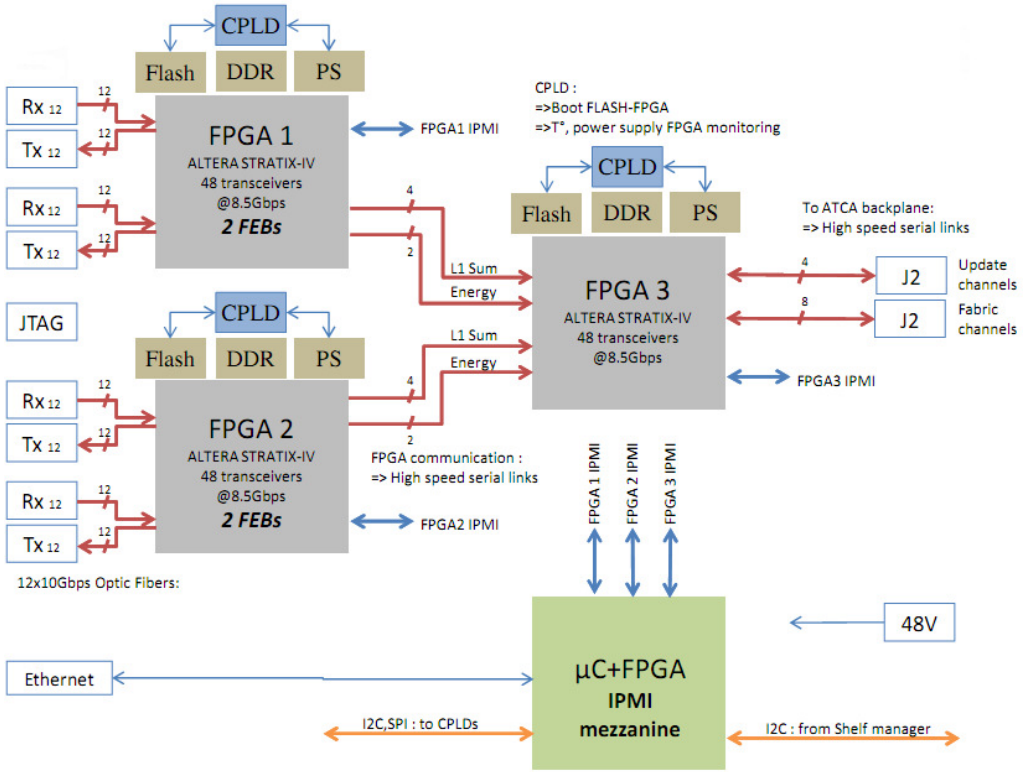

ROD and ROB communication in ATCA System (TU Dresden)

- read-out driver and read-out buffer test system:
- evaluate communication between ROD and ROB using 10 Gb Ethernet
  - server PC is a ROB prototype
    → planned to be a commercial ATCA blade in future
- hardware used:
  - Radisys ATCA crate with dual star backplane
  - 1 ATCA Sub-ROD equipped with Xilinx Virtex-5
  - 24 x 10 Gbps switch with XFP transceiver
  - optical 10 GbE connection to server PC
  - PC with dual Xeon 5550 CPU, 24 GB RAM, 4 Myricom 10 GbE dual port NICs

  - ROB server performance reached:
    - 79 Gbps @ 30% CPU load
  - ROD prototype goals reached:
    - 10 Gb XAUI interface implemented in Virtex-5 FPGA
- currently working on:
  - 10 Gb Ethernet MAC in FPGA
  - File system on PC-RAM via direct network access

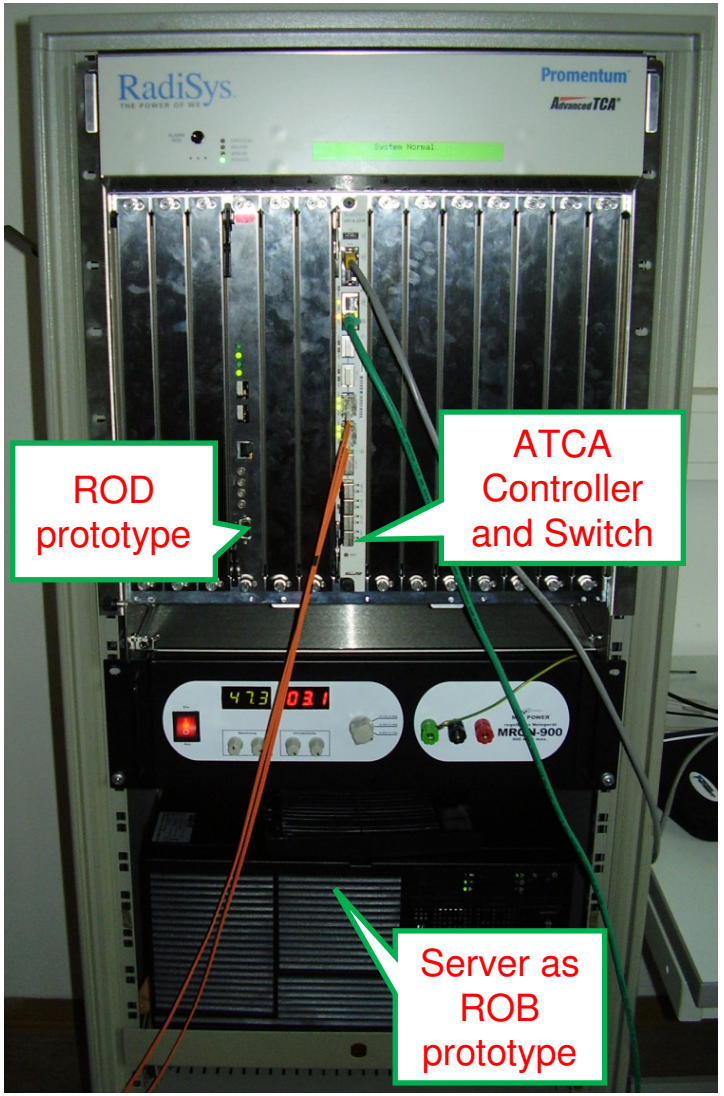

Development of Digital FIR Filter in ROD FPGA

- FIR filter prototype as core data processing unit

- incoming raw data
  - 2 bit gain factor and 12 bit for energy → gg xx EEEE EEEE EEEE

- outgoing data: energy sum

$$S = \sum_{i=1}^{5} a_i^{(gain)} \cdot (E_i^{(gain)} + Ped_i^{(gain)})$$

  - 2 bit range factor and 14 bit for energy sum → rrsS SSSS SSSS SSSS

- test hardware
  - FPGA: Xilinx Virtes 5 XC5VFX70T
  - testboard: Xilinx ML507
  - Communication PC ←→ FPGA:
    - 1 GbE, UDP/IP

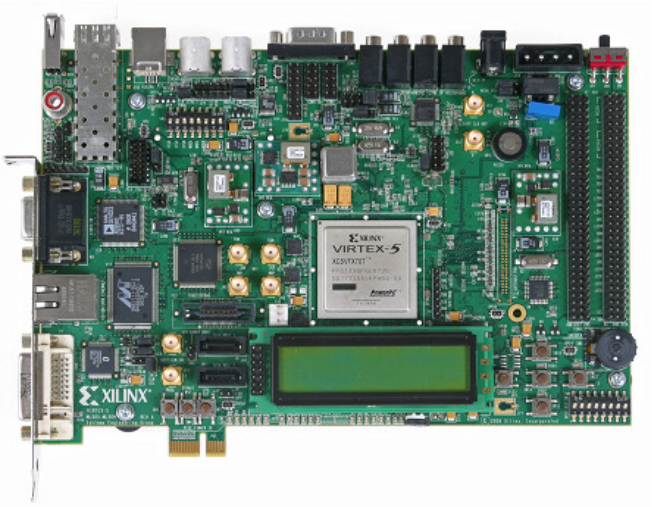

ROD FPGA – Digital Filter Prototype

- DSP modules used for energy sum calculation
- filter operates correctly up to 350 MHz
- filter internal latency is 3+2 clock cycles (due to DSP)
  → much less than 1 bunch-crossing (40 MHz)
- dynamic loading of calibration constants possible w/o FPGA firmware update

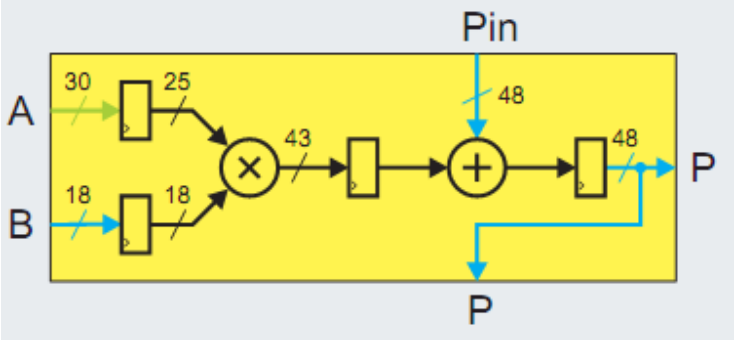

- Datasheet: DSPs can operate at up to 550 MHz → full speed not yet reached
  - could process 12 interleaved channels
- 128 DSP slices allow implementation of 25 filter units → 2 FEBs with 128 channels

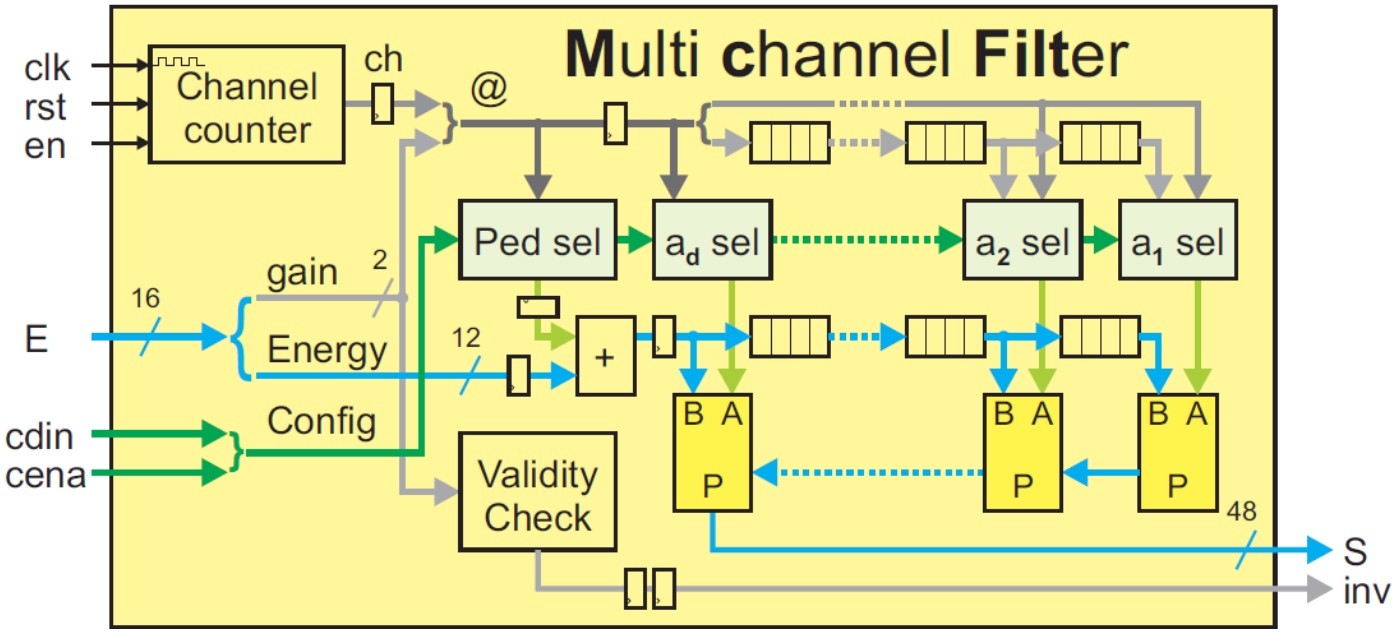

Summary and Outlook

- several reasons for updating / upgrading current LAr read-out environment
  ➔ a HL-LHC with advanced trigger logic with respect to granularity and flexibility

- new read-out scheme ➔ change in functionality of FEB and ROD/ROB systems

- ROD planned to have additional capabilities to generate L0/L1 digital sums, perform optimal filtering based on digitized raw data and digital buffer of raw data
  - ➔ first ROD prototypes in ATCA format
  - ➔ SERDES test results with Xilinx/Altera FPGAs: 75 Gbps ➔ 100 Gbps required
  - ➔ FIR energy filter: 4 channels at 350 MHz in FPGA ➔ need Nx128 at 500 MHz
  - ➔ evaluate 10 GbE communication to ROB

  ➔further studies:
  - ➔ event buffering and identification in FPGA
  - ➔ evaluate integration of all tasks in FPGA(s) on ATCA ROD
  - ➔ evaluate resource limits

Thank you for your attention!

LAr Calorimeter Signal

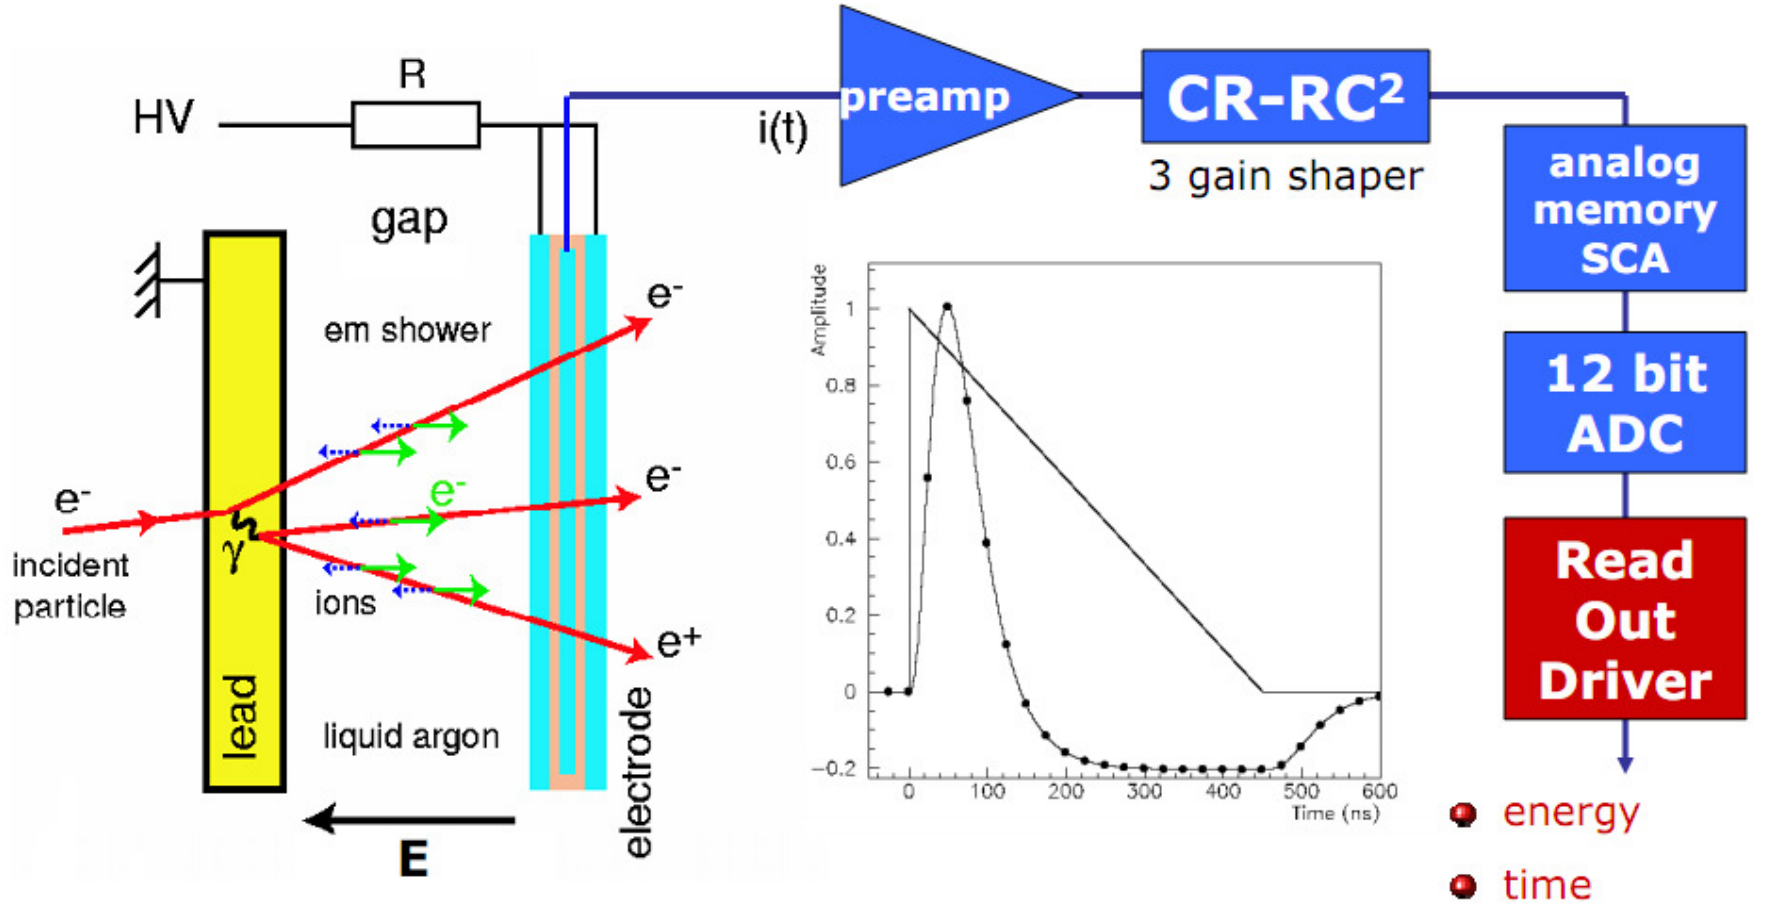

- LAr detector gives a triangular signal
- on-detector: pre-amplifying, shaping, analog buffering, analog-digital-conversion, transmission to ROD
- off-detector: optimal filtering applied to obtain optimum S/N ratio

Readout scheme for HL-LHC

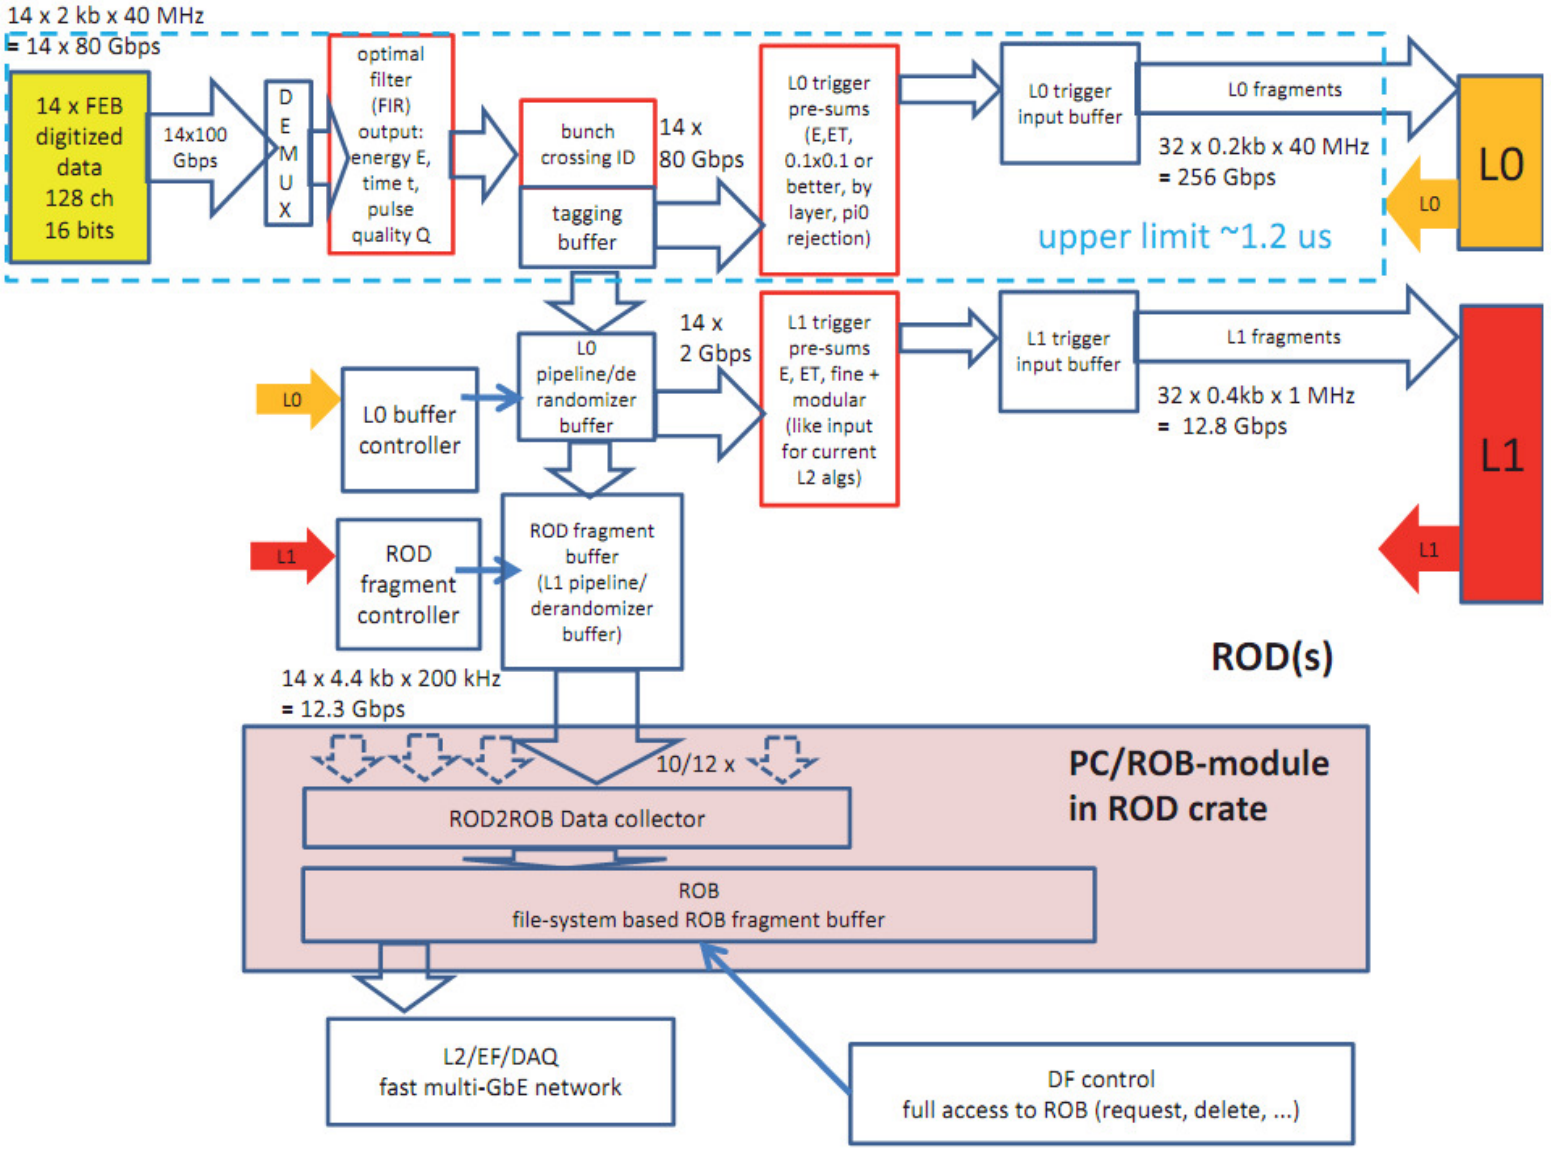

ROB development: status + Outlook

- read-out buffer: high-level-trigger and DAQ buffer (as today)
  - ROD is planned to be 10GbE interface to ROB

- new:
  - event data is stored as files directly in RAM of ROB (use proper file format)
  - Idea: use of standard components and protocols (ATCA, commercial computing boards, 10GbE, UDP, …)

- Status:
  - Dual Xeon 5550 PC (4 cores, 6.4GT, 24GB RAM)
  - 4 x Myricom 10GbE dual port NICs

- performance:
  - ~80Gbps @ 60% CPU (9.9Gbps, MTU9000)
  - HT on: ~30% CPU

- to be done:
  - data transfer from ROD FPGA to ROB PC
  - TDAQ software tests

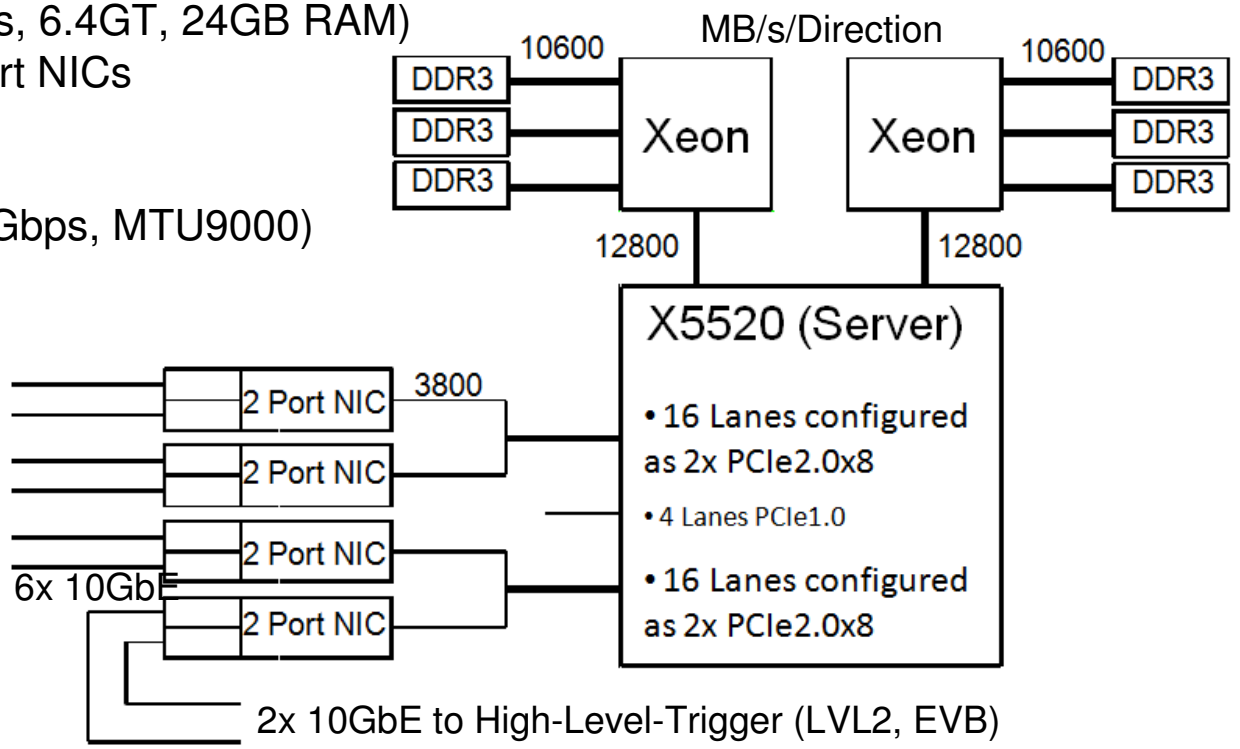

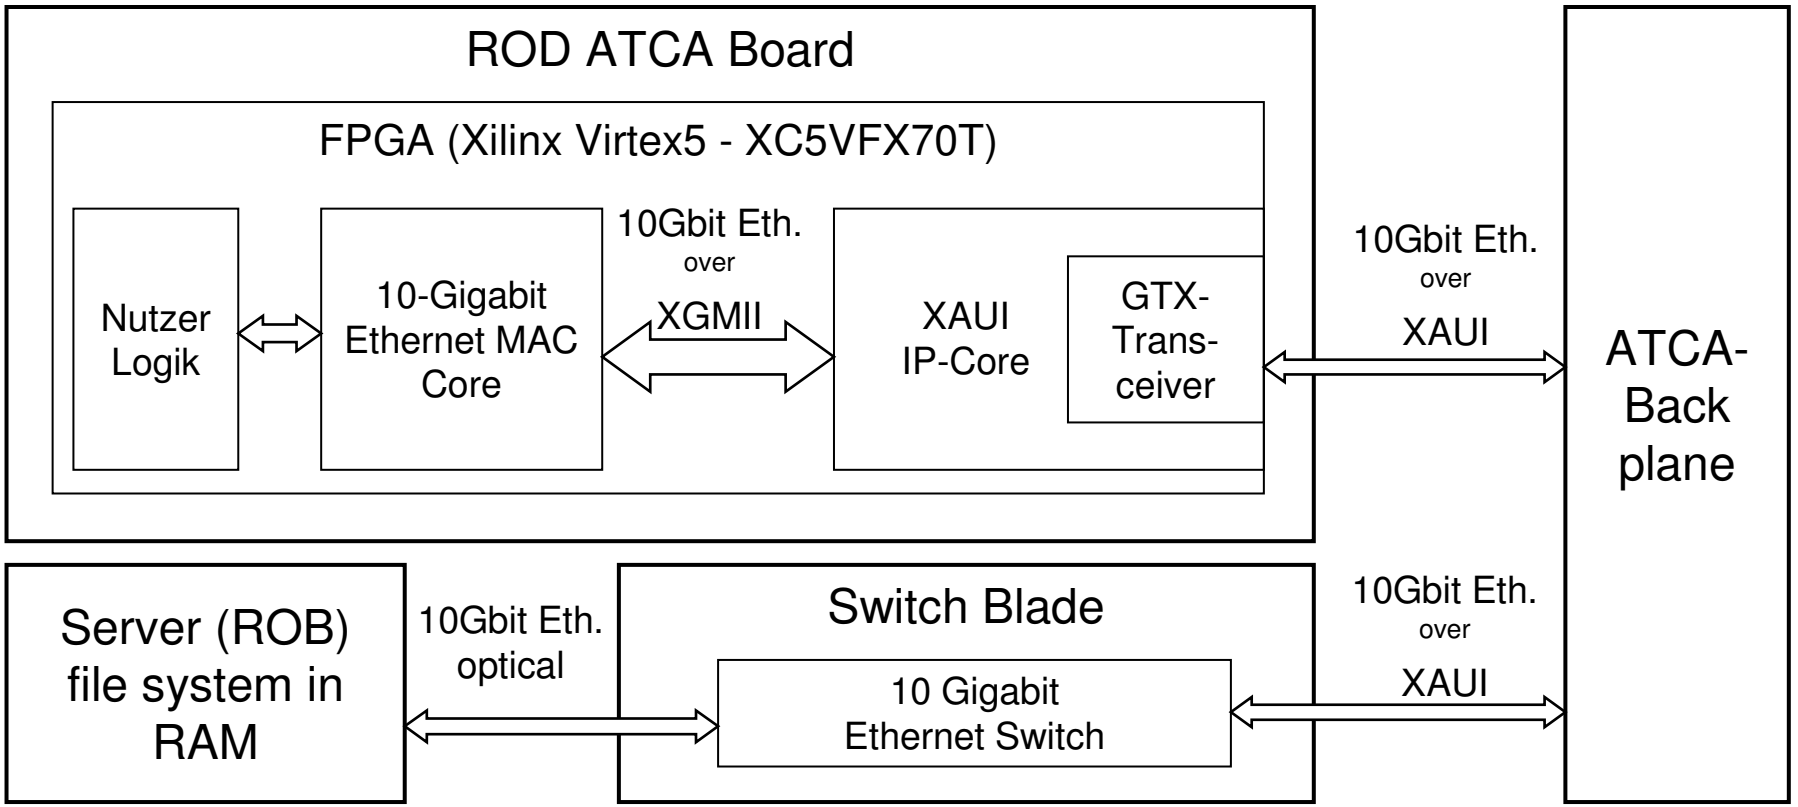
ROD ATCA Board
FPGA (Xilinx Virtex5 - XC5VFX70T)
Nutzer Logik
10-Gigabit Ethernet MAC Core
10Gbit Eth. over XGMII
XAUI IP-Core
GTX-Trans-ceiver
10Gbit Eth. over XAUI
ATCA-Back plane
Server (ROB) file system in RAM
10Gbit Eth. optical
Switch Blade
10 Gigabit Ethernet Switch
10Gbit Eth. over XAUI