



Hardware Acceleration
LAB



PANDA DAQ Summary 2022

Grzegorz Korcyl

Panda Collaboration Meeting 22/2

30.05.2022



SIM-DAQ

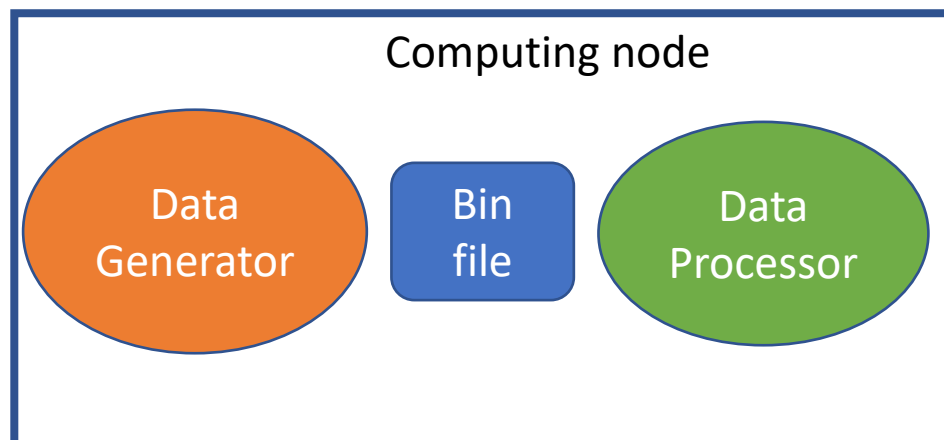
- System composed out of currently available components (simulations, dev. boards, platforms from other systems, ...)
- Subsystems imitated by computers/dummy boards with simulated data
 - Victor Rodin research on time-based simulations
- Framework for:
 - Key DC firmware components development
 - Protocols and data formats
 - Development of software processing pipeline
 - Evaluation of networking
 - Gathering of data flow characteristics

SIM-DAQ


- PandaROOT segmentation:
 - Data generator:
 - Produces binary data similar to what hardware would
 - Data organized according to subsystem electronics scheme and continuous readout
 - Data stored in files, to be loaded to emulation hardware
 - Data processor:
 - Software analysis pipeline for data selection
 - Benchmarking of the implementation

SIM-DAQ

- Step 1

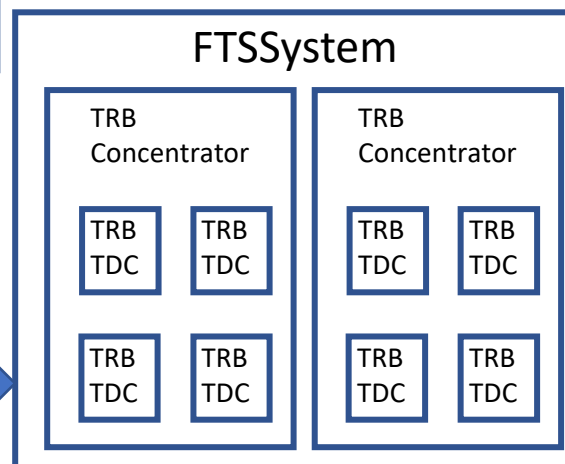


Binary HW data format for TRB-based subsystems



GetEntry()

FTSHit_

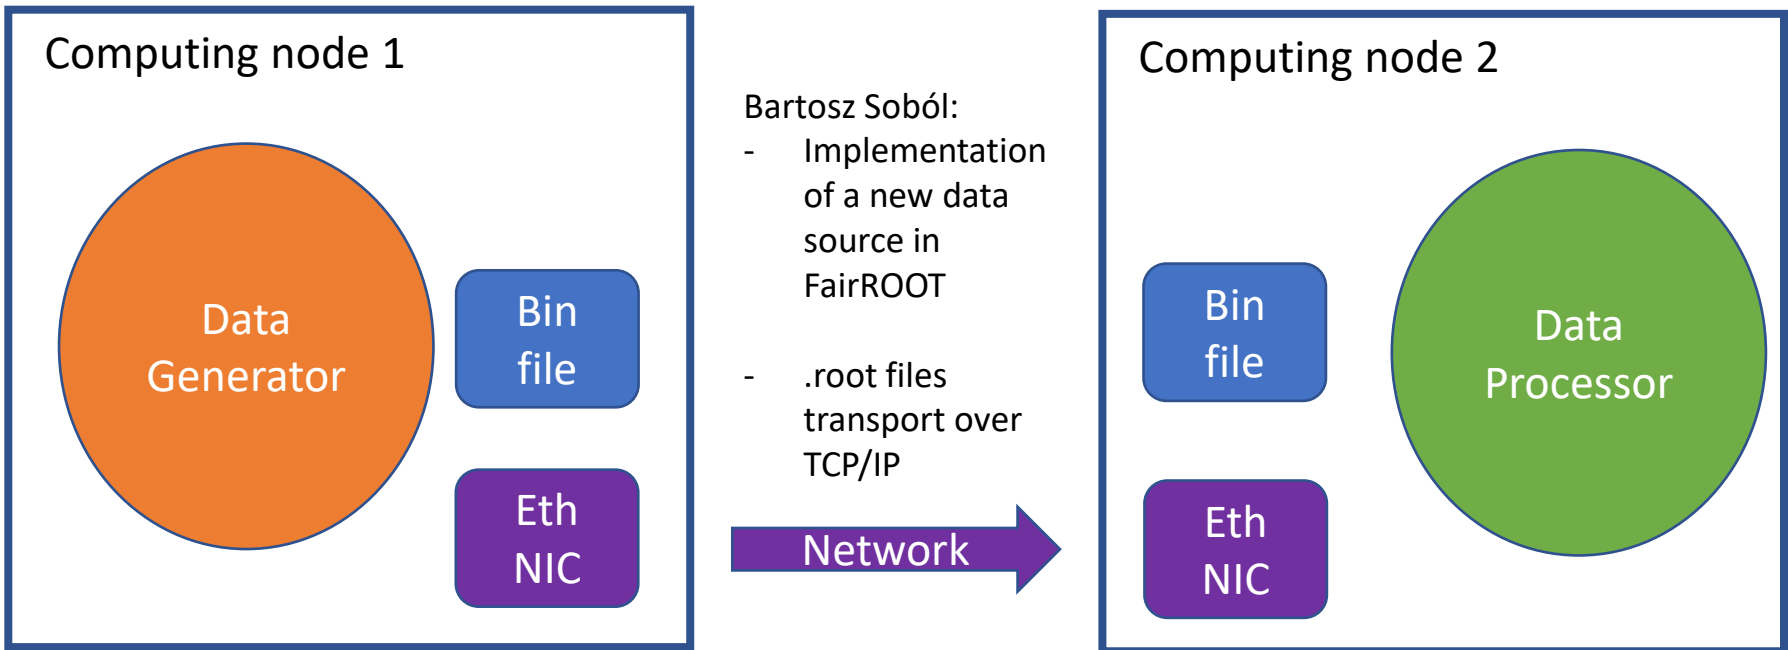


```

0000 ff ff ff ff ff ff 00 00
0010 01 1c 0b 00 00 00 ff 11
0020 ff ff c3 50 c3 50 01 08
0030 00 02 01 03 00 00 00 00
0040 00 02 06 01 00 02 06 02
    
```

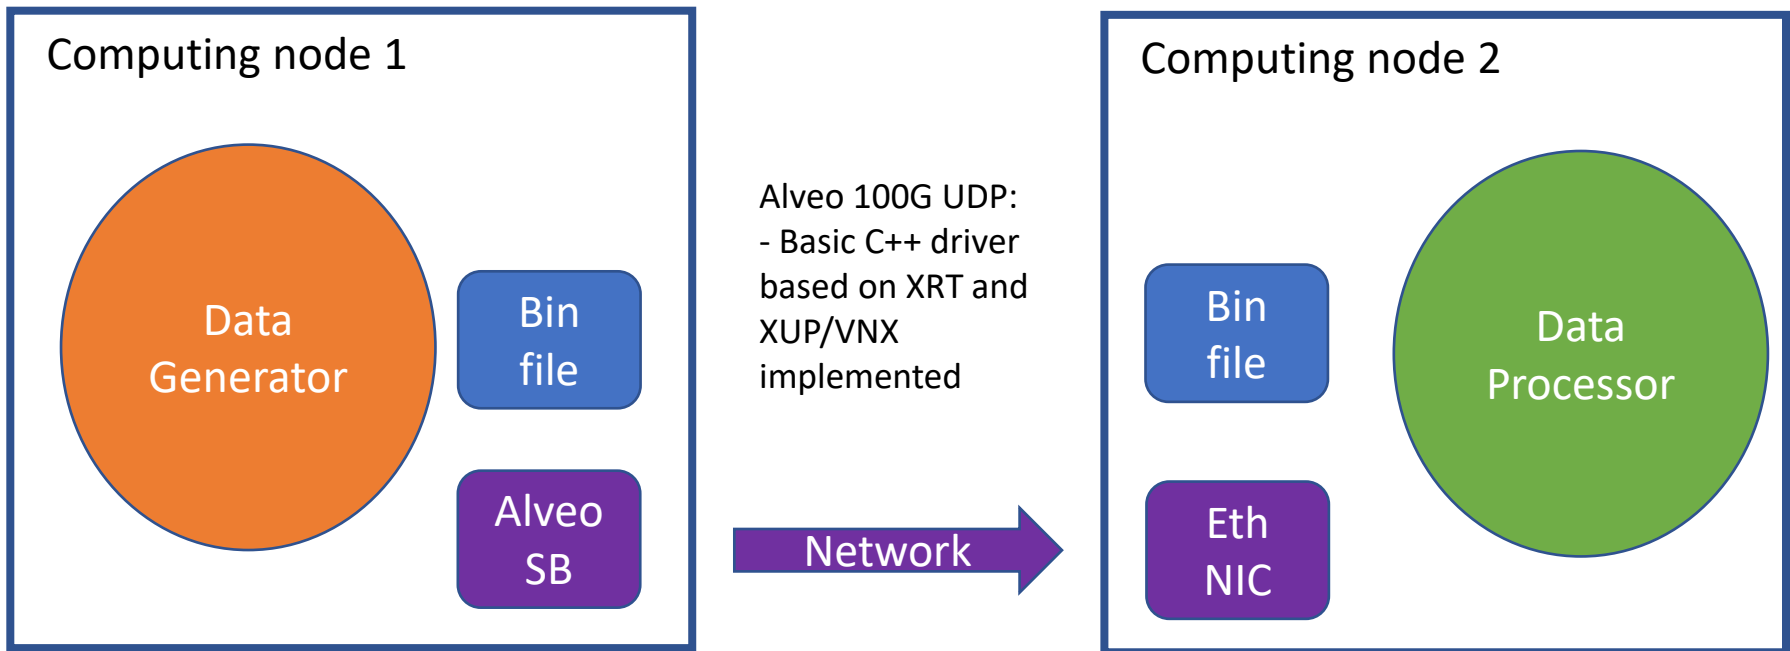
SIM-DAQ

- Step 2



SIM-DAQ

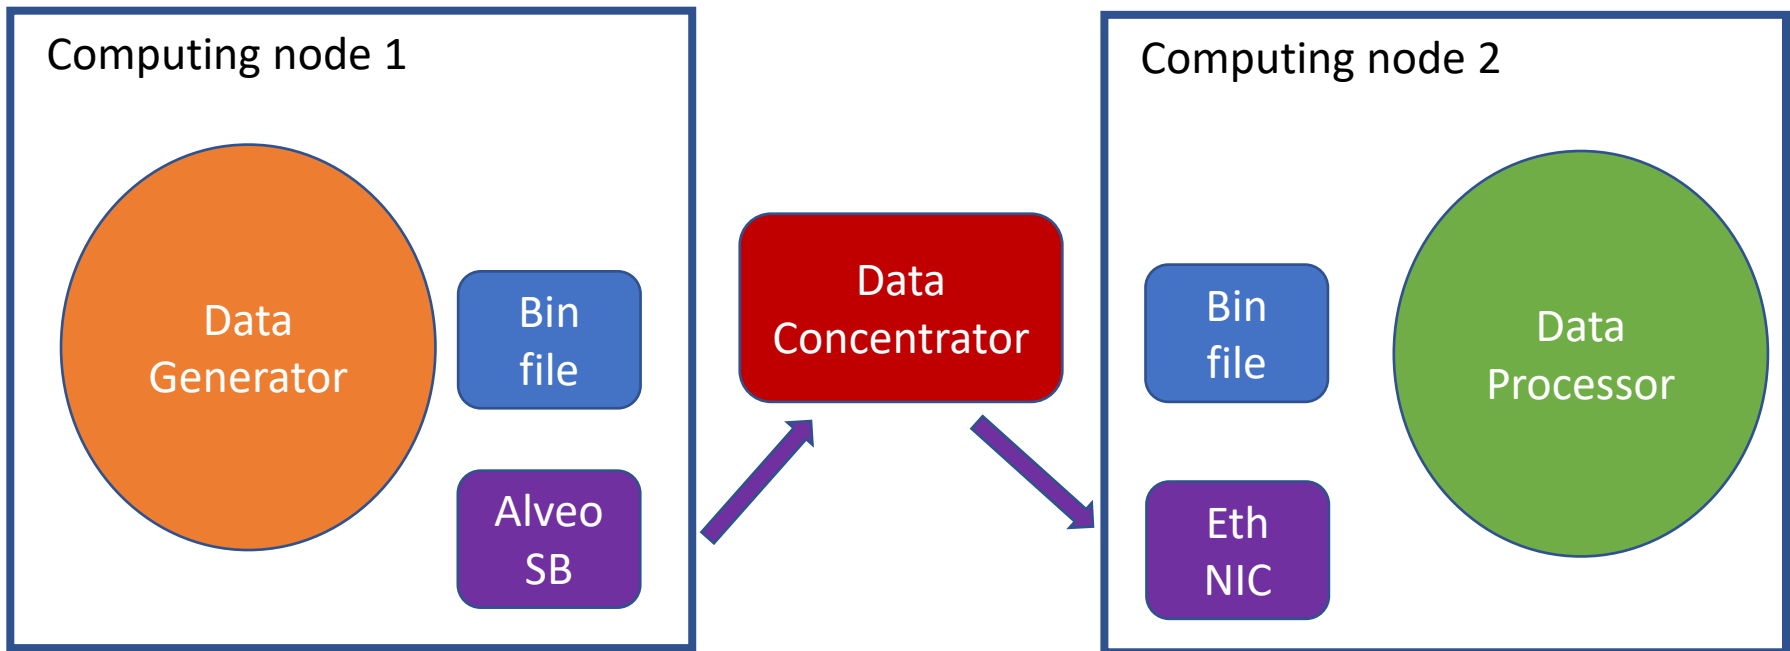
- Step 3



Alveo SB: preloaded data grouping into SuperBursts and blasting at SodaNet frequency

SIM-DAQ

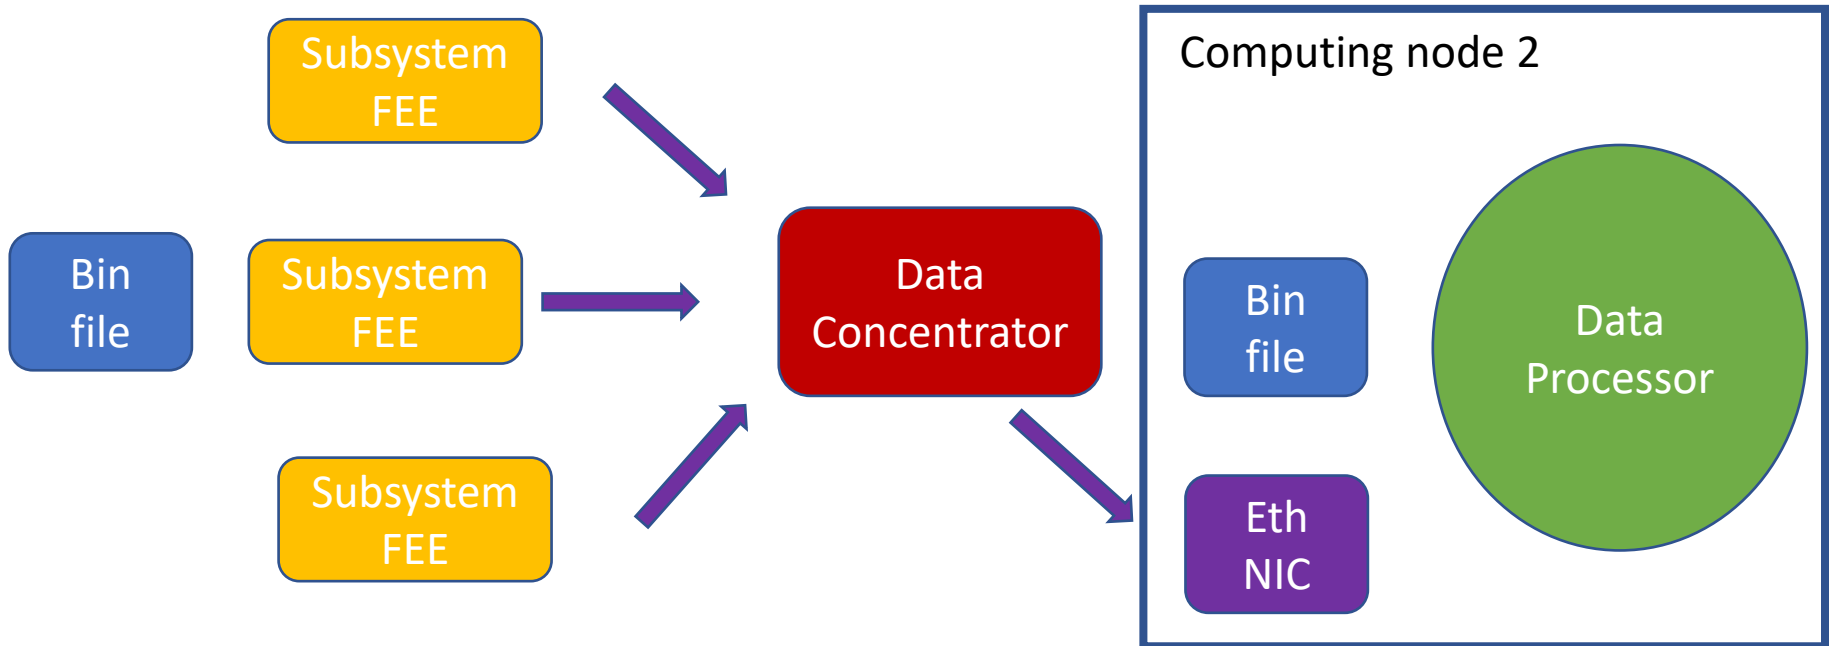
- Step 4



Data concentrator: firmware evaluation, framework for subsystems preprocessing

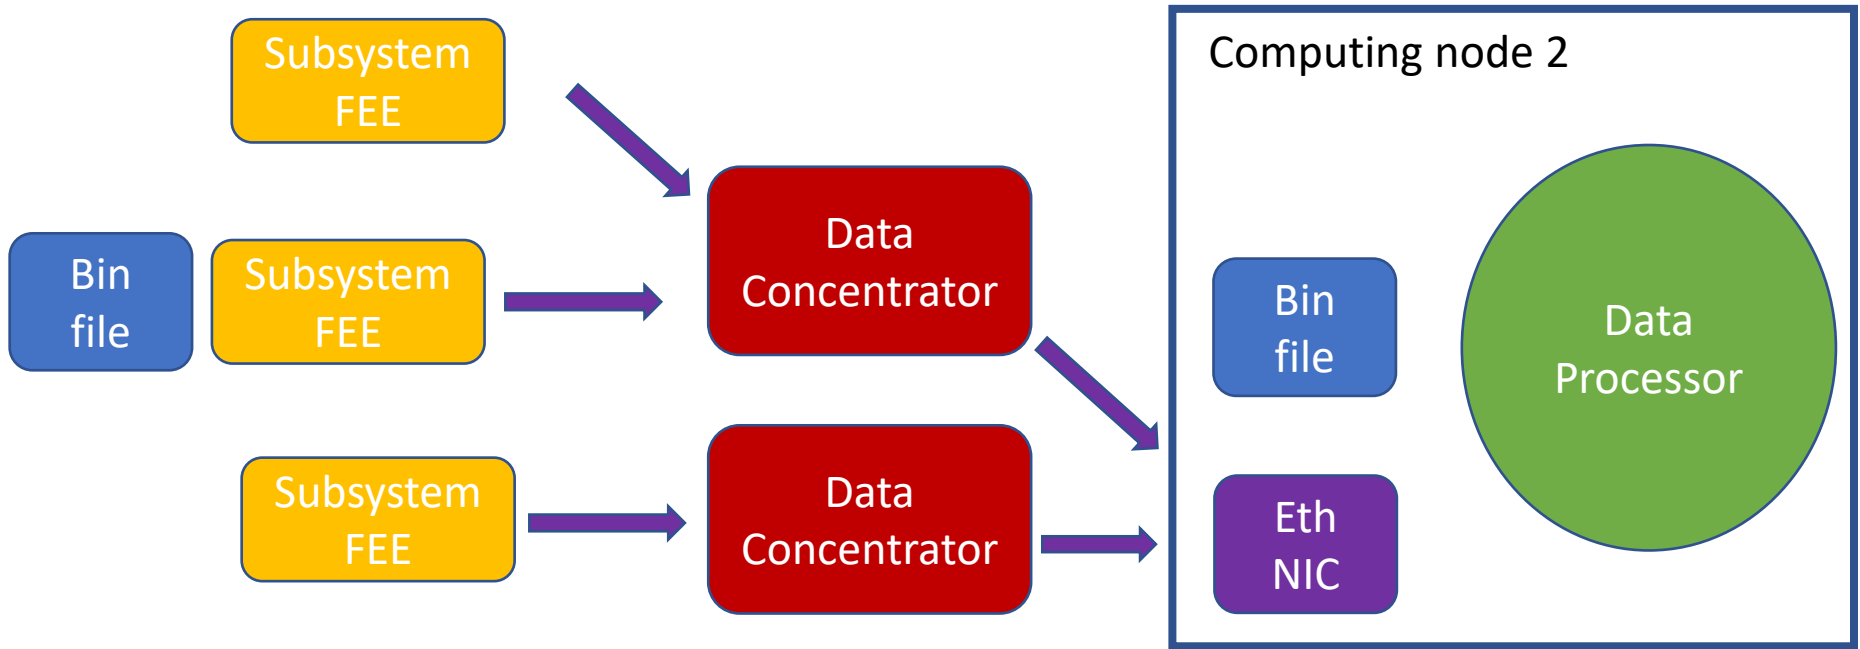
SIM-DAQ

- Step 5



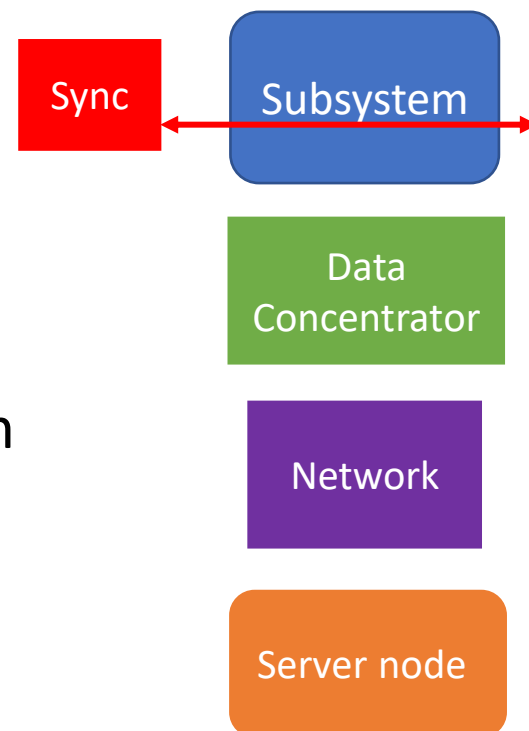
SIM-DAQ

- Step 6



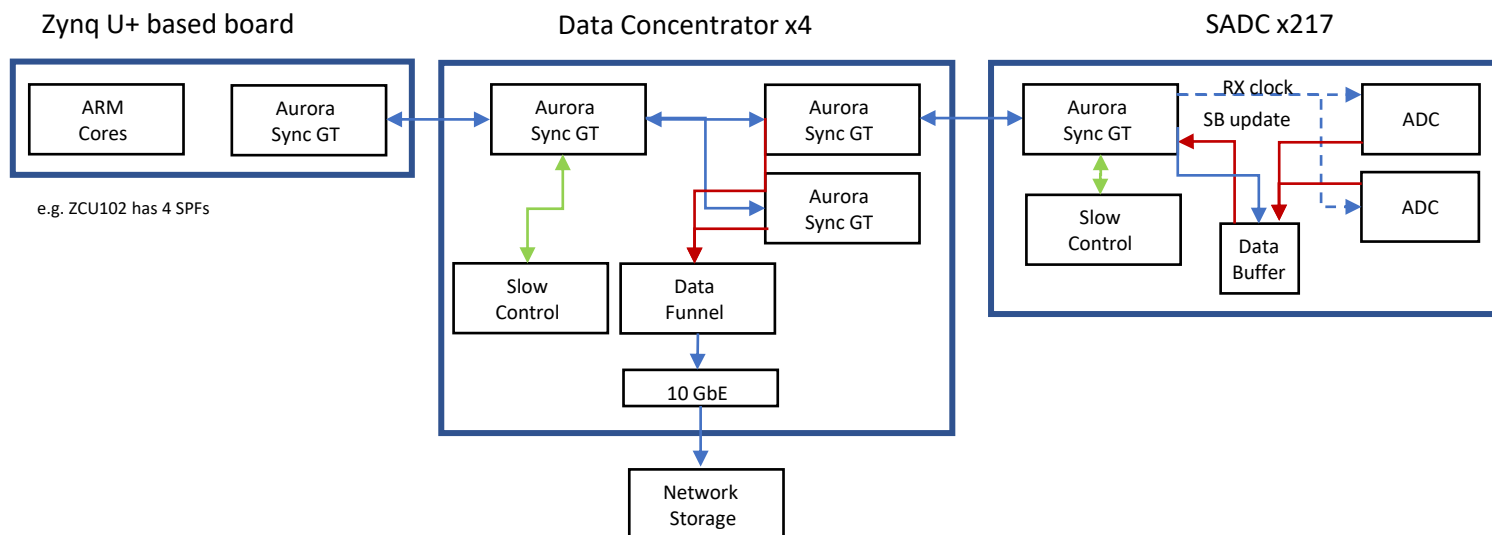
DAQ-0

- Actual subsystem setup
 - A single segment or more complex setup
- Synchronization
- Actual Data Concentrator hardware platform
- Framework for:
 - Subsystem integration with the DAQ
 - Development of subsystem preprocessing on DC
 - Development of control and monitoring procedures
 - DAQ performance and scalability evaluation
 - Available for beamtests of subsystems



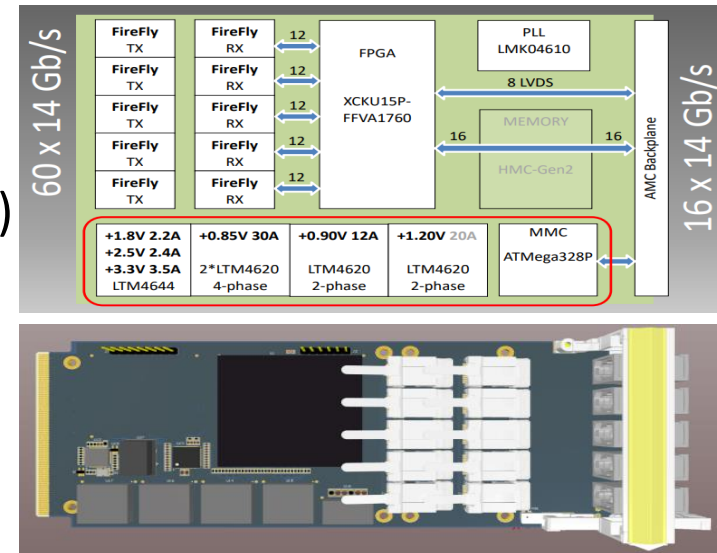
DAQ-0 candidate EMC

- Beamtime in Q1 2023 – Lukas Linzen presentation
- Preliminary HW/FW/SW setup ready by the end of 2022
 - Some temporary solutions to secure data taking




Data Concentrator

- 2 units produced and equipped
 - 1 u. operational, 1 u. power supply under investigation
- FPGA programmable
- Module Management Controller operational
- Initial firmware (P. Marciniewski & M. Caselle)
 - PLL HDL controller – clock source for MGTs
 - Evaluation of MGT channels
- FPGAs purchased for more units
- Firefly transceivers sets for 2 units
 - Purchase required for additional 2 units



P. Marciniewski

DC Firmware



Particular DC firmware has to be tailored to a given subsystem:

- In/Out link configurations (resource balance)
- Preprocessing algorithms
- Supporting services

Data receivers region:

- Logic common to all DCs (configurable num. of links)
- Common link type and protocol
- Unified interface to Preprocessing region

Preprocessing region:

- Fixed, unified data in and out interfaces
- Region available to Subsystem devs.

Data transmitters region:

- Logic common to all DCs (configurable num. of links)
- Unified interface to Preprocessing region
- Commercial network interface

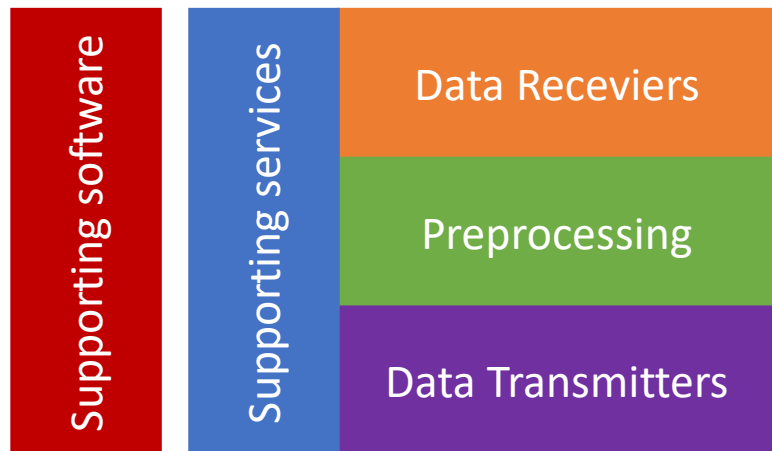

Supporting services:

- Control and monitoring
- Synchronization
- ATCA management
- ...

Supporting software:

- Board management
- Data QA
- ...

DC Protocols

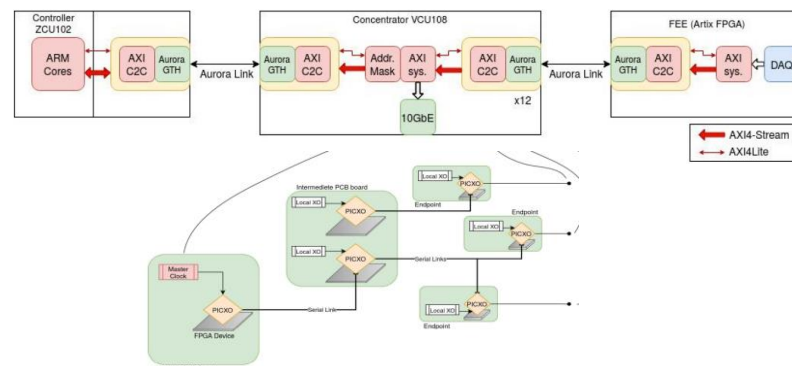


10 Gb Ethernet



Data Receivers:

- All Subsys. end up with Xilinx FPGAs
- Wide range of link speeds
- Aurora IPs are well tested and easy to use
- Lightweight
- Can transport AXI transactions
- Can be synchronous

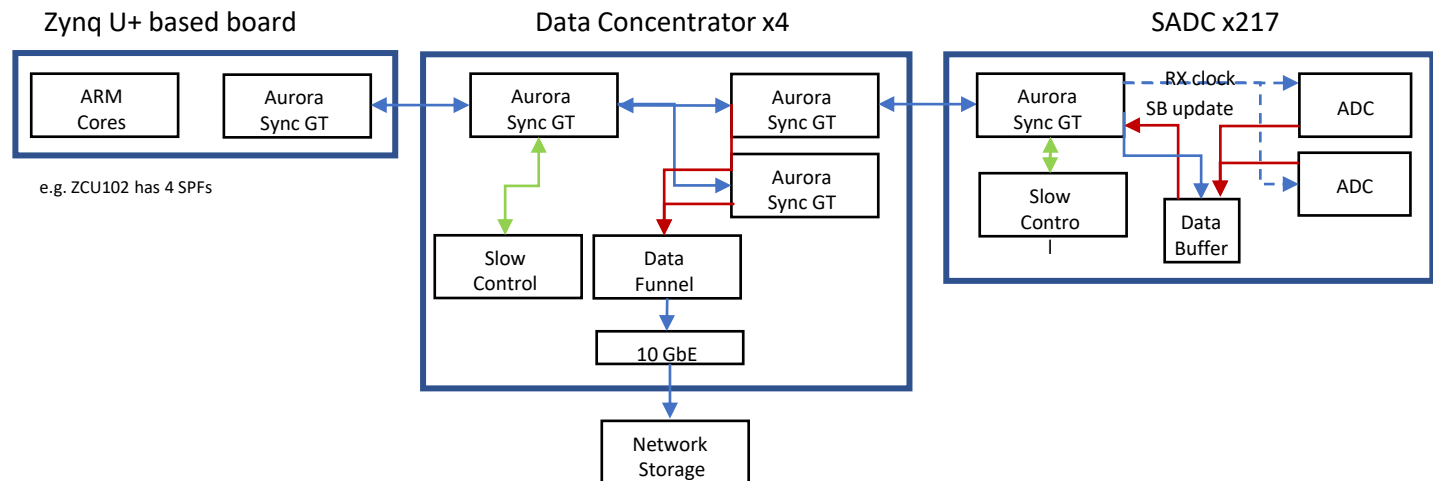


Data Transmitters:

- Common networking standard
- Easy to set up a simple UDP blaster for basic operation and evaluation of the readout chains
- Works with DABC and GO4

DC Firmware

- Aurora Sync 5Gbps GT components available
- Master clock propagated to ADCs
- TX synchronous , RX asynchronous
- Slow control from Linux on ARM through AXI components
- Requires offline time offset correction
- Critical points:
 - SADCs read-out one by one by the Data Funnel (54 to one) – can take long time, multiple GbE output
 - Resources consumption by over 50 data links



DC Workgroup

- Paweł Marciniewski, Michele Caselle, Olena Manzhura, Lukas Linzen, Grzegorz Korcyl
- Source code and documentation repository
- Set of ready to use FW components
- Regular meetings, 2-weeks
- Additional ad-hoc trainings

- Everyone is welcome to join in!

The screenshot shows the GitHub repository page for the **DataConcentratorGroup**. The group name is displayed with a shield icon, and the group ID is 310. Below the group name, there are tabs for **Subgroups and projects**, **Shared projects**, and **Archived projects**. The **Subgroups and projects** tab is active, showing a list of subgroups and projects:

- dc_doc**: Documentation and board files for the Data Concentrator projects
- dc_soc**: SoC-based Data Concentrator project
- dc_fpga**: FPGA-based Data Concentrator project
- dc_fw_ips**: Collection of IPs required to build the complete Data Concentrator

On the right side of the screenshot, a file browser view is shown with a table of files:

Name
PICXO_FRACXO
aurora_rx_sync_slide_gth
aurora_rx_sync_slide_gtp
aurora_status_reset
axi_address_mask
axis_cmd_gen
djpet_trigger_receiver
gth_rx_sync_slide
gtp_rx_sync_slide
transit_fifo
usr_interface_to_axis
README.md
build_project.sh

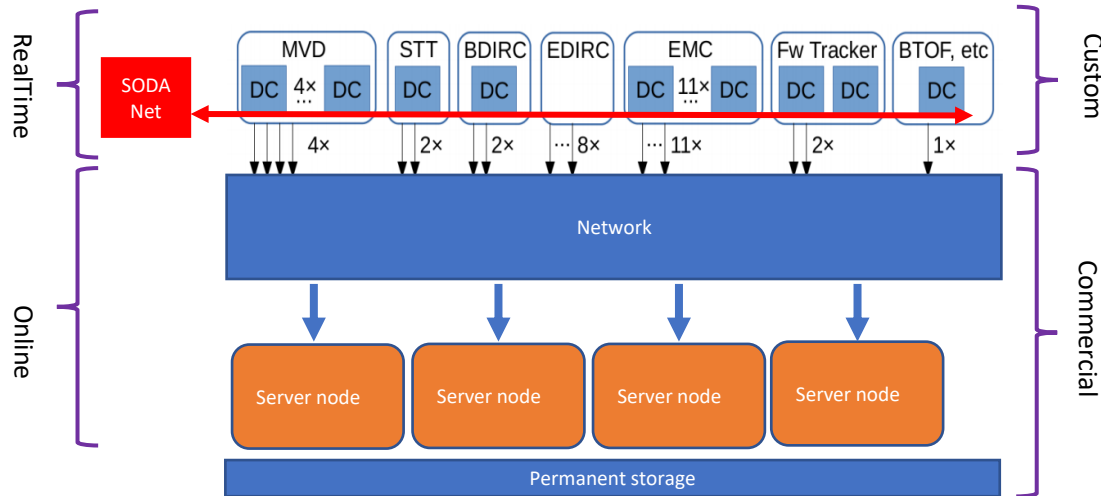
Summary

- Development of SIM-DAQ initiated
 - PandaRoot segmentation
 - Subsystem-specific binary HW data representation
 - Basic networking (TCP file transport, Alveo UDP data blaster)
- Data Concentrator firmware development initiated
 - HW evaluation ongoing
 - Set of firmware components for DAQ-0 and EMC beamtime defined
 - Workgroup established
 - Detailed workplan to be defined

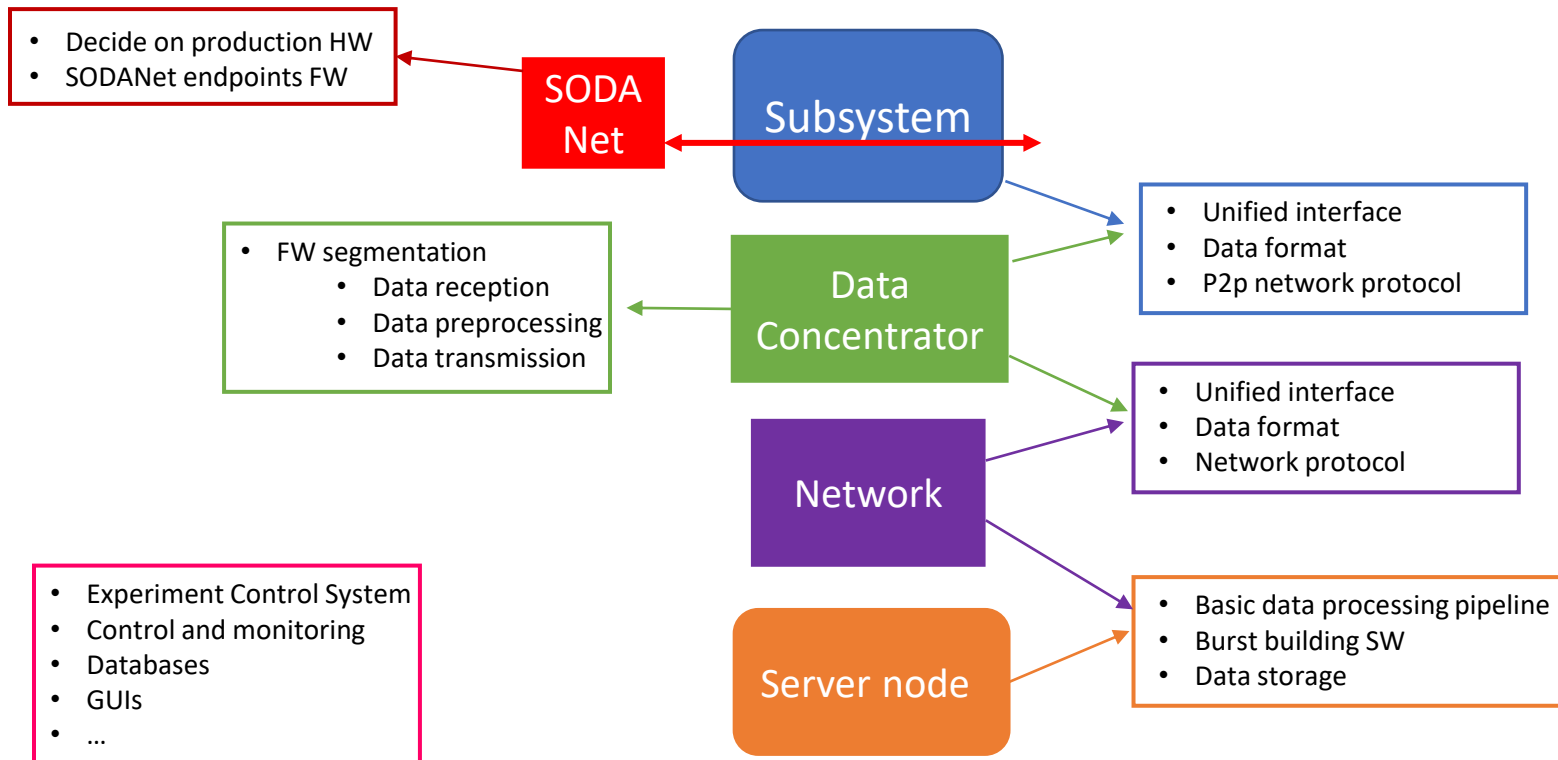
System development strategy

3. DAQ-1

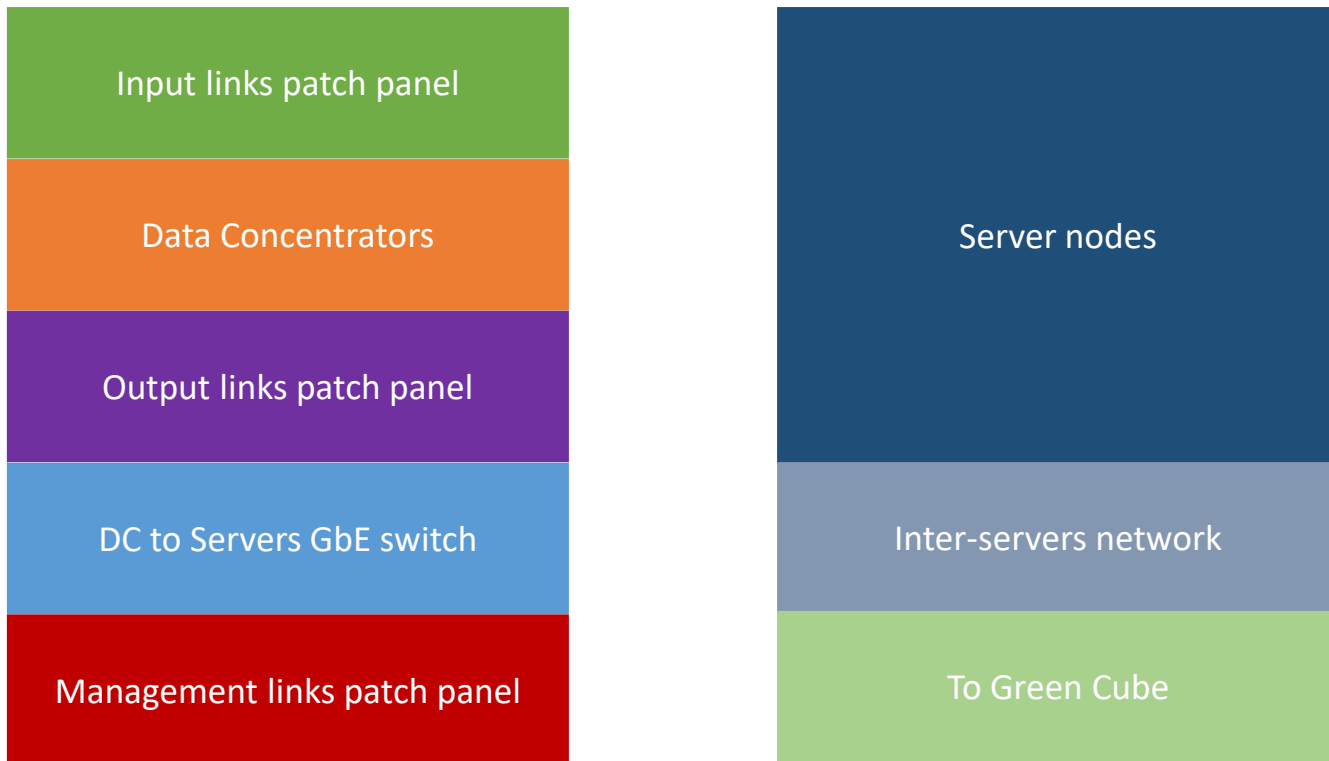
- DAQ-0 scaled up to accomodate all subsystems



Towards DAQ-0



System composition





Subsystems Overview

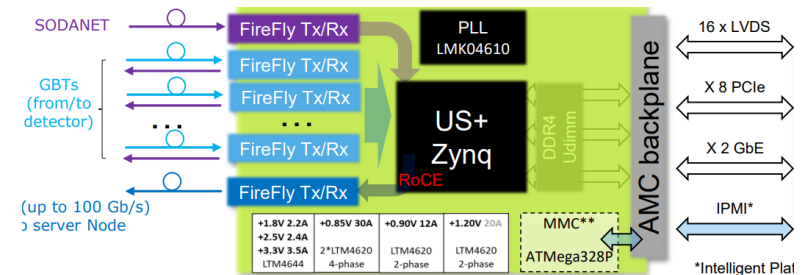
FT			
PASTTREC (352)	TRB5SC (88)	Crate Master (10)	DC (1-?)
STT			
PASTTREC (270)	TRB5SC (68)	Crate Master (8)	DC (1-?)
MVD			
TOAST ()	MDC (296)		DC (5-?)
Endcap Disc DIRC			
TOFPET2	RDB (96)		DC (2-?)
Barrel DIRC			
DIRICH	TRB3/TRB5SC		DC (?)
EMC			
SADC	Crate Controller (41)		DC (1-?)
HDA	LVDS DC (40)		DC (1-?)
Lumi			
MuPIX	LVDS DC ()		DC (?)

Data Concentrators

- ATCA – AMC platform
 - 60x bidirectional optical links (FireFly)
 - Backplane I/O for management and board-board communication
- Processing unit:
 - **Xilinx Kintex: FPGA-DC**
 - Regular FPGA resources
 - **Xilinx Zynq MPSoC: SoC-DC**
 - Regular FPGA resources
 - Quad-core ARM Cortex A53
 - Dual-core ARM Cortex R5F
 - Mali 400MP2 GPU



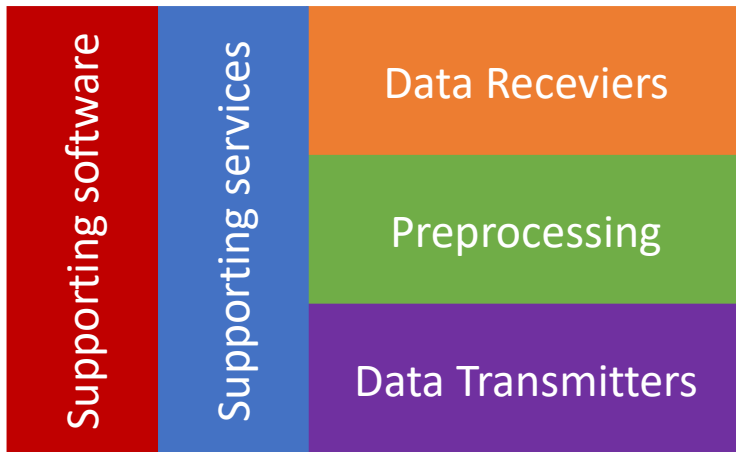
P. Marciniwski



M. Caselle

*Intelligent Plat

DC Firmware



Particular DC firmware has to be tailored to a given subsystem:

- In/Out link configurations (resource balance)
- Preprocessing algorithms
- Supporting services

Data receivers region:

- Logic common to all DCs (configurable num. of links)
- Common link type and protocol
- Unified interface to Preprocessing region

Preprocessing region:

- Fixed, unified data in and out interfaces
- Region available to Subsystem devs.

Data transmitters region:

- Logic common to all DCs (configurable num. of links)
- Unified interface to Preprocessing region
- Commercial network interface



Supporting services:

- Control and monitoring
- SODANet
- ATCA management
- ...

Supporting software:

- Board management
- Data QA
- ...

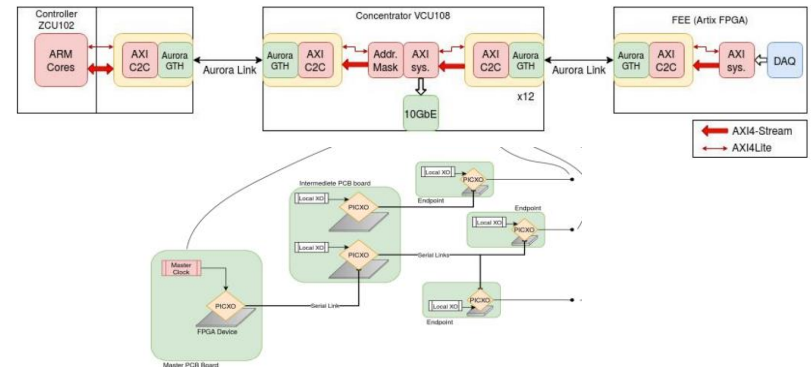
DC Protocols



Data Receivers:

- All Subsys. end up with Xilinx FPGAs
- Wide range of link speeds
- Aurora IPs are well tested and easy to use
- Lightweight
- Can transport AXI transactions
- Can be synchronous

M. Bakalarek talk



Data Transmitters:

- Common networking standard
- Easy to set up a simple UDP blaster for basic operation and evaluation of the readout chains

Data Formats

DC Output Data-format



- DC can start transmitting FEE data once it is available
(without waiting till the end of a super-burst)
- If no data are available –
DC sends an empty package at the end of the Super-burst

Data-package

31	16	15	0
last-packet flag; packet number		data size in bytes	
Not used (same as HADES)		Not used (same as HADES)	
Status and error		System ID	
Super-burst number			
Data			

All FEE modules should issue “empty packet” in case no data is measured – assures readout integrity

Data formats

Data format for transport between FEE/Data Concentrators and CN layer

- Proposal: **define universal packet format**, independent of detector subsystem
 - Advantage: we can **use the same set of IP cores** to handle data streams independent of detector and DAQ layer
 - Data is streamed and organised in packets
 - Push architecture, but with support for back pressure
- **Data Origin Definitions:**
 - “detector”: Panda Subsystem ID (EMC, STT, MVD etc.)
 - reserve one byte for unique detector definition
 - “module” : section ID of a detector (EMC Forward endcap, STT stereo layer, MVD pixels etc.)
 - reserve one byte for unique module definition
 - “location”: unique identifier for the geographical location of a hit within a module (could be STT wire number, MVD pixel ID etc.)
 - reserve 4 bytes (need addressing space for MVD)

Data formats

Universal Packet Format (UPF)

- Each packet contains a **packet header** with the following information (distributed in part by SODANET to FEE and/or set by slow control in FEE for local runs):
 - ID for “experiment” number (should be unique over the lifetime of Panda) (2 Bytes ?) (set by run control)
 - Run number (is reset at the start of a new experiment) 2 Bytes (set by run control)
 - Run type (physics, calibration, test/debug, local run, etc.) 1 Byte (set by run control)
 - Event selection identifier or has code pointing to data base(4 bytes ?)
 - Superburst ID: 4 Byte (to be distributed by SODANET) (use also as last packet flag)
 - Payload size (in bytes): 3 Bytes, payload item size in bytes (1 byte) (detector specific)
 - **Payload header** : Number of items in payload
 - **Payload items:**
 - Time stamp
 - Data Origin
 - Data value(s) (detector-specific)
 - **Payload trailer:** repeat payload size, add checksum (CRC)
- **Packet trailer**
 - repeat payload size (for redundancy, debugging and recovery)
 - CRC (packet checksum)

DC to Servers network

- **Maintain subsystems isolation**

- Subsystem -> DC -> dedicated server
- Server instances optimized for particular subsystem data characteristics
- Superburst/burst building done manually
- Communication between nodes required
- Distributed processing


- **Superburst round-robin**

- Complete superburst -> particular server
- Superburst building done by network „for free”
- No communication between nodes required
- Possible bottlenecks (all DCs -> one server at a time)

Processing pipeline

Data Flow

Subline



Mitglied der Helmholtz-Gemeinschaft

Reconstruction Tasks

Subline



1. Course event building and t0 determination in fast detectors (MVD, TOF, EMC)
2. Combination of different detectors into one event candidate
3. Tracking
4. PID/EMC combined with tracking data
5. Precise t0 determination and event building
6. Event filter

Mitglied der Helmholtz-Gemeinschaft

Selle 17

1. Subsystem independent processing

- Coarse event building / clustering
- Calibrations, mappings, ...
- FT, STT tracking with coarse T0
- PID on EMC

2. Burst building

- Combination of tracks and PID
- Precise T0 determination
- Complete event building

3. Event filter

- Software trigger
- Event tagging
- Event rejection

Processing pipeline

1. Subsystem independent processing

- Subsystem independent
- Processing of time-ordered stream
- Course event building – clustering
- Calibrations, mappings, ...

- Clustering efficient on FPGA
 - Preprocessing region on DC
 - FPGA accelerator in server node
 - Direct network connection

- Coarse tracking on GPU
 - Intermediate storage host DDR, CPU interaction required
 - RoCE from Data Concentrators

- Product subevents stored in host DDRs


1. **Subsystem independent processing**
 - Coarse event building / clustering
 - Calibrations, mappings, ...
 - FT, STT tracking with coarse TO
 - PID on EMC

2. **Burst building**
 - Combination of tracks and PID
 - Precise TO determination
 - Complete event building

3. **Event filter**
 - Software trigger
 - Event tagging
 - Event rejection



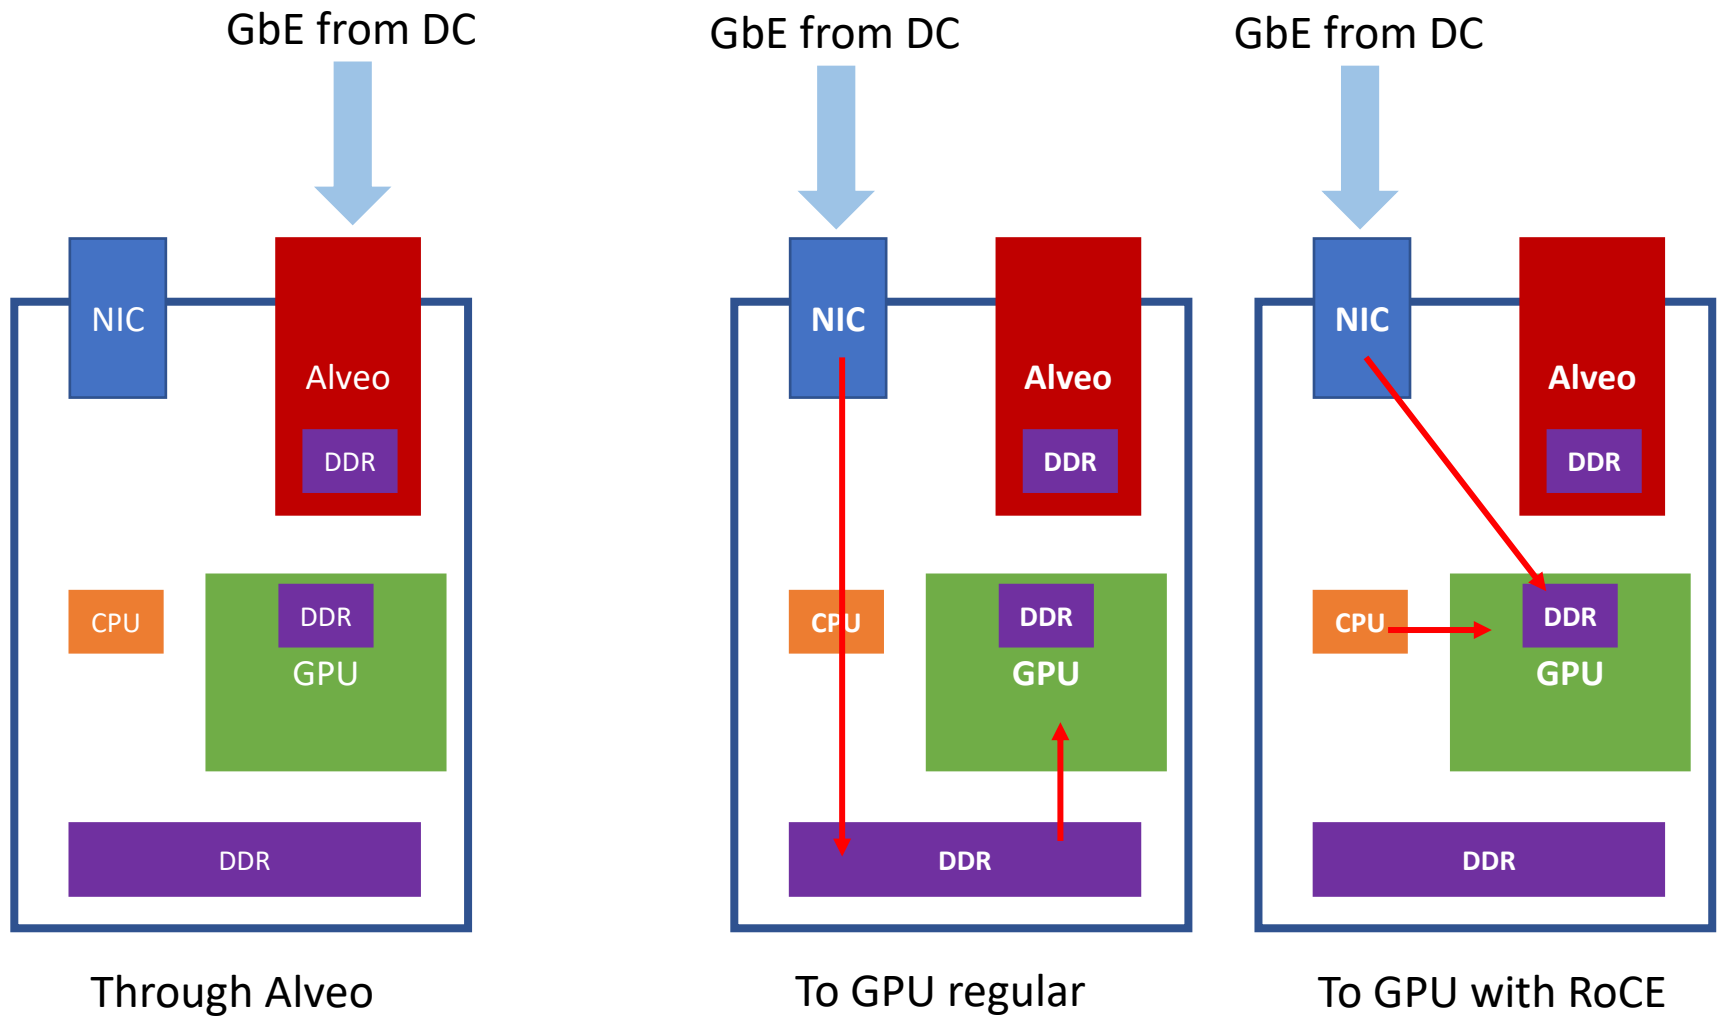
VNx (XUP)



- ▶ CMAC kernel
 - Card and interface specific configuration
 - Built in layers
 - Easily adaptable
- ▶ Network Layer
 - UDP as transport protocol
 - Support for multiple interfaces and cards
 - Single channel application interface
 - Currently U50, U250 and U280

https://github.com/Xilinx/xup_vitis_network_example

Processing pipeline



Processing pipeline

2. Burst Building

- Combination of tracks and PID
- Precise T0 determination
- Complete event building

- Mostly CPU intensive memory scans and data matching
- Inter-server communication required
- Fast interconnect each-each

- Burst building processes on subsystems processing nodes
- Round-robin job distribution

1. **Subsystem independent processing**
 - Coarse event building / clustering
 - Calibrations, mappings, ...
 - FT, STT tracking with coarse T0
 - PID on EMC

2. **Burst building**
 - Combination of tracks and PID
 - Precise T0 determination
 - Complete event building

3. **Event filter**
 - Software trigger
 - Event tagging
 - Event rejection

Processing pipeline

3. Event filter

- Software trigger for event tagging or filtering
- Promising implementation as a set of Neural Networks (P. Jiang)
- Very fast inference ~1.5M candidates per second on a single RTX3090
- Transmission to Green Cube for storage or further processing

1. **Subsystem independent processing**
 - Coarse event building / clustering
 - Calibrations, mappings, ...
 - FT, STT tracking with coarse TO
 - PID on EMC
2. **Burst building**
 - Combination of tracks and PID
 - Precise TO determination
 - Complete event building
3. **Event filter**
 - Software trigger
 - Event tagging
 - Event rejection

Towards DAQ-0

- Dummy boards with data from simulations
- Data concentrator:
 - Target DC board or any development board with Xilinx and transceivers
- Server node:
 - Any modern PC
- Development work groups:
 - Data Concentrator firmware
 - Software processing pipeline

Summary

- Basic concepts and elements are defined, require to be more and more detailed
- All input is appreciated:
 - Updates on subsystems data rates
 - Updates on numbers of modules and links
 - Potential preprocessing methods/algorithm on levels: DC and Servers