7.5a: TDAQ Overview

Yun-Tsung Lai, Alex Keshavarzi

on behalf of 7.5a group

KEK IPNS, the University of Manchester

ytlai@post.kek.jp

3rd DRD7 workshop

10th Sep., 2024

7.5a TDAQ Overview

- The 7.5a project is for:
  - TDAQ backend in HEP experiments
  - Hardware device: keep pace with COTs technologies (CPU, GPU, FPGA).
  - Software/firmware development: generic algorithms, workflows, or design tools for TDAQ real-time processing.
  - Technique sharing: Open-access, repository-hosted infrastructure for these commonly used tools and algorithms.
  - In each experiment, usually not so many TDAQ experts, especially their expertise are separated for various techniques: Collaboration crossing experiments can help to collect the technical knowledge complementary to each other, and to form a comprehensive database for future activity.

- COTs:
  - New technology with better performance.
  - But we (TDAQ people in exp. HEP) usually requires customized design for experimental purpose.
    - In some senses, more effort in R&D process.
  - Benefit: IP, libraries, and package directly from vendors provides more convenience in manual software/firmware design. Technical supports from industry.

7.5a members

- 14 sub-groups in 7.5a

Activities within 7.5a

- In principle, everyone in 7.5a is working on **Trigger**.
  - The development involves
    - **Algorithm/Physics**
    - **Software/Firmware framework**
    - **Hardware device**.

Relatively coupled

Physics/ Algorithm	**Software/ Firmware**	**Hardware**
• Tracking • Calo clustering • Topology • Anomaly detection •	• Software coding • ML framework • HDL • HLS • ML inference •	• CPU • GPU • FPGA

- Other than trigger, there are also
  - DAQ: readout, FEE, etc
  - Generic development for fundamental purpose: **firmware optimization/maintenance**, etc.

Activities within 7.5a (cont'd)

- Considering the trend, two major types of **Trigger** systems for real-time processing/filter:
  - A very personal point of view, as the boundary in between is getting vague nowadays.

- **Hardware/low-level**: Rely on the FPGA programmable logic (PL) design based on HDL/RTL logics to achieve short-latency quick processing (in O(μs)).

Keywords: FPGA, PL, HDL, RTL, HLS, ML inference (hls4ml, etc), Versal ACAP, Versal AI engine, short latency in O(μs).

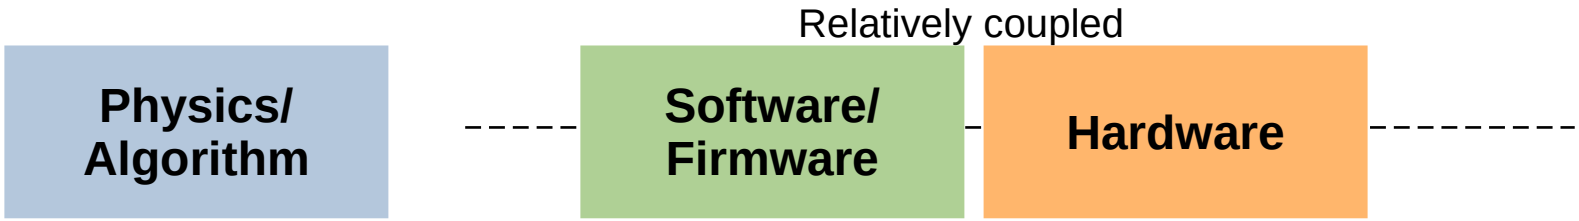

- **High-level**: Software-based algorithm development with computing farm. In addition to CPU/GPU, FPGA acceleration is a new application.

Keywords: Hardware acceleration platforms with CPU/GPU/FPGA, Alveo, Versal acceleration card, Versal DPU.

Collaboration within 7.5a

- How do we work together?
- The physics/algorithm part is mostly dedicated to experiment purpose, so the source codes might not be suitable to be shared.
- The **framework of software/firmware + hardware** is expected part to share.
  - The major core technique for TDAQ experts.
  - The unique core technique from each sub-group.
  - A git repository has been prepared for maintenance purpose.
    The contribution from each sub-group is under discussion within 7.5a.

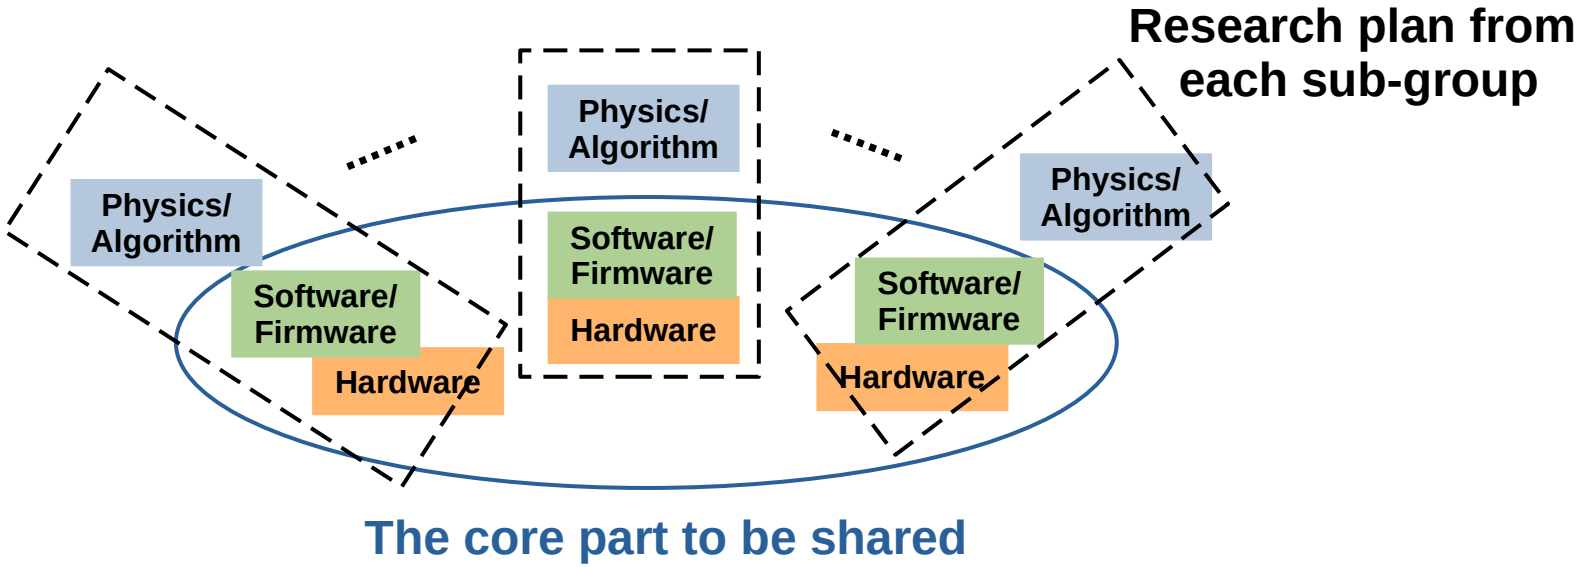

Take FPGA as example

- This technical roadmap was proposed by KEK IPNS group. We are trying to cover most of the methodologies here to build a knowledge database.
- In addition to sharing the framework by git, holding workshop for hand-on practice is also one way to consider as a 7.5a activity.

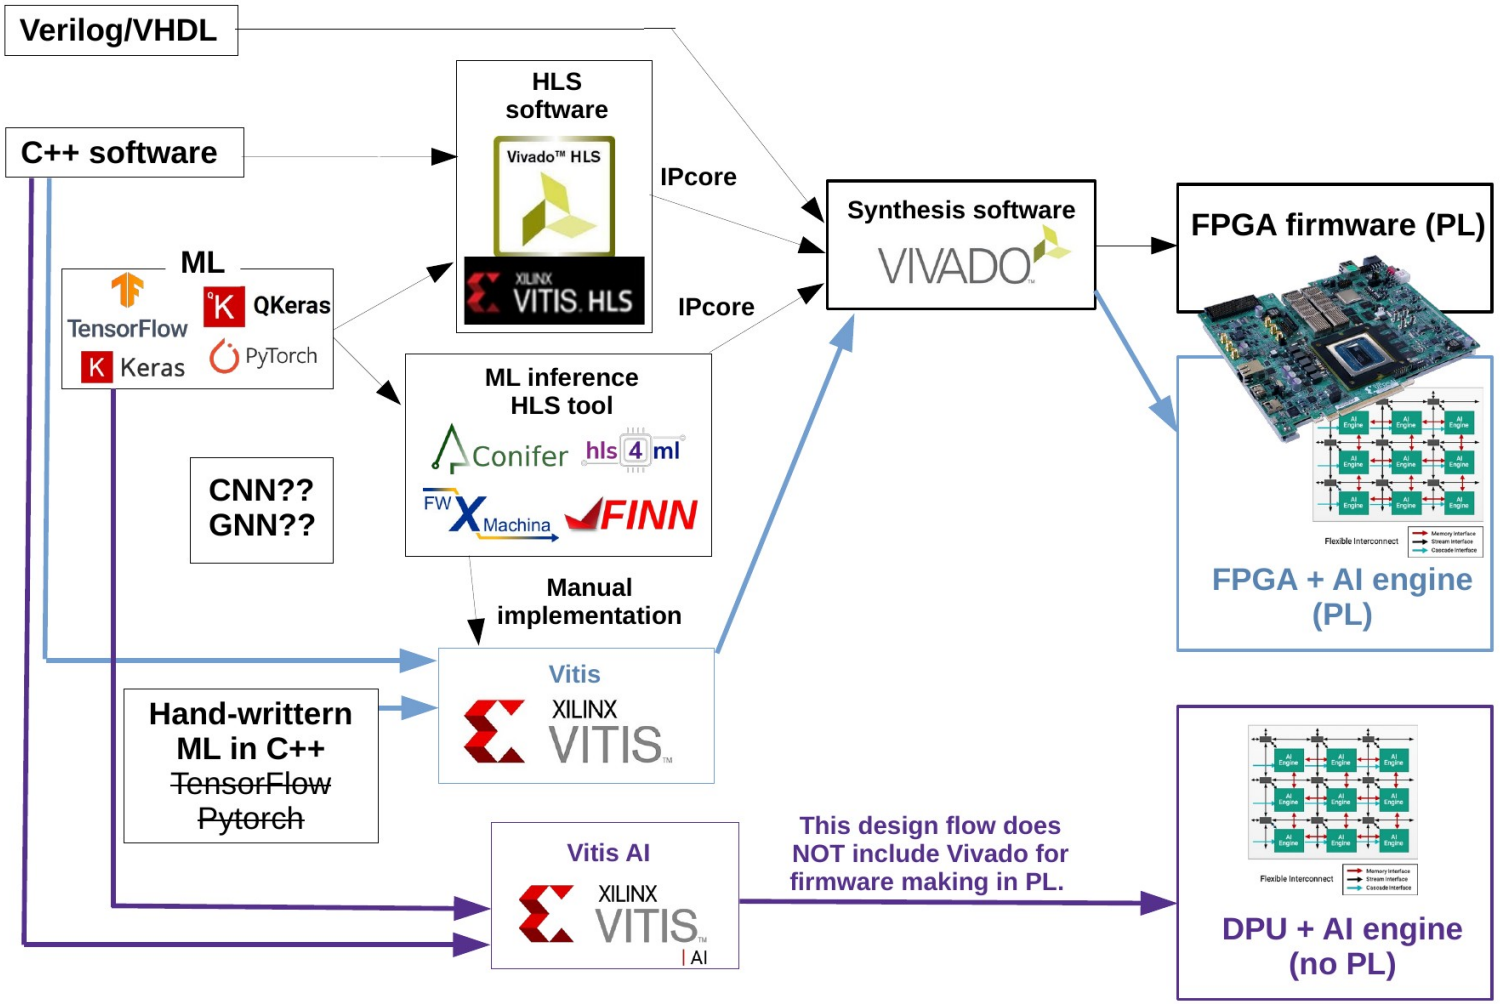

Overview on our activities

- A rough categorization

Low-level: short latency with FPGA (PL)

Birmingham: ATLAS L1/global

RAL: e/gamma and tau trigger btw. BDT and hand crafted algorithm

KEK IPNS: Versal for Belle II and ATLAS trigger

UCM-GAE: ML in FPGA for CTAO camera

Bristol: CMS L1, medical imaging, anomaly, CNN at Versal

Oviedo: GNN for real-time muon reconstruction algorithm with Versal

Geneva: QDIPS FPGA b-tagging

CIEMAT:CMS muon trigger, short-latency, AI in FPGA

High-level: Acceleration with GPU/FPGA

KEK IPNS: Versal DPU for HLT application

IFIC Valencia: ATLAS and LHCb trigger with different acceleration platforms

UCL:ML-based trigger/tracking with GPU/FPGA
ATLAS global electron trigger

Birmingham: ATLAS L1 with FPGA/GPU

Manchester: ML data compression on FPGA with neutrino data and ATLAS jet data

Geneva: ATLAS calos topoclustering FPGA acceleration

RAL: traccc with Intel/Altera FPGA

DAQ/readout

KEK IPNS: PCIe readout with Versal

ANL: Full system for future detector readout

CIEMAT: CMS Drift Tube readout

Other than Trigger/DAQ

Birmingham: HOG for HDL in git

RAL: IPBus, IPPB, and Versal techniques

IFIC Valencia, Oviedo: MC generator with GPU/FPGA

UCL: COTs PMT

Manchester: Q-Pix for FEE with COTs CMOS

CIEMAT: FPGA firmware optimization

ANL: DAQ system

- Experience in various ATLAS systems:
  - ATLAS Readout (aka FELIX) for the Run-3 and Run-4
  - Readout firmware of the ATLAS Global Trigger System for the Run-4
  - Data processing firmware for the Event Filter Track trigger
  - ATCA board design and firmware for FTK to Level-2 Interface Card (FLIC)
  - ATLAS Region of Interest Builder Board (RoIB): board design and firmware
  - PC-based RoIB: software development

- Interests:
  - Full system and digital design for the future detector readout systems
    - High-speed links
    - Low-latency data processing
    - Intelligent high-bandwidth data processing
  - Data flow and online software

CIEMAT: Firmware optimization

- CMS Drift Tube readout and Muon Trigger for the HL-LHC:
  - Very low latency complex trigger algorithms development
  - AI optimization on FPGAs

- **Highly-optimized FPGA algorithms** for performance.
  - High clock speeds (~500 MHz) timing closure
  - Knowing how to write "good firmware code" is critical for developer!

- Features:
  - TCL constraints file
  - FPGA resources and modules utilization, auto size pblocks
  - Customized application for the entire user design
  - Maintainable for user design's upgrade
  - Mostly ready for the community.

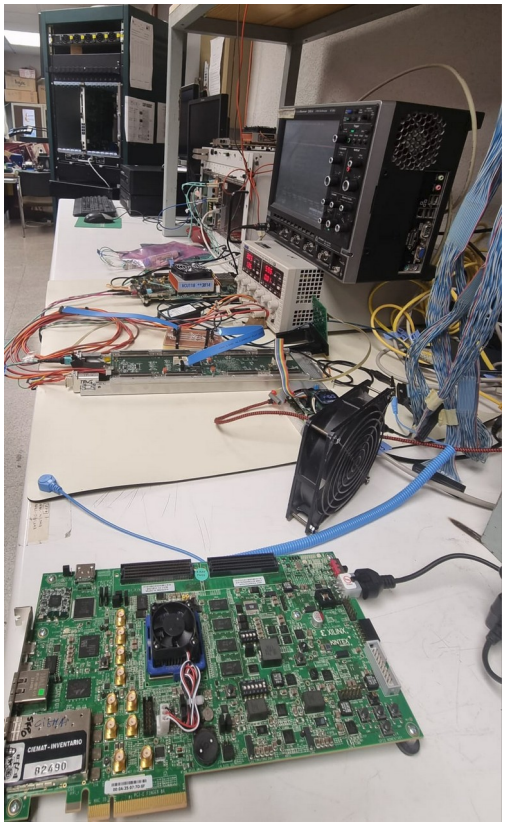

```
link_frame_monitoring
     SSDSSSSSSSSRSSSSSRSSDSSSSSSSSSSSDSSSDSSRSSDSSSSS     USSDSSSSSSSSSS
  Slice shape: 40x60, boundaries=(0,47,0,3)
  Resource utilization:
           dsp        0 /     120 (0.0%)
            ff     2650 /   38400 (6.9%)
           lut      530 /   19200 (2.8%)
          bram       19 /      36 (52.8%)
      ultraram        0 /       0 (0.0%)
```

```
create_pblock S1_link_frame_monitoring
resize_pblock S1_link_frame_monitoring -add { SLICE_X0Y480:SLICE_X39Y539 DSP48E2_X0Y192:DSP48E2_X4Y215 RAMB18_X0Y192:RAMB18_X2Y215 RAMB36_X0Y96:RAMB36_X2Y107 }
set_property IS_SOFT FALSE [get_pblocks S1_link_frame_monitoring]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/d_muxed* ]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/clkstrobes_clkp ]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/pipe_phase_strobes_inst ]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/obdt_decoder_monitor ]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/mortadelo_frame/placeable* ]
add_cells_to_pblock [get_pblock S1_link_frame_monitoring] [get_cells payload/ly1/g[1].g.sector/f.links_to_streams/p40* ]
```

IFIC Valencia: Trigger with FPGA/CPU/GPU

- Computing acceleration for reconstruction algorithms in trigger and generator software in HEP.
  - Low Power consumption
- Various devices for study: Nvidia (Cuda) for GPU, Xilinx Alveo, Xilinx Versal, etc.
- Activities: mainly for ATLAS and LHCb
  - NN on Versal accelerator
  - Downstream tracking with GPUs
  - Matrix multiplication with ALVEO U250
  - Hardware power consumption
  - Signal reconstruction with ML with pile-up
  - “Faraway” tracking reconstruction
  - High-speed simulations with hardware accelerators
- More than trigger: **Implement GPU and FPGA solution in MC generators.**

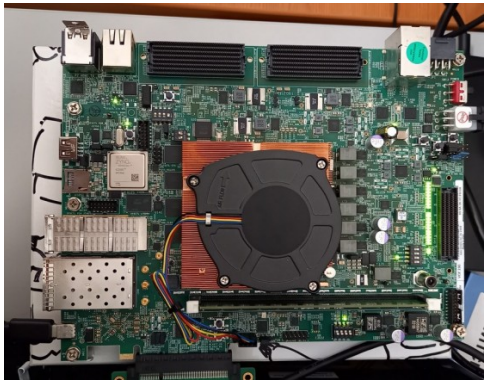

APC Metered Rack PDU ZeroU 2G AP8
SWITCH D-LINK DXS-1210-28T 24x 10GB
GPUs server + CPU : T10G Dual Xeon Scalable HPC 10xGPU PCIe

KEK IPNS: Versal for future device

- KEK IPNS together with Nagoya Univ.

- Using **Xilinx Versal for Trigger application** for both L1 and HLT.
  - Target: **new universal trigger device for Belle II and ATLAS**.
  - Hardware acceleration with Versal DPU for HLT.

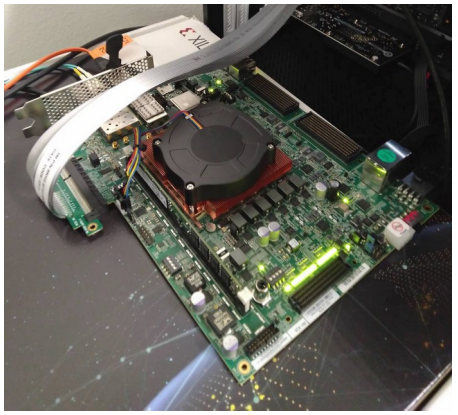

- Test benches with VPK120 and VCK190.
  - Study with PAM4 data link, PCIe Gen5, and AI engine.

- Trigger logic implementation for **Belle II and ATLAS**:
  - Tracking, clustering, topology, anomaly detection.
  - **Deployment in AI engine** is the main technical goal.
  - Comprehensive Methodology study for FPGA:
    HLS-based ML inference, AI engine, DPU, etc.

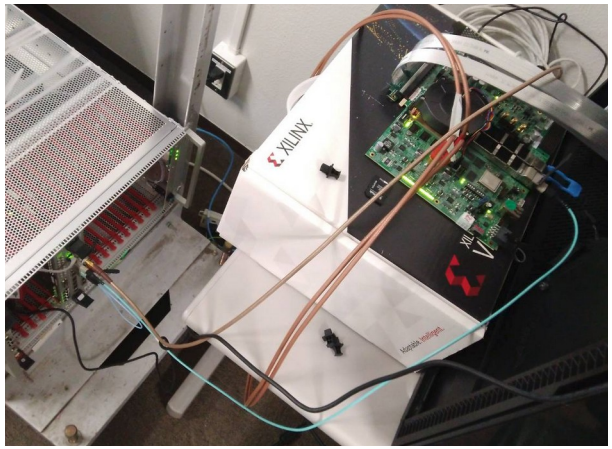

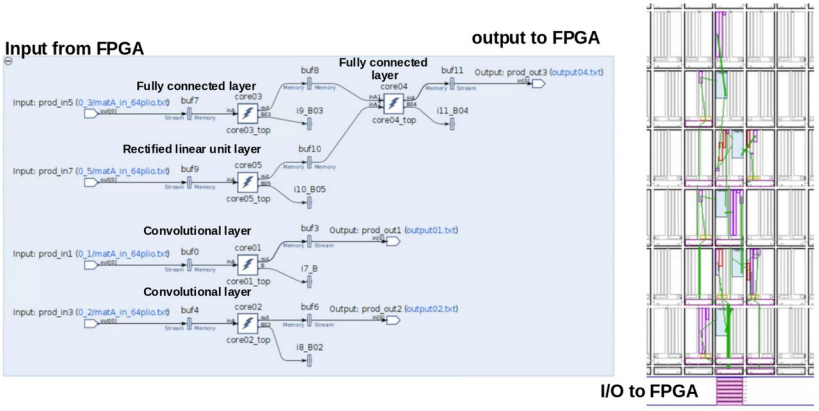

A CNN model deployed in AI engine

RAL: ATLAS L1 trigger

- Long contributions in ATLAS, CMS (both L1 and HLT), and DUNE with collaboration with many groups.

- Device:
  - FPGA: Xilinx Alveo and Intel Arria 10, Stratix 10
  - GPU: Various Nvidia

- Project:
  - **e/gamma and tau low-level trigger algorithms**
    - Improve the present existing one based on BDT and compare the performance with the hand crafted version.
  - SYCL-based tracking algorithm from **traccc with Intel/Altera FPGA**
  - IPbus: **Ethernet** protocol.
  - **IPBB firmware build tool**.
  - Extensive experience with Xilinx Versal: Plan to share the expertise with the colleagues.

UCL: GPU/FPGA ML trigger

- Experience in Data Acquisition in ATLAS for SCT, ITk, global, g-2 and mu2/3e, DUNE
- Target: ML-based trigger and tracking with GPU/FPGA
- Projects:
  - **ML algorithms on Versal to improve ATLAS global electron trigger**
    - BDT using XGBoost
    - Feature generation to ensure sym input
    - Currently passed simulation stage
  - High bandwidth readout for ITk strip and pixel
- Other study with COTs: PMT device

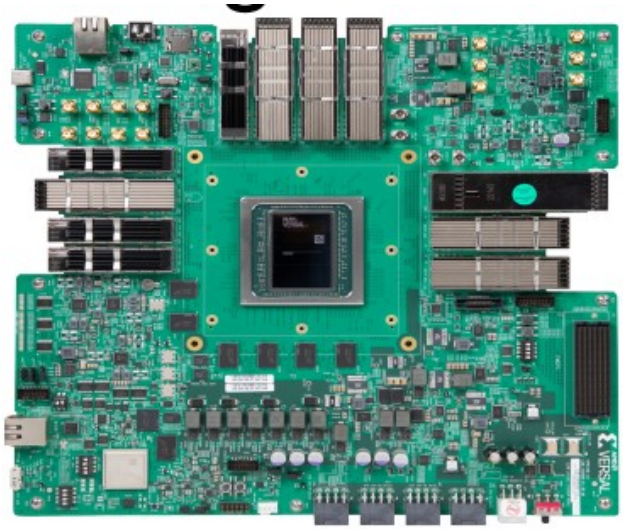

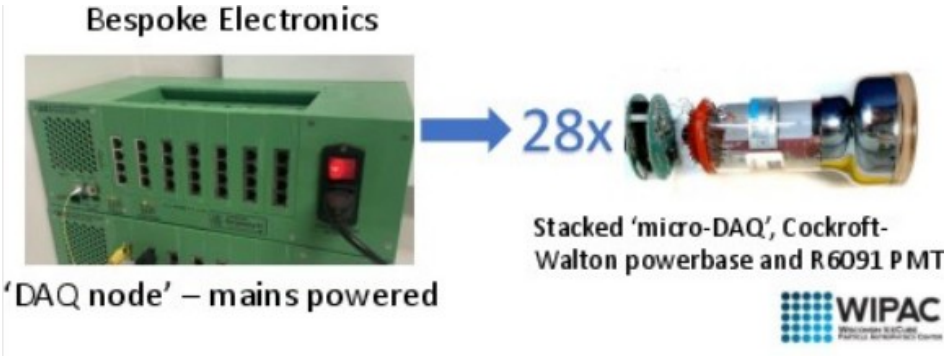

University of Birmingham: ATLAS L1 and HOG

UNIVERSITY OF BIRMINGHAM

- Long-term contribution on Trigger/DAQ. Mainly for ATLAS L1/global trigger, and involved in EIC and muon collider.

- Interested in general infrastructure for hardware trigger (e.g. hog, module control, etc) and efficient (ML-based) selection and compression algorithms
  - Various kinds of algorithms for the **ATLAS L1 trigger**: FPGA and GPU
  - **Anti-kt like algorithm** for various platform for ATLAS Global
  - **Data compression algorithm** with links to future tracking projects
  - **ML-based ATLAS L1 calo**.

- Other than trigger: **HOG**
  - Maintaining HDL in git repository
  - Available: https://gitlab.com/hog-cern/Hog
  - Documentation: http://cern.ch/hog

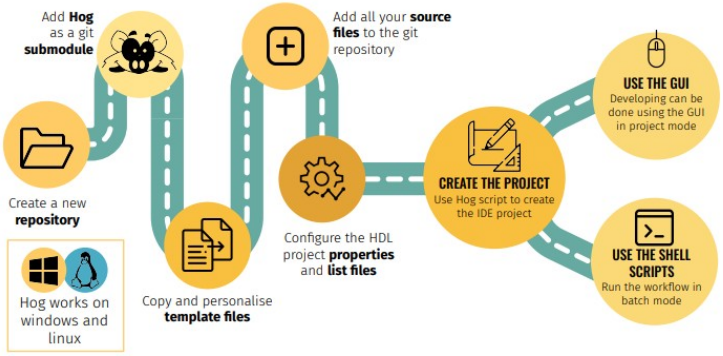

University of Bristol: CMS L1, ML@FPGA

- Experience in ML in FPGA.
- Device: Xilinx Versal and UltraScale
- Projects:
  - **CMS L1 trigger**
    - Short latency
  - Collaboration with North Bristol Trust for **medical imaging/anomaly detection**
  - ML algorithms for data selection for 2D image
    - hls4ml with Xilinx Virtex UltraScale
  - **CNN on Xilinx VCK5000 v.s. CMS ATCA card**

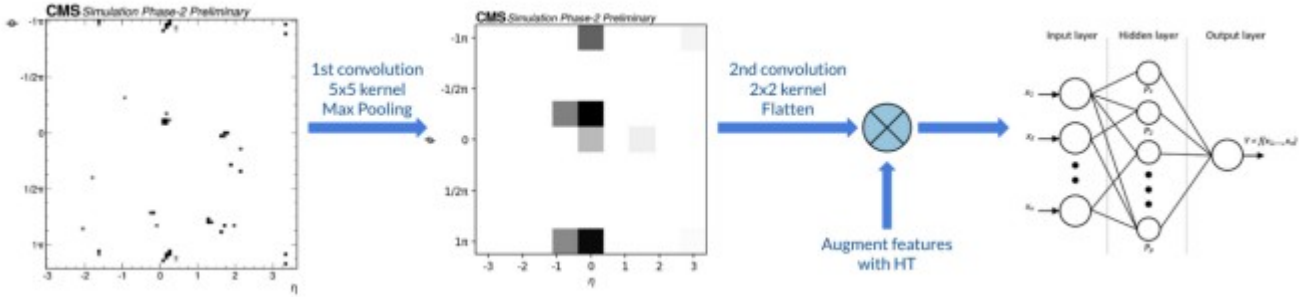

University of Geneva: ML with COTs

- Device: AMD CPU, Xilinx Versal and Alveo, Nvidia GPU

- Projects:
  - **QDIPS**: hls4ml FPGA demonstrator of Deep Sets **b-tagging**
    - Model used for early selection at the HLT on ATLAS
  - **ATLAS calos topoclustering FPGA acceleration**
    - Acceleration with FPGA for EFCalo (Run4)
    - FPGA integration in Athena framework

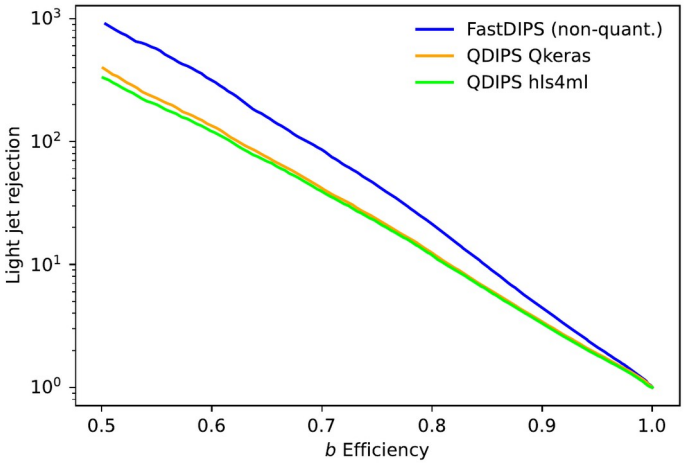

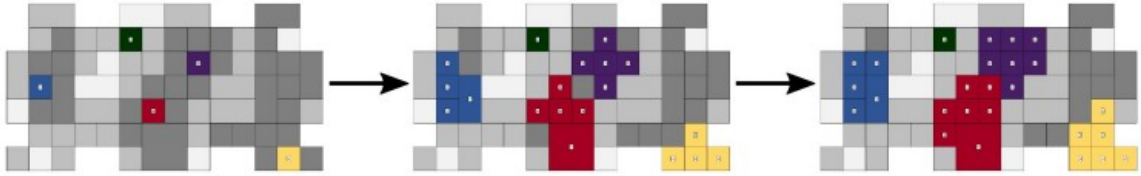

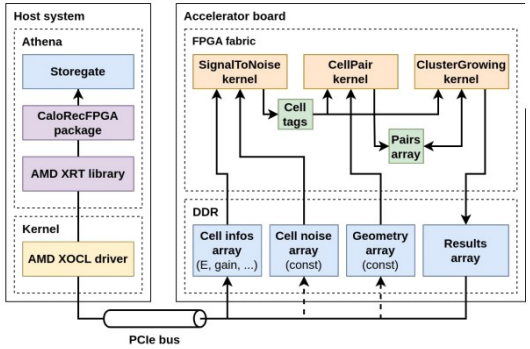

University of Manchester: ML data compression

- Experience in trigger object analysis, tracking@GPU, ML@FPGA
- Device: Xilinx Versal and UltraScale
- Projects: **ML-based data compression on FPGA with neutrino data and ATLAS jet data**
  - Performance metrics for image compression
  - Anomaly detection
  - Data compression and zero suppression
  - Started with CPU/GPU, then plan at FPGA
- Other than trigger:
  Q-Pix. On chip digitization scheme for kTon TPC
  - On-chip "zero suppression" scheme
  - Front-end architecture with COTs CMOS

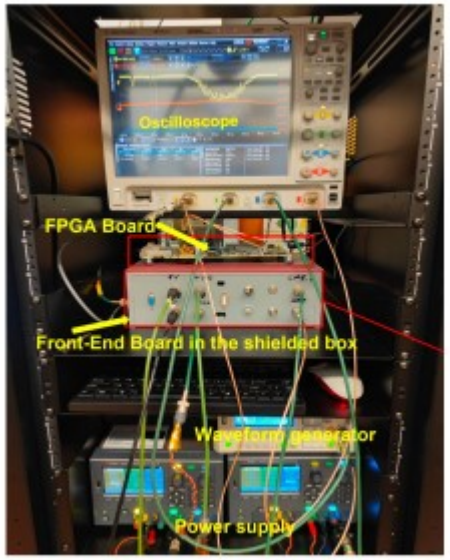

University of Oviedo: GNN in ACAP COTS

Universidad de Oviedo

- Project: **GNN for real-time (O(μs)) muon reconstruction algorithm**
  - Targeting on implementation in Xilinx Versal ACAP and the AI engine.

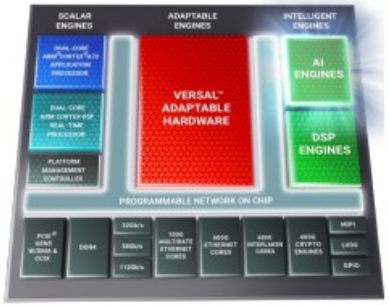

- Other then trigger: **MC generation with accelerated platforms**
  - Parallelization on the CPU-intense steps
  - With GPU and FPGA
  - Compare the time and power consumption

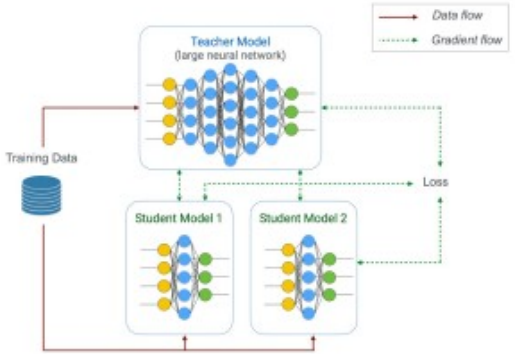

UCM-GAE: ML@FPGA for CTAO Telescopes

- Project: **ML-based hardware trigger with FPGA for camera: EM Calorimeter.**
- CNN models for IACT array-based gamma/hadron/Night-Sky-Background separation.
  - Based on hls4ml.
- Test benches with with KCU kits.
- Home-made prototype board is also under design.

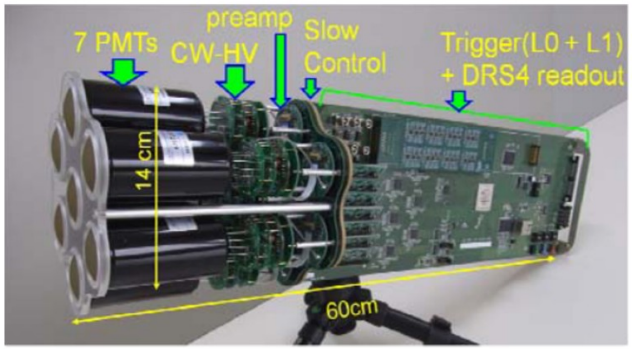

~2000 PMT-based camera

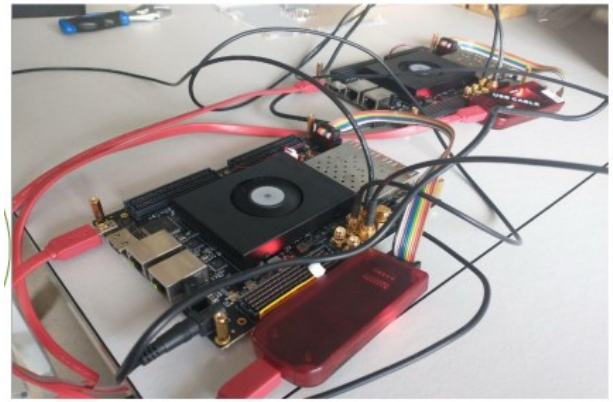

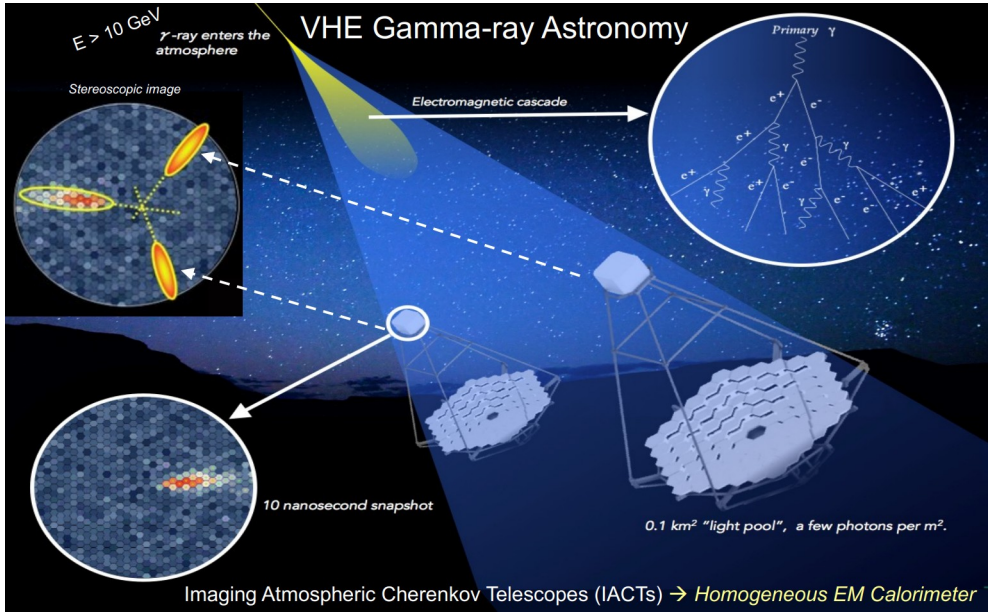

Summary

- 7.5a is aiming for common TDAQ development with COTs.
  - Mainly focus on trigger with various platforms.
  - Also generic development for fundamental purpose.

- Our collaboration will be based on the sharing of technical knowledge of
  - Software/firmware framework
  - Hardware device
  - Using git repository, or even hand-on workshop in the future

- Present activities:
  - PI: Alex Keshavarzi (Manchester), Yun-Tsung lai (KEK IPNS)
  - Monthly meeting on the first Friday of every month
  - Mailing list:
    ECFA-DRD7-PROJECT7_5_a-Contributors@cern.ch
    ECFA-DRD7-PROJECT7_5_a-Observers@cern.ch

- If you are interested in relevant research works and would like to join us, please do not hesitate to let us know!