UC2.2.3 towards MS8 reports

Simone Gennai & Bernardino Spisso

MS8 Deliverables

☑ TAR2.12 [UC2.2.3: Trigger, DAQ, on-line processing: Development of digital trigger logic for a "missing energy" experiment with a positron beam at CERN (POKER/NA64): porting of a missing momentum reconstruction algorithm to FPGA and its testing on a dedicated testbed].

- KPI:

✅ TAR2.13 [UC2.2.3: Developing FPGA tools: Development and testing of RDMA over converged ethernet (ROCE) on FPGA for data transfer from detectors' front-end to computing servers: adapting the system developed the past year to the FPGAs being used at the moment]

- KPI:

Simone.Gennai@cern.ch

FPGA flagship live document

- Link to the flagship document
- Trigger, DAQ and on-line processing
  - Development of trigger and anomaly detection algorithms based on FPGAs at event level for the Atlas experiment.
    - No report expected for MS8
  - Development of track reconstruction on FPGA for LHC-b data acquisition:
    - No report expected for MS8
  - Development of digital trigger logic for a “missing energy” experiment with a positron beam at CERN (POKER/ NA64):
    - Still waiting for openCalls to be finalized …
    - Postponed to MS9
  - Scouting and processing of Level-1 trigger data using FPGA to run on-the-fly momentum object calibration with ML:
    - A Versal Evaluation Board VCK190 has been received and works on this just started
    - NGT doctoral of valentina for H->tautau mass regression on low pt taus for trigger
    - First implementation of W->3pions HLS code is started

Simone.Gennai@cern.ch

FPGA flagship live document

- Link to the flagship document
- Developing FPGA tools
  - Development of a Customizable Framework foar Multi-FPGA Accelerator Generation via architectures:
    - No report expected for MS8
  - Development and testing of RDMA over converged ethernet (ROCE) on FPGA for data transfer from detectors' front-end to computing servers
    - the activities proceed We created a first FW implementation and made direct transfers from the VCU118 to a connect-X NIC mounted on a "receiver" server without interposing switches. Tests have shown that we are capable of saturating the 100 Gbps bandwidth. Several changes to the ROCE stack are planned for a more optimized implementation in FW. We subsequently proceeded in the direction of sending, aggregation and congestion monitoring tests on the networking side of the project, initially through a commercial 10G switch. We are currently testing sending data simultaneously from a NIC mounted on a "sender" server and from the first FW implementation on an "equivalent front-end" evaluation board, to be aggregated by the switch and sent to one or more NICs on "receiver" server. We are currently installing a 100G switch, and will proceed with testing on this.
- Terabit (FPGA clusters) : order sent the past week

Simone.Gennai@cern.ch

Scouting

- W->3pi HLS project
  - “Bug” in xilinxhls backend in Conifer
    - details in conifer/issues#66 → conifer/pull#68 → merged!
  - Working on the code managing the pion pT sorting and triplet creation

Results from `csynth_design` step using 208 input Puppi candidates:

Kernel	LUT	FF	DSP	Latency	CPU time
masker2	214058 (24%)	35780 (2%)	0	3	60 s
orderer3	-	-	-	-	FAIL!
selector2bis	93119 (10%)	23114 (1%)	0	134	1532 s
triplettor2	11020 (1%)	6912 (0%)	0	326	40 s

pivoter	928167 (106%)	4153 (0%)	0	6	30.5 s

Simone.Gennai@cern.ch

Update from ATLAS Group

- From a person power perspective, Salvatore Loffredo from Federico II, University of Naples has started to collaborate to the ATLAS L0 trigger and HLT with the Rome-1 group.
- In particular, he will work on the low-level VHDL implementation of Graph Neural Network architecture in FPGA.
- The target architecture are various types of Xilinx FPGAs, such as the FPGA that will be adopted for the ATLAS Sector Logic (Xilinx Virtex Ultrascale+).

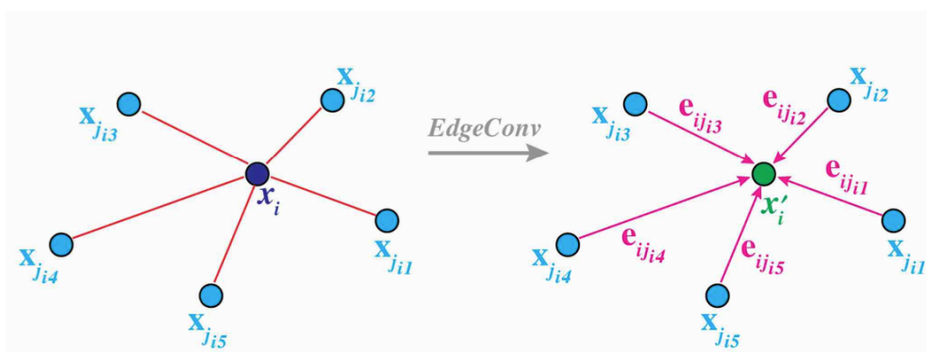

Simone.Gennai@cern.ch

Updated list of papers and contributions

Tabella 1

Titolo	Autori	Link a paper/conference	Accepted for PUB/PRoceeding
Fast Neural Network Inference on FPGAs for Triggering on Long-Lived Particles at Colliders	Andrea Coccaro Francesco Armando Di Bello Stefano Giagu Lucrezia Rambelli and Nicola Stocchetti	https://arxiv.org/pdf/2307.05152.pdf	
Sviluppo di acceleratori per il Machine Learning e sistemi di Inference as a Service su FPGA	Daniele Spiga, Diego Ciangottini , Giacomo Surace, Giulio Bianchini, Loriano Storchi , Mirko Mariotti	Workshop Loano	
KServe inference extension for a FPGA vendor-free ecosystem	Daniele Spiga, Diego Ciangottini , Giacomo Surace, Giulio Bianchini, Loriano Storchi , Mirko Mariotti	CHEP 2023	
Deep Learning techniques for reconstruction on ASTRI Mini-Array Monte Carlo data	Saverio Lombardi, Francesco Visconti, Michele Mastropietro	https://pos.sissa.it/444/713/pdf	
A novel explainable approach in radiomics pipeline for local recurrence prediction of lung cancer: a feasibility study exploiting high energy physics potential to evaluate the model	Mariagrazia Monteleone, Simone Gennai, Pietro Govoni, Chiara Paganelli	ACM ISBN 979-8-4007-0815-2/23/09. https://doi.org/10.1145/3632047.3632074	ACM ISBN 979-8-4007-0815-2/23/09. https://doi.org/10.1145/3632047.3632074
Triggerless data acquisition pipeline for Machine Learning based statistical anomaly detection	Gaia Grosso, Nicolò Lai, Matteo Migliorini, Jacopo Pazzini, Andrea Triossi, Marco Zanetti, Alberto Zucchetta	CHEP 2023	G. Grosso et al EPJ Web of Conf., 295 (2024) 02033
40MHz Triggerless Readout of the CMS Drift Tube Muon Detector	Matteo Migliorini, Jacopo Pazzini, Andrea Triossi, Marco Zanetti	TWEPP 2023	M. Migliorini et al 2024 JINST 19 C02050
Front-End RDMA Over Converged Ethernet, real-time firmware simulation	Gabriele Bortolato, Antonio Bergnoli, Damiano Bortolato, Daniele Mengoni, Matteo Migliorini, Fabio Montecassiano, Jacopo Pazzini, Sandro Ventura, Andrea Triossi, Marco Zanetti	TWEPP 2023	G. Bortolato et al 2024 JINST 19 C03038
Front-End Rdma Over Converged Ethernet, real-time firmware simulation	Gabriele Bortolato, Antonio Bergnoli, Damiano Bortolato, Daniele Mengoni, Matteo Migliorini, Fabio Montecassiano, Jacopo Pazzini, Sandro Ventura, Andrea Triossi, Marco Zanetti	TIPP 2023	
The CMS Level-1 trigger data scouting for LHC run 3 and the CMS phase-2 upgrade	Sabrina Giorgetti (Matteo Migliorini, Rocco Ardino, Jacopo Pazzini, Andrea Triossi, Marco Zanetti) on behalf of the CMS Collaboration	La Thuile 2024 - YSF	
Hardware implementation of quantum machine learning predictors for ultra-low latency applications	Lorenzo Borella, Alberto Coppi, Jacopo Pazzini, Andrea Stanco, Andrea Triossi, Marco Zanetti	EuCAIFCon 2024	
Quantum machine learning classifiers implemented on FPGA for ultra-low latency applications	Lorenzo Borella, Alberto Coppi, Jacopo Pazzini, Andrea Stanco, Andrea Triossi, Marco Zanetti	ICHEP 2024	

Simone.Gennai@cern.ch

Status KPIs

Link to the flagship document

KPI ID	Description	Acceptance threshold	Status up to today
KPI2.2.3.1	Development of triggering algorithms, on-line analyses, data acquisition on FPGA	Submission of 1 paper to a peer-reviewed journal	1 paper already accepted
KPI2.2.3.2	Online scouting	Submission of 1 paper to a peer-reviewed journal	Abstract being submitted to ichep 2024 about scouting
KPI2.2.3.3	Development of tools to integrate several FPGAs together	Submission of 1 paper to a peer-reviewed journal	G. Bortolato et al 2024 JINST 19 C03038
KPI2.2.3.4	Organizing courses about FPGA programming on low and high level	At least two courses organized	1 course done at the end of 2023 1 VHDL course done in February 2024

Simone.Gennai@cern.ch