

DRD7 - Electronics

L. Fiorini (IFIC-Valencia)

DRD Days @ IFIC
09/03/2026

with inputs or contributions from:

Sonakshi Ahuja, Fernando Carrió, Francesco Curcio, Alvaro Fernandez, Héctor Gutierrez, Valerii Kholoimov, Miriam Lucio, Laura Molina, Arantza Oyanguren, Ximo Poveda, Arantxa Ruiz, Karan Singh, Alberto Valero, Rui Wang, Jiahui Zhuo.





<https://drd7.web.cern.ch>
68 Collaborating
Institutes
registered.

- Spokesperson: Frank Simon (Karlsruhe)
- Deputy Spokesp.: Angelo Rivetti (INFN Turin)
- Collaboration Board (CB) Chair: Attilio Andreazza (U. Milan)
- Deputy CB Chair: Cristina Fernandez (CIEMAT)
- **IFIC is part of the DRD7 Collaboration**
- Institute Representative (IR): L.F.
- Deputy IR: Alberto Valero

WP7.1: Data density and power efficiency

WP7.2: Intelligence on the detector

WP7.3: 4D and 5D techniques

WP7.4: Extreme environments

WP7.5: Backend systems and COTS components

WP7.6: Complex imaging ASICs and technologies

WG7.7: Tools and Technologies

- A discussion about the MoU of DRD7 is in progress.
- Membership to the Collaboration is linked to the participation in at least one WP Project.
- No financial commitment is required from the Funding Agency*
- The signature of the MoU is not mandatory to be part of the Collaboration.

*: Under discussion

Draft 11 – status: 5 December 2025 16:01

DD Month YYYY

Memorandum of Understanding
for the
Detector Research and Development **n (DRD**n**)**
Collaboration
Concerning **XXXX**

between

THE EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH, “CERN”, an Intergovernmental Organization having its seat in Geneva, Switzerland, as Host Laboratory

on the one hand,



and

the Collaborating Institutions/Funding Agencies of the DRD7 Collaboration


on the other hand,

Rack K RETEX LOGIC-2 A600 42U F1000 PH
APC Metered Rack PDU ZeroU 2G AP8
SWITCH D-LINK DXS-1210-28T 24x 10GB

- HL01: T10G Dual Xeon Scalable
 - 1 x NVIDIA RTX A5000 24GB GDDR6
 - 1 x NVIDIA RTX A6000 Ada Generation 48GB DDR6
- HL02: NG8 Dual Epyc 9004-8B
 - 2 x NVIDIA H100 NVL 94GB
 - 1 x NVIDIA RTX A6000 Ada Generation 48GB DDR6
- ARM Ampere Altra



FLEXIBILITY ← → EFFICIENCY



ATLAS TileCal Preprocessors

The ATLAS TileCal PreProcessor (TilePPr): is the interface between the “on-detector electronics” and the trigger & data acquisition system, 40 MHz (40 MHz x 1 Mb event size = 40 Tb/s)


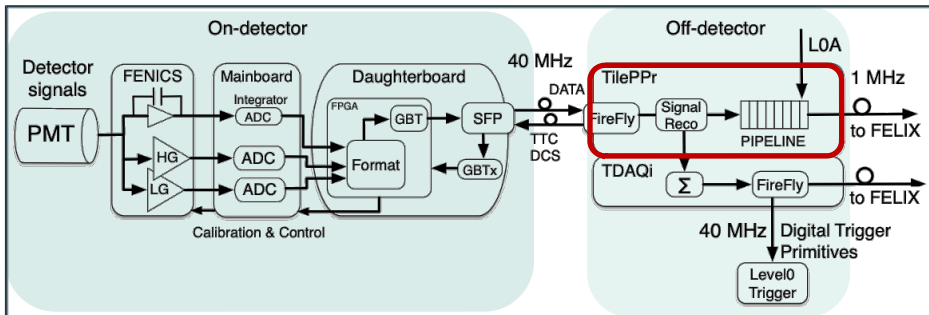
A total of 32 TilePPrs are needed to read the TileCalorimeter. Each PPr is formed by

- 1 ATCA Carrier board (including IPMC, TileCoM, GbESwitch)
- 4 CPMs (Compact Processing Module)

Final designs approved. **First lot now in production (<25%)**



The TilePPr main tasks are:

- Real time signal reconstruction to obtain cell energy depositions
- Provides digital trigger primitives to Level 0 Trigger Systems
- Readout events after LO Acceptance
- LHC clock and control commands transmitted to on-detector electronics



ATLAS TileCal Signal Reconstruction

- During HL-LHC, ATLAS Tile Calorimeter signals will be processed in real-time at 40 MHz with fixed latency before sending it to the L0 trigger.
- Higher pileup and rates require better algorithms than linear filters used for LHC
- Signal processed in Xilinx Kintex Ultrascale KU115
- Input per FPGA is 2×77 channels = 154
- Several Machine Learning algorithms based on MLP, CNN and RNN have been tested.
- Number of parameters is a bottleneck (FPGA resources)

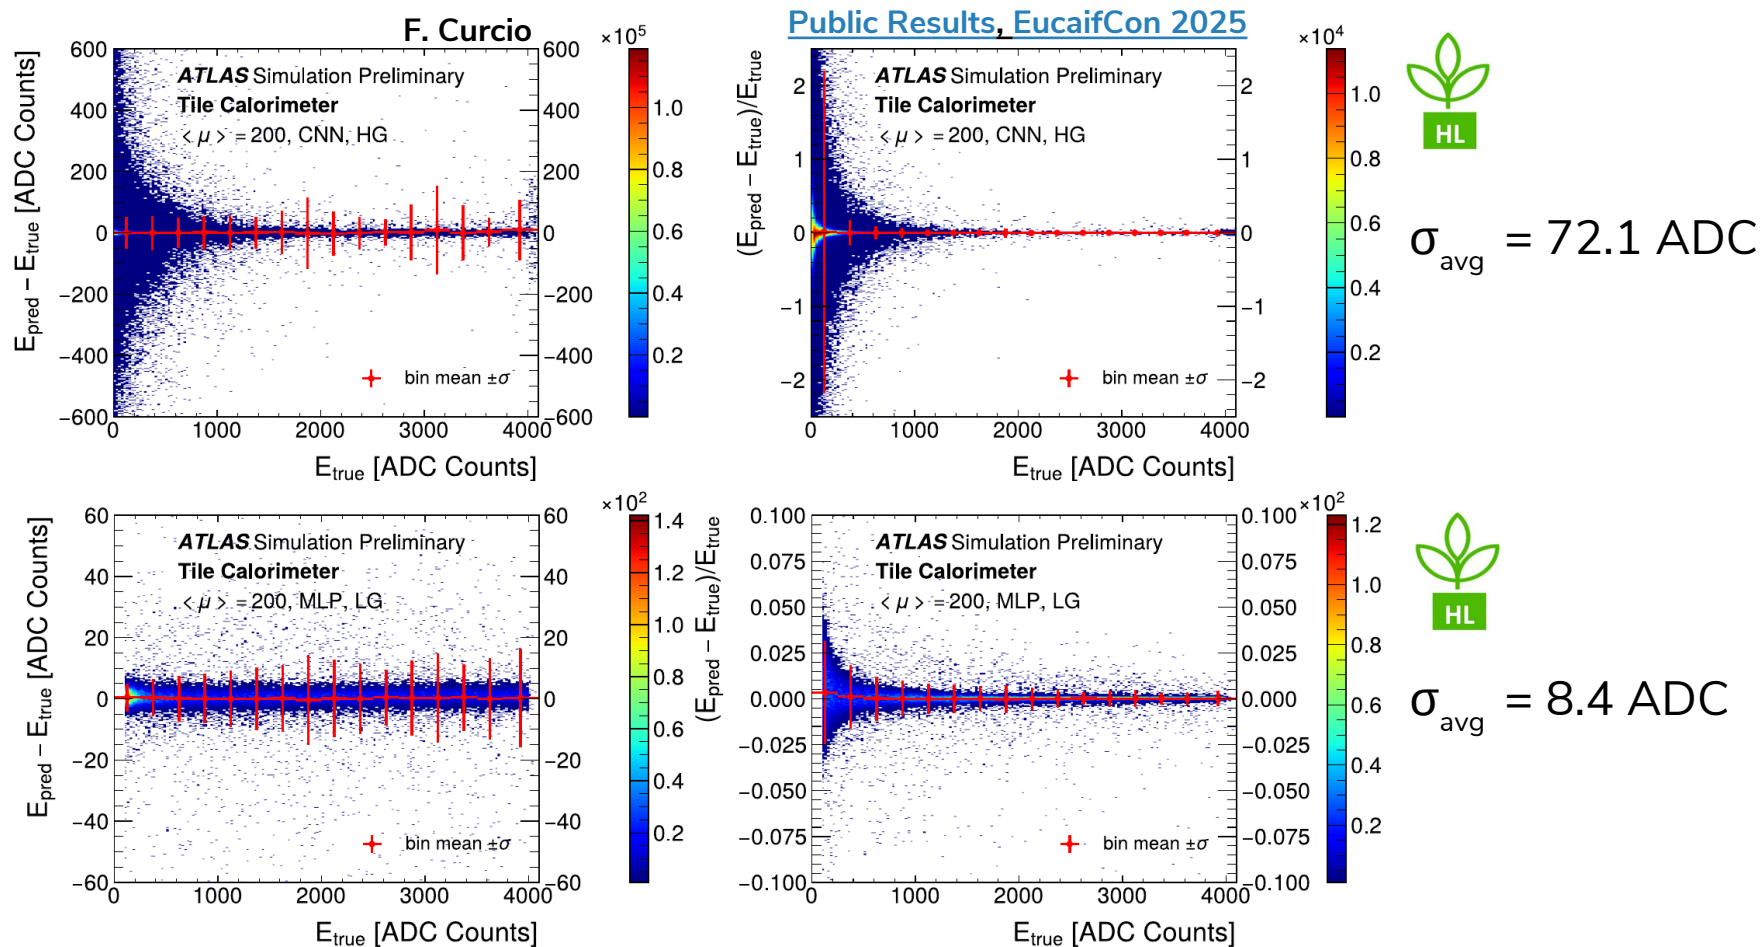


Compact Processing Module (CPM)

Firmware Block	Latency
Uplink	50 ns, 2 BC
Data Decoder	12.5 ns, 0.5 BC
Energy reconstruction + sample delay	225 ns, 9 BC
Trigger Packer	12.5 ns, 0.5 BC
Trigger Interface	25 ns, 1 BC
Total	325 ns, 13 BC

ATLAS TileCal Signal Reconstruction

- Input dataset based on specific simulation replicating HL-LHC conditions.
- Algorithms trained with PyTorch.
- CNN trained with hybrid loss ($0.5 \cdot \text{MAE} + 0.5 \cdot \text{RMSE}$) gives best trade-off.



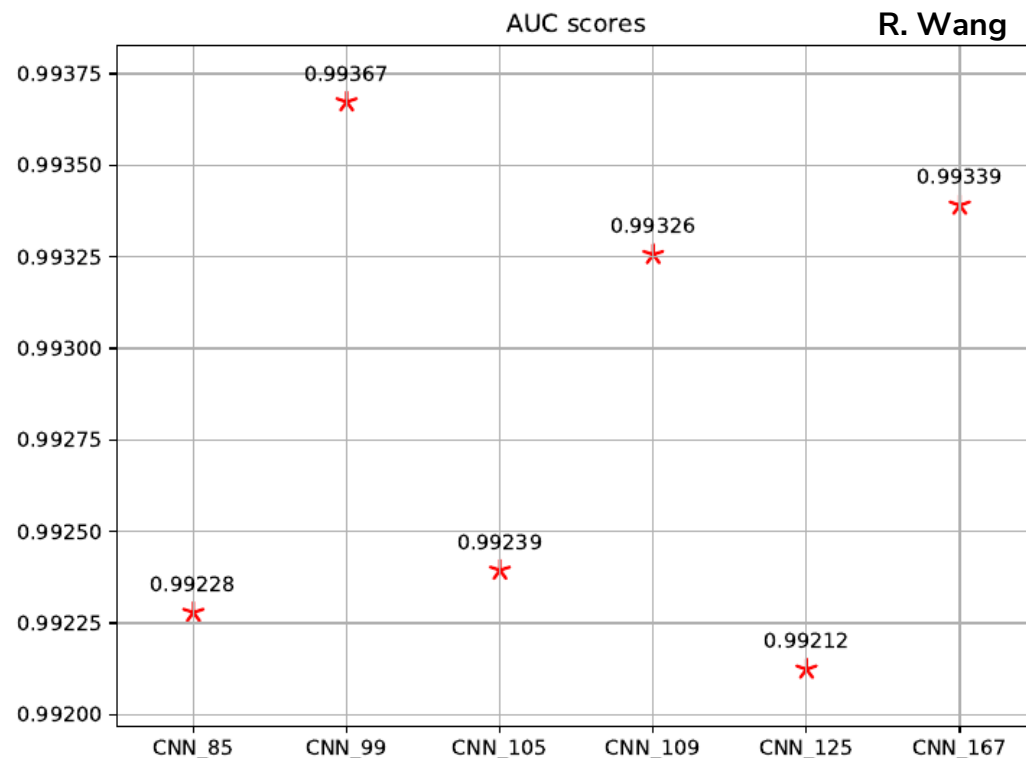
ATLAS TileCal Signal Detection

- There is little gain to reconstruct samples without signal.
- The ratio of noise over signal for the preprocessed samples is checked for individual cells
- Even if high pileup the majority of the PMTs have no signal in a given BC.

Cell	Ratio
A1	1.97
A2	1.86
A3	1.70
A4	1.68
A5	1.49
A6	1.40
A7	1.26
A8	1.04
A9	0.83
A10	0.84
BC1	5.51
BC2	5.33

Cell	Ratio
BC3	5.22
BC4	4.84
BC5	4.58
BC6	4.06
BC7	3.59
BC8	2.11
B9	2.68
D0	53.70
D1	53.88
D2	50.78
D3	10.52
A12	0.42
A13	0.14

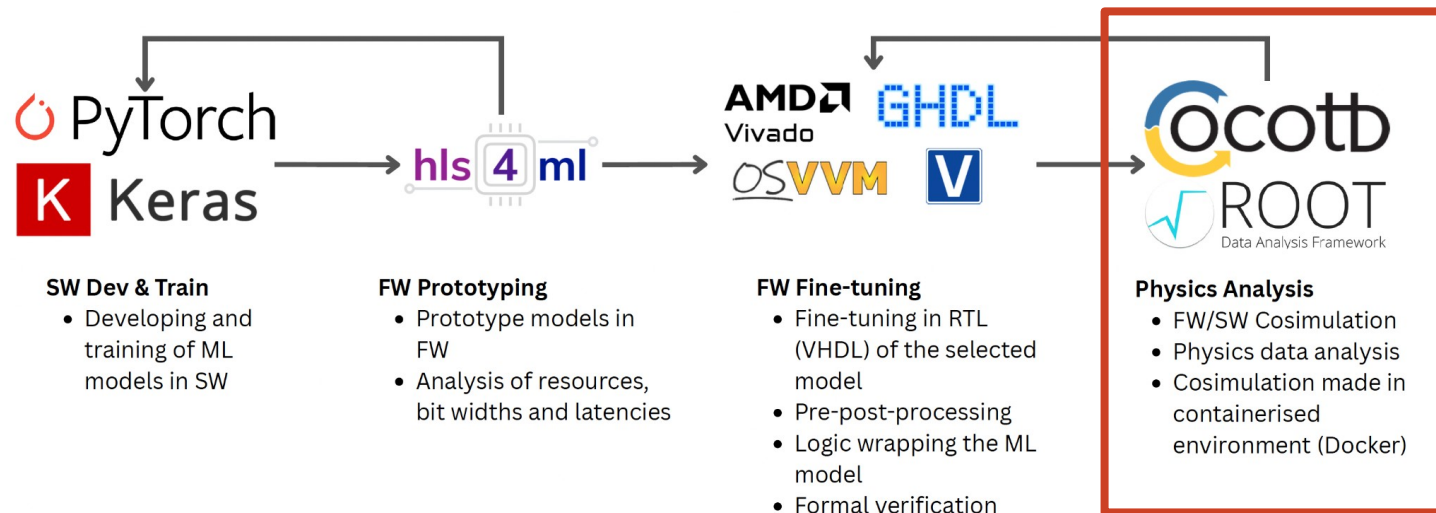
Cell	Ratio
A14	0.34
A15	1.34
A16	5.24
B11	0.94
B12	1.34
B13	2.37
B14	6.04
B15	18.52
C10	0.65
D4	3.73
D5	2.52
D6	24.13



VHDL simulation framework


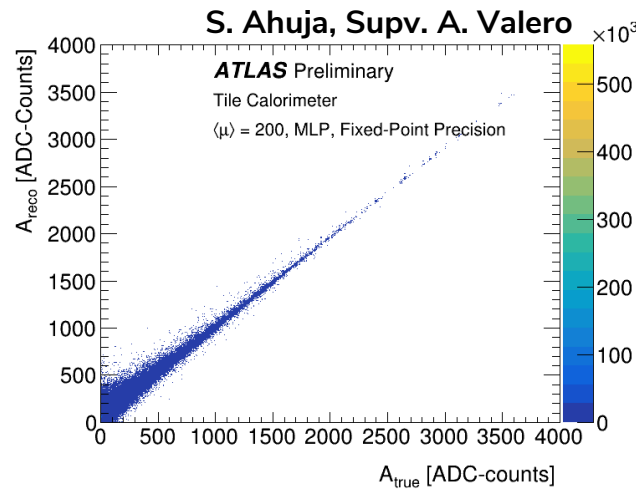
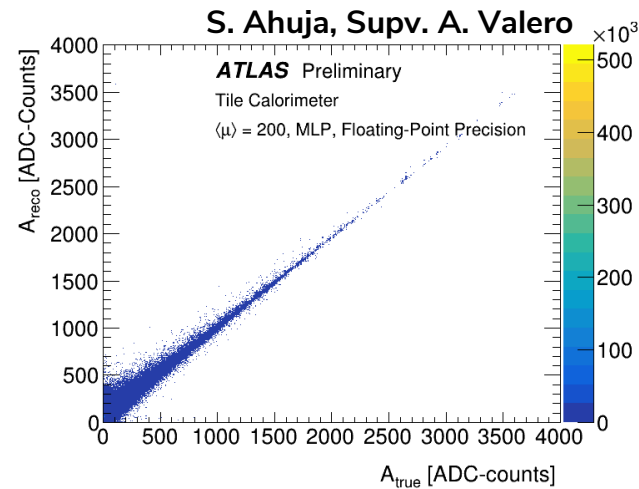
- Python-based simulation framework to validate signal reconstruction algorithms.
- Input simulated MC or detector data.
 - It can be adapted for any VHDL block.
 - It accepts ROOT ntuples as input/output.
- Implemented as a Docker container including the standalone framework.
- Firmware/Software co-simulation
- Useful for fw validation

[EPJ Web Conf. 338 \(2025\) 01004](#)



Use case: Neural Network on FPGA

- Implementation of a neural network for signal reconstruction in high pileup environments.
 - Pytorch/Keras (Software) → VHDL (Firmware).
- Event-by-event comparison between software and VHDL firmware.
- Validation of the algorithm and the implementation in a single framework.

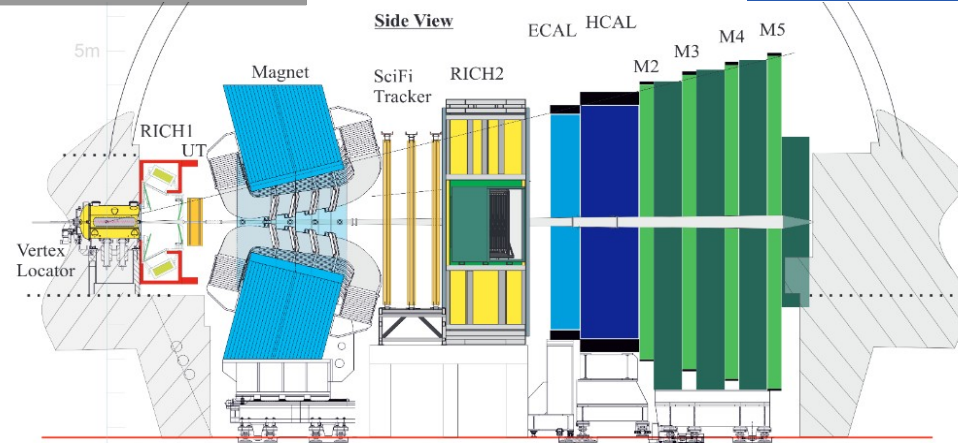




Deliverables:


- Code and documentation of the framework.

Real Time Analysis @ LHCb

- **Allen**: the LHCb “trigger-less” high-level trigger 1 (HLT1) application on GPUs. [LHCb-TDR-021] Fast detector reconstruction in O(500) Nvidia RTX A5000.



	LHCb	ATLAS	CMS	ALICE
$\mathcal{L} [cm^{-2} s^{-1}]$	2×10^{33}	2×10^{34}	2×10^{34}	6×10^{29}
pile-up	5	60	60	1
reconstruction rate	30 MHz	100 kHz	100 kHz	50 kHz
reconstructed tracks/s	1800 M	90 M	90 M	10 M



NA64 experiment is exploring as well the idea of a trigger-less readout → software-only HLT.

Allen is open-source project

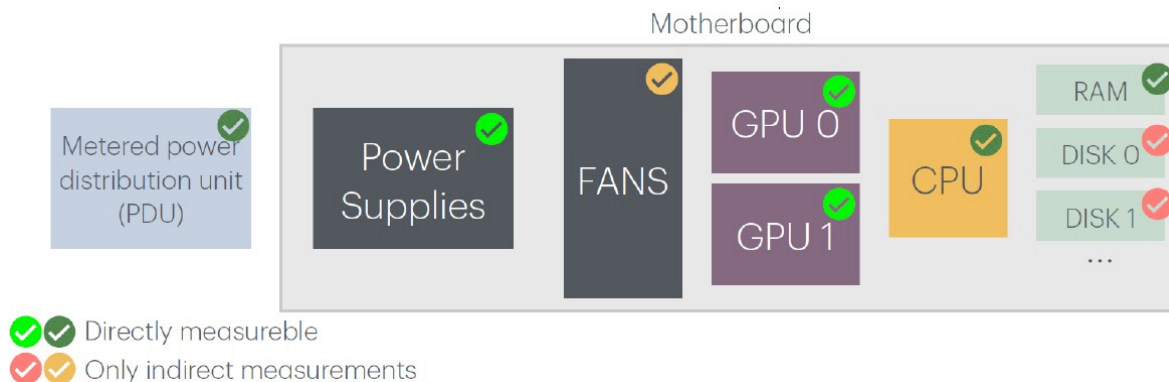
- Portable: executed on several architectures: CPU, GPU
- Modular: design allows various execution sequences
- Total of approx. 250 algos used in data-taking
- It has to reduce in real time the input rate of 40 Tbits/s by a factor 50
- Available in LHCb simulation

Real Time Analysis @ LHCb

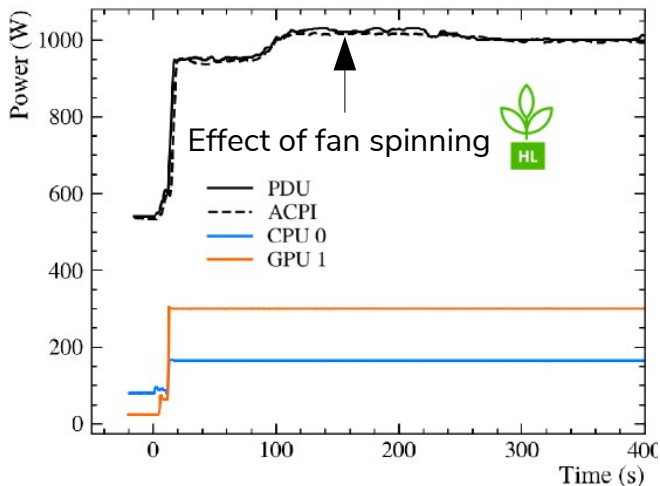
- Power consumption measurements of **Allen algorithms** performed using dedicated external hardware and specific software to access the built-in sensors.



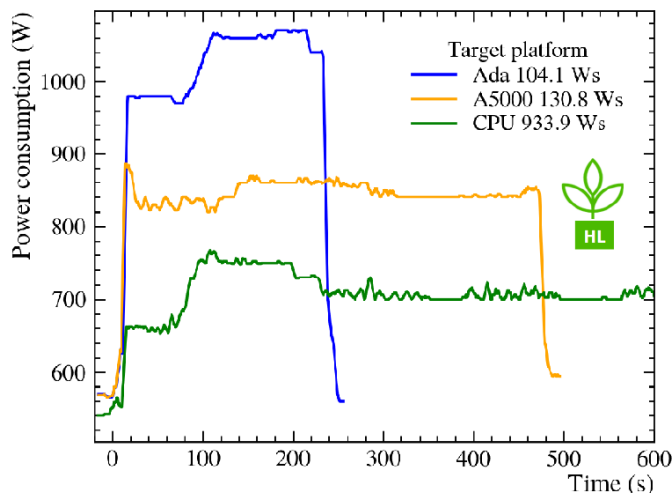
Fan speed measurement device



CPU + NVIDIA RTX 6000 Ada 320M events



40 M events, PDU power usage



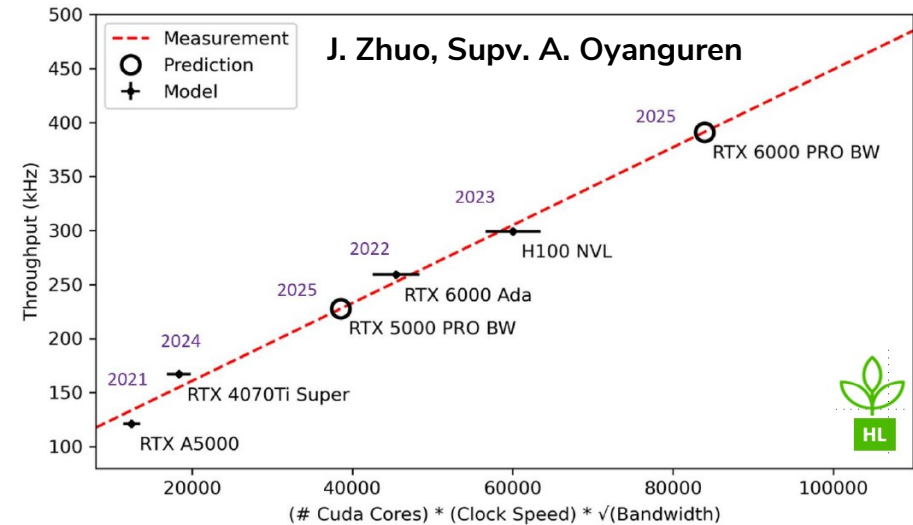
[EPJ Web of Conferences 337, 01272 \(2025\)](#)

More powerful devices with faster execution time usually exhibit less power consumption:
(Throughput \uparrow Energy \downarrow)

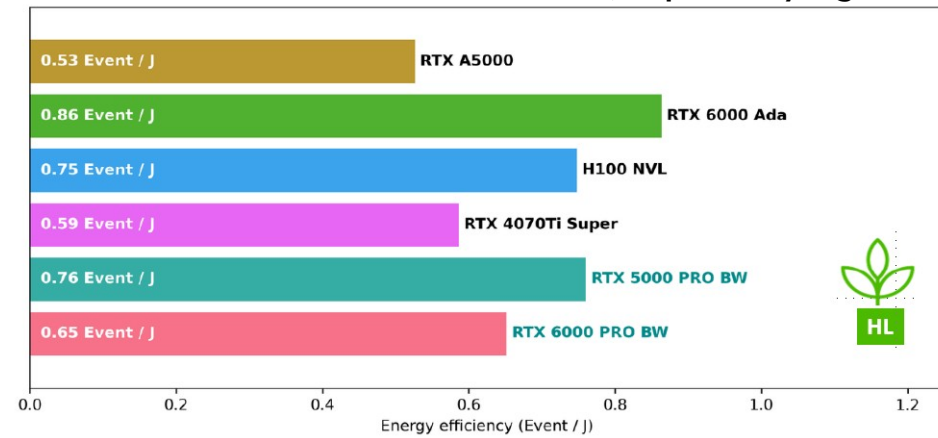
Power Consumption: Hardware optimization

- **Recent Results!**
 - Work in Progress
- Define **Energy Efficiency (EE)** for a given task:
 - $EE = \text{Throughput} / \text{TDP} (\# \text{events} / \text{J})$
 - **TDP (W)** in a GPU is the **Thermal Design Power**: the maximum amount of heat the cooling system is designed to dissipate during normal operation.
- Measured throughput (Allen sequence, 5M evts) vs Predicted one from GPU parameters (available in data sheets)

Evaluating the Energy Efficiency (EE) of a GPU:



J. Zhuo, Supv. A. Oyanguren



- Including **power consumption** as figure of merit in addition to **efficiency, throughput and physics performance**

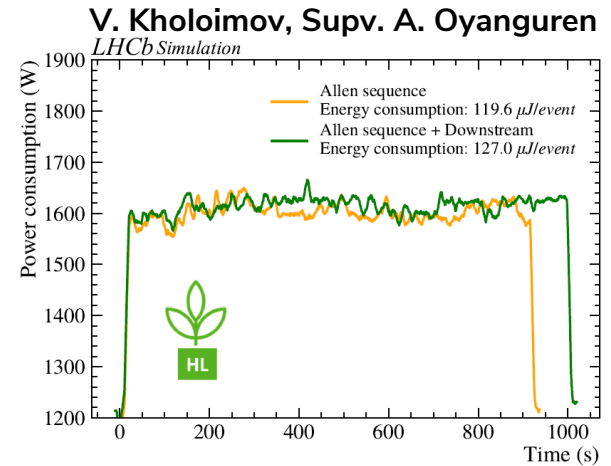



Fig. 17: Power consumption with **Allen** software running over 3.2M $B_s \rightarrow \phi\phi$ events without (blue) and with (orange) **Downstream** algorithm. The power

Seeding and Matching algorithms:
 only GPU measurements (paper in preparation)



A **Downstream** and vertexing algorithm for Long Lived Particles (LLP) selection at the first High level trigger (HLT1) of LHCb

V. Kholoimov¹, B. Kishor Jashal^{1,2}, A. Oyanguren¹, V. Svintozelskyi¹ and J. Zhuo¹
¹Instituto de Física Corpuscular (IFIC), University of Valencia- CSIC, Valencia, Spain.
²Rutherford Appleton Laboratory (RAL), Oxford, United Kingdom.

2.4.5 Power consumption






The effect on the power consumption from the execution of **Downstream** algorithm in the HLT1 sequence is studied in the following and shown in Fig. 17. Several techniques are employed to measure the power consumption including the use of a metered power distribution unit (PDU¹) within the rack, analysis of device driver outputs (e.g., Nvidia

¹An APC PDU AP8858EU3 is used in this work.

[Comp. and Sof. for Big Science 9,10, (2025)]

Calorimetry and Electronics R&D for FCC


- New Project at IFIC “Advanced Data Processing Technologies for Exploring New Physics in Future Particle Colliders”
 - Funded with 600k € during 2025-29 by G. Valenciana under the Prometheus program
- Researchers (physicist / engineer):
 - IFIC: Arantxa Ruiz (PI), Alberto Valero, Ximo Poveda, Fernando Carrió, Juan Valls, David Hernandez
 - CIEMAT: Cristina Fernández, Ignacio Redondo
 - Students: Ana Arranz, Ivan Burriel
- Certification of MPPC/ SiPM
 - Hamamatsu MPPC (16Ch, 64Ch)
 - Single channel (OnSemi, Avago)
 - Different cell side (number of pixels)
- Now with an LED → we have ordered a Laser
 - Thorlabs NPL45C - Nanosecond Pulsed Laser Diode System
 - To be used with Integrating sphere to characterize the light





Read-out for Calorimeters R&D

Development of read-out for the R&D of FCC Calorimeters (TileCal-like Allegro hadronic calorimeter):


- Currently, the certification of SiPM is performed using commercial modular electronic modules previously acquired.
- Plan to develop a fully custom DAQ system for the readout of the SiPM arrays.
- Evaluation of RADIOROC ASIC with readout from FPGA (testboard ordered)
- Replicate in a custom scalable system with high throughput DAQ



Current readout modules



Planned setup


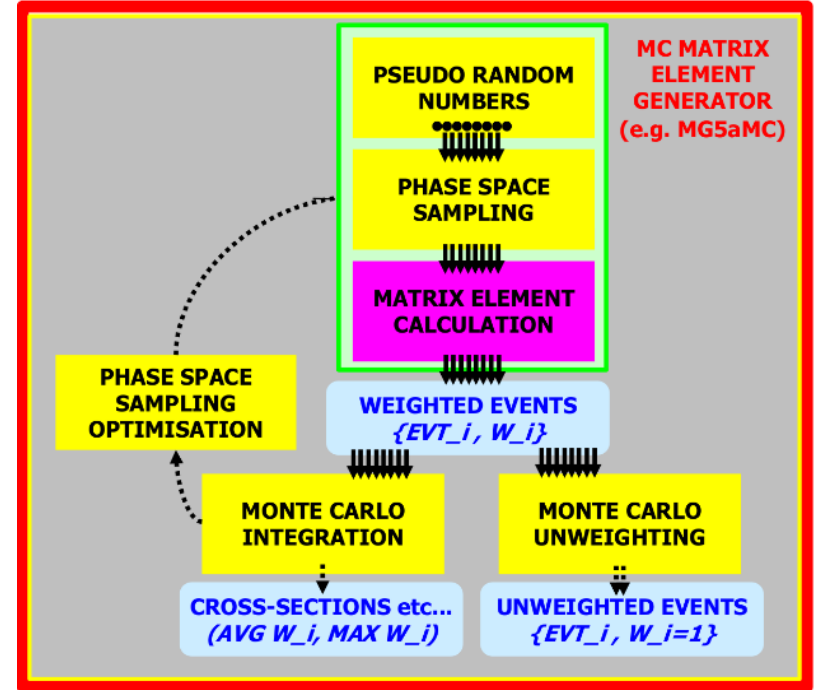


More info in [DRD Calo talk](#)

Porting MADGRAPH to FPGA using HLS

The Computational Challenge

- Event generation in HEP consumes significant CPU resources.
- As the event generation becomes more and more precise ($N^{\times}LO$) the Matrix Element calculation becomes a bottleneck for HEP computing.
- Explore new architectures (GPUs, FPGAs) to accelerate these calculations in Event Generators like MADGRAPH.
- Projects like [Madgraph4GPU](#) already showed nice progress.
- We are testing a further step by porting the ME calculation to an FPGA (Alveo U250) using [HLS](#) for several processes
 - $e^+e^- \rightarrow \mu^+\mu^-$ process.
 - $gg \rightarrow t\bar{t}+jets$



Platform and Tools



- **Hardware:** AMD Xilinx Alveo U250 FPGA operating at 120 MHz.
- **Software:** Vitis HLS (C/C++) and Xilinx Runtime (XRT) for host-FPGA communication.
- HLS Adaptation (from GPU to FPGA)

Performance tested on $17M e^+e^- \rightarrow \mu^+\mu^-$ generated events. EPJ Web Conf. 337 (2025) 01173

- The FPGA (Alveo U250) completed the simulation in **9.63 s**, $1765 \cdot 10^3$ ev/s.
- **~19x faster** than the GPU (NVIDIA RTX 3050 8 GB) $93 \cdot 10^3$ ev/s.
- **~51x faster** than the CPU (Intel i7-13700), $35 \cdot 10^3$ ev/s.

Energy Measurements:

- CPU: 2.287 mJ/ev
- GPU: 1.553 mJ/ev
- FPGA: 0.0042 mJ/ev (0.0017 mJ/ev Host consumption excluded)



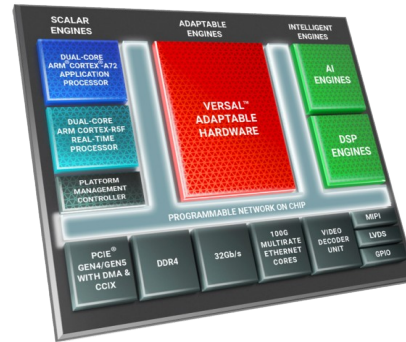

Amplitude and Color in hadronic processes

Platform and Tools

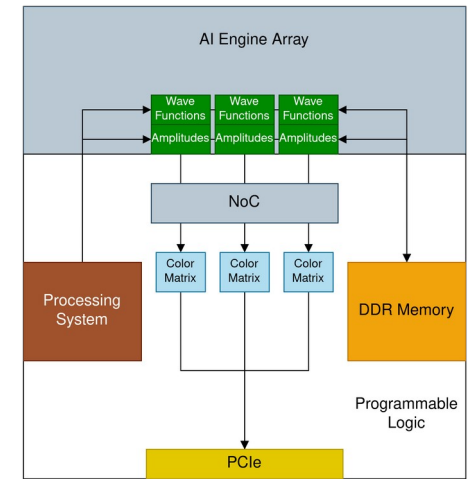
- For complex hadronic processes, time spent by color recombination grows exponentially
- Use Versal AI engines for computation of the color part

	$gg \rightarrow t\bar{t}$	$gg \rightarrow t\bar{t}gg$	$gg \rightarrow t\bar{t}ggg$
madevent	13G	470G	11T
matrix1	3.1G (23%)	450G (96%)	11T (>99%)


- color
- amplitude
- int/propagator
- external
- not ME



VCK190




Complete process



$gg \rightarrow tt\sim g$

Pipeline Level	AIE HWSim [ns/event]	Max. Cores	* Aggregated HWSim [ns/event/max_cores]	CPU [ns/event]	Speedup Factor
17	1300	5	260	600	x2.3
10	1300	10	130		x4.6
7	1750	14	125		x4.8

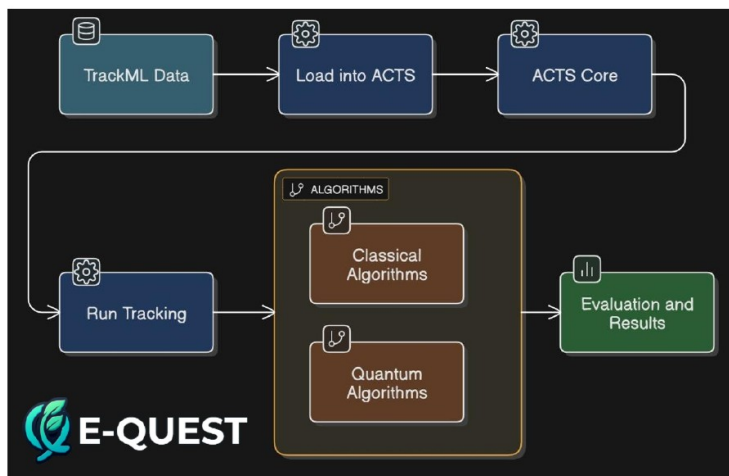
Color Matrix (FPGA-HLS)



Process	t_CPU/evt (ns)	t_FPGA/evt (ns)
$gg \rightarrow tt\sim g$	24	13
$gg \rightarrow tt\sim gg$	120	16
$gg \rightarrow tt\sim ggg$	4400	170

Sustainability in Quantum Computing


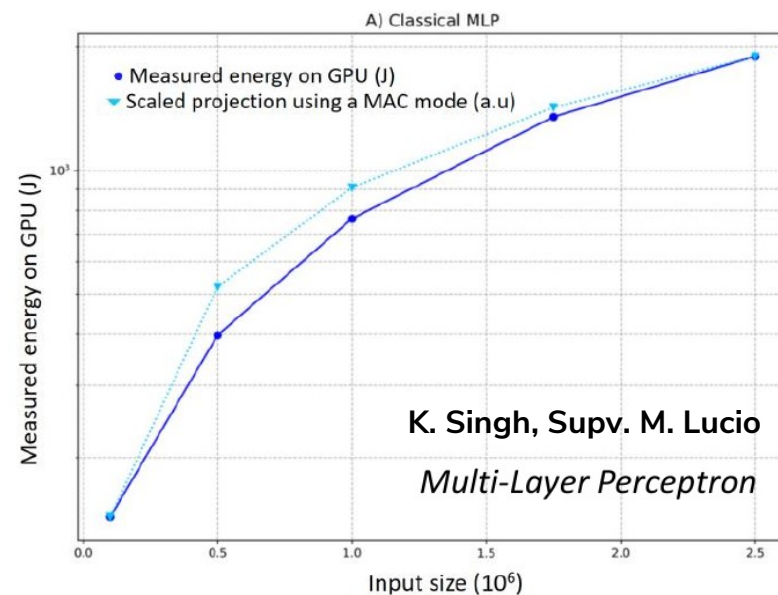
- Quantum computing is an emerging technology.
- Testing quantum computing on complex reconstruction tasks, like tracking.



- General framework for studying the energy consumption of quantum and classical computation being established to compare power consumption with classic algorithms.



[PRX Energy 4 \(2025\) 2, 023008](#)

Is complexity a good f.o.m. for power consumption?



Conclusions

- Computing in HEP is evolving fast.
- DRD7 focuses on the development of Electronics for next-generation experiments in high-energy physics. Heterogeneous computing systems allows to optimize the resources and power consumption.
- Main activities @IFIC:
 - Online processing
 - Power consumption
 - DAQ Readout with Commercial electronics
 - VHDL Simulation framework
 - Hybrid systems for MC generation
- Additionally, Quantum Computing as an emerging technology.



UNIVERSITAT DE VALÈNCIA





Backup

LHCb Hardware optimization

- Optimization of hardware utilization
- GPU parallelisation (# CUDA streams)
- Optimize the GPU usage for higher performance

- Checking other hardware architectures in tracking reconstruction:
- Real-time reconstruction on FPGAs with the “artificial retina” architecture
- Clustering of the VELO detector already running for Run3 in FPGAs
- SciFi tracking in development for Run4 (2030) [CERN-LHCC-2024-001]

Total power consumption 10M events, RTX 6000 Ada



(J. Zhuo)

[<https://cds.cern.ch/record/2888549>]

Activity Example: DNN Signal Reco

arXiv.org > physics > arXiv:1903.02439

Search or Article

(Help | Advanced se

Physics > Instrumentation and Detectors


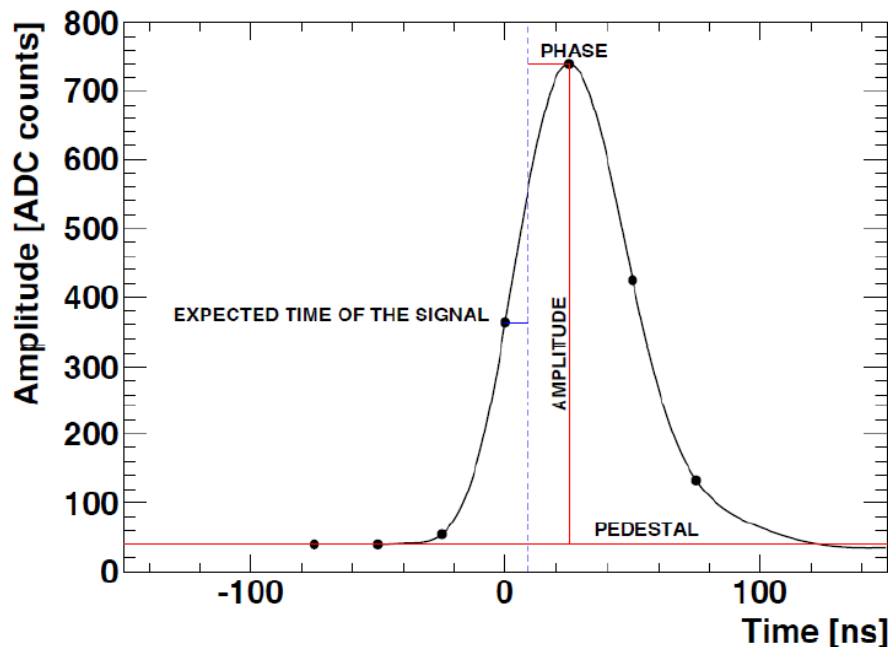
FPGA implementation of a deep learning algorithm for real-time signal reconstruction in radiation detectors under high pile-up conditions

J. L. Ortiz, F. Carrió, A. Valero

(Submitted on 6 Mar 2019)




- **Deep NN for real-time signal reconstruction** at the HL-LHC.
- HL-LHC pileup degrades the pulse quality, LHC algos performance deteriorates.
- 128 FPGAs to process the full ATLAS Tile Calorimeter.
- Each FPGA process 96 different signals with 40 MHz rate







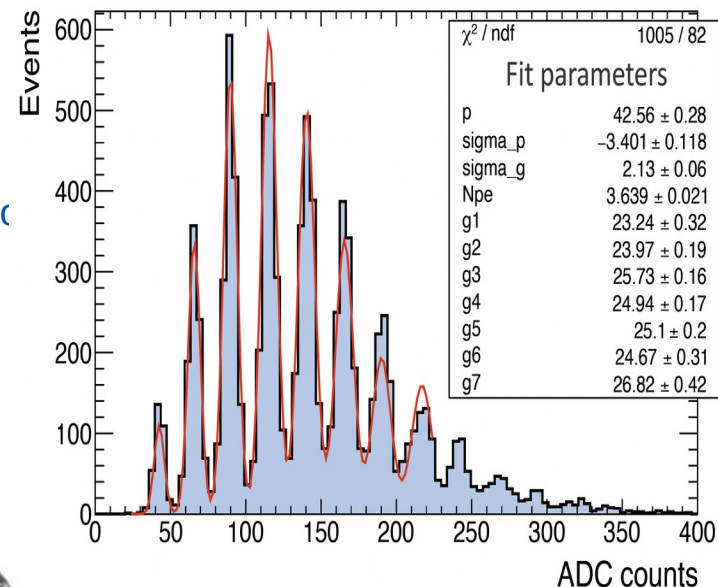
Activity Example: Machine & Deep Learning

- Application of **Machine Learning** in LHC searches for new physics.
- **Computer Vision** (ConvNets) Network for event classification, muon tomography and search for long-lived particles
 - Application as well for satellite imaging, medical diagnosis, etc.



Calorimetry and Electronics R&D for FCC

- New Project at IFIC “Advanced Data Processing Technologies for Exploring New Physics in Future Particle Colliders”
 - Funded with 600k € during 2025-29 by G. Valenciana under the Prometheus program
- Researchers (physicist / engineer):
 - IFIC: Arantxa Ruiz (PI), Alberto Valero, Ximo Poveda, Fernando Carrió, Juan Valls, David Hernandez
 - CIEMAT: Cristina Fernández, Ignacio Redondo
 - Students: Ana Arranz, Ivan Burriel
- Certification of MPPC/ SiPM
 - Hamamatsu MPPC (16Ch, 64Ch)
 - Single channel (OnSemi, Avago)
 - Different cell side (number of pixels)
- Now with an LED → we have ordered a Laser
 - Thorlabs NPL45C - Nanosecond Pulsed Laser Diode System
 - To be used with Integrating sphere to characterize the light



Resource Optimization

$$e^+ e^- \rightarrow \mu^+ \mu^-$$

double → **ap_fixed <24,8>**

Nevts = 16777216

f FPGA = 100 MHz

Interval_fp = 22

Interval = 1

Resource	Utilization	Available	Utilization(%)
LUT	192101	1759631	10.9
FF	218211	3660140	5.96
DSP	1129	12424	9.08
BRAM	31	5442	0.6

Time (s)	Platform	SpeedUp (Platform/FPGA_imp)
4.623	CPU(13th Intel(R) Core (TM) i7-13700) (1 thread)	22
5.32	CPU (AMD EPYC 9474F 48- Core Processor) (1 thread)	25
3.3	RTX 3050 8GB	16
0.814	ADA 6000	3.87
0.645	H100	3.07
3.69	FPGA (floating point)	17.57
0.21	FPGA (fixed point) (1CU)	1

Resource Optimization

$$e^+ e^- \rightarrow \mu^+ \mu^-$$

double → ap_fixed <24,8>

Nevts = 16777216


Resource	Utilization	Available	Utilization(%)
LUT	573077	1759631	32.5
FF	335408	3660140	9.16
DSP	8544	12424	68.77
BRAM	128	5442	2.35

Time (s)	Platform	SpeedUp (Platform/FPGA_imp)
4.623	CPU(13th Intel(R) Core (TM) i7-13700) (1 thread)	128.41
5.32	CPU (AMD EPYC 9474F 48-Core Processor) (1 thread)	147.7
2.43	CPU (AMD EPYC 9474F 48-Core Processor) (multi thread)	67.5
3.3	RTX 3050 8GB	91.6
0.814	ADA 6000	22.61
0.645	H100	17.91
3.69	FPGA (floating point)	102.5
0.036	FPGA (fixed point) (8CU)	1

$$t_{FPGA}(N_{CU}) \approx \frac{t_{FPGA(1CU)}}{N_{CU}}$$

Resource Optimization

Momenta Relative Error Distribution



Matrix Element Relative Error Distribution

