Design Ideas for an Online Data Reduction System for the ePIC dRICH Detector

Alessandro Lonardo, Cristian Rossi

INFN Roma, APE Lab

for the ePIC Roma1/2 team

dRICH Meeting - Irradiation tests and Data filter
January 29th 2025

ePIC General DAQ Scheme

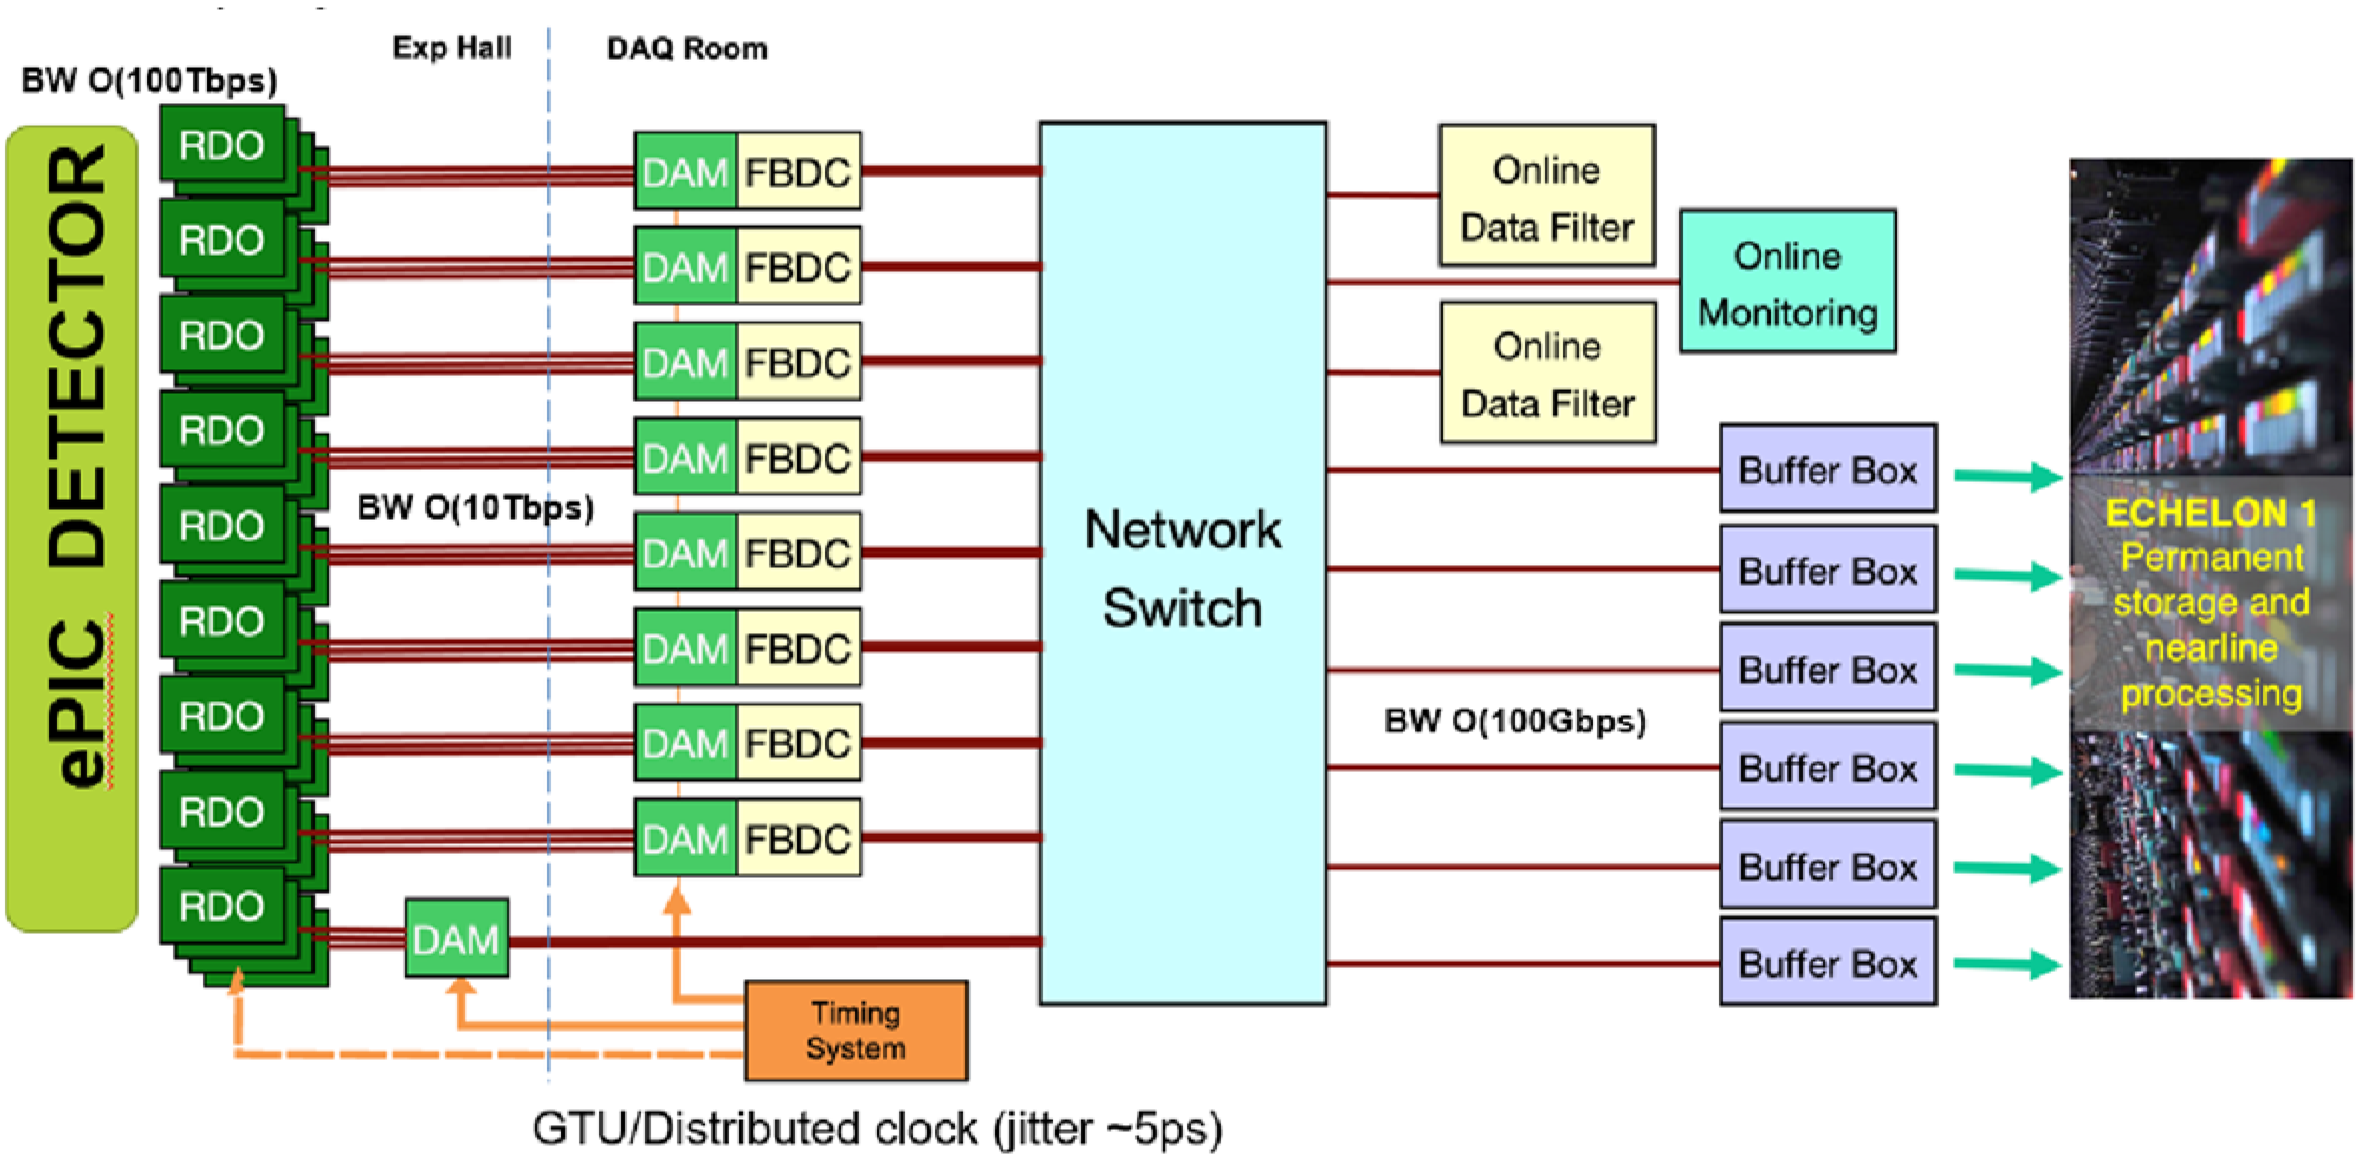

Dual Radiator RICH (dRICH)

Forward cap

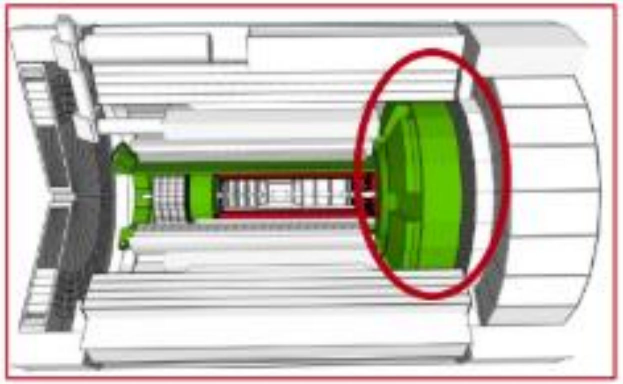

SiPMs

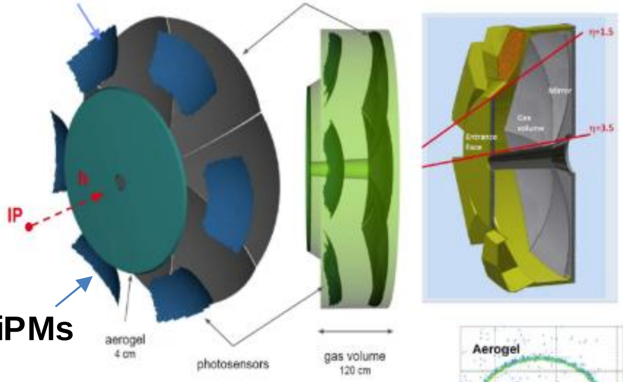

SiPMs

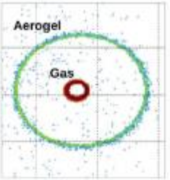

Aerogel: n = 1. 020 – 1.026

Gas: C_2F_6 (hexafluoroethane)

PDU

- **1 photodetector unit PDU: 4x64 SiPM array device (256 channels), 4 FEBs, 1 RDO**
- **1248 PDUs for full dRICH readout**
- **319488 readout channels divided in six sectors**

FELIX

DAM
Data Aggregation Module

Assembled FLX-155

SERVER

- **42 links from PDUs to Felix-155 board**
- **30 Felix-155 boards in total**

ePIC processing and storage system

Analysis of dRICH Output Bandwidth

The dRICH DAQ chain in ePIC ➔ **the throughput issue**

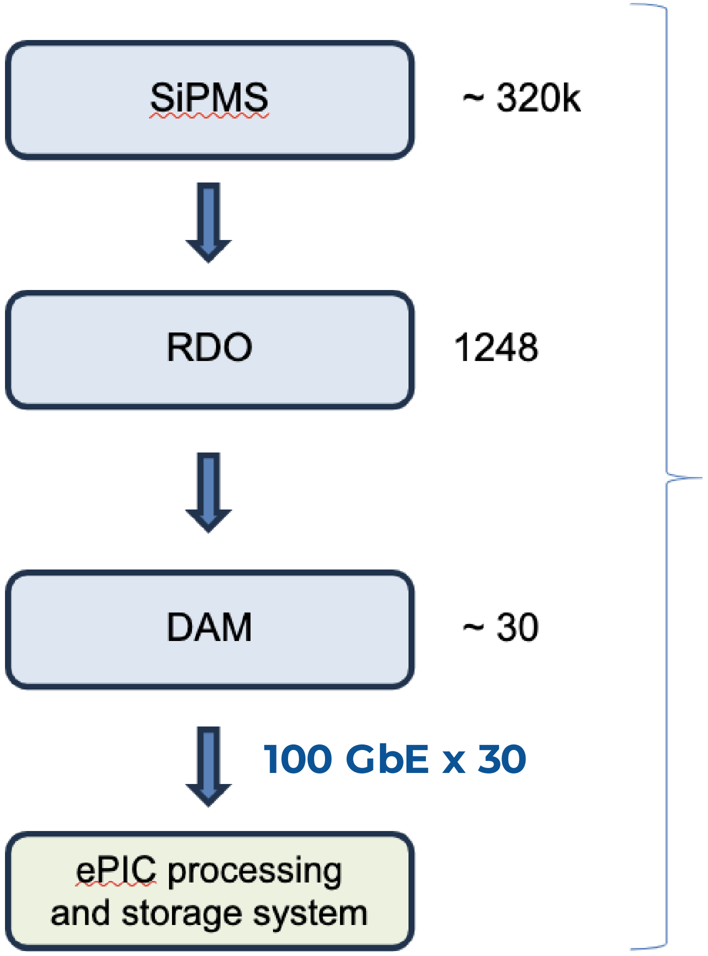

dRICH DAQ parameters	
RDO boards	1248
ALCOR64 x RDO	4
dRICH channels (total)	319488
Number of DAM L1	27
Input link in DAM L1	47
Output links in DAM L1	1
Number of DAM L2	1
Input link to DAM L2	27
Link bandwidth [Gb/s] (assumes VTRX+)	10
Interaction tagger reduction factor	1
Interaction tagger latency [s]	2,00E-03
EIC parameters	
EIC Clock [MHz]	98,522
Orbit efficiency (takes into account gap)	0,92

Bandwidth analysis		Limit
Sensor rate per channel [kHz]	300,00	4.000,00
Rate post-shutter [kHz]	55,20	800,00
Throughput to serializer [Mb/s]	34,50	788,16
Throughput from ALCOR64 [Mb/s]	276,00	
Throughput from RDO [Gb/s]	1,08	10,00
Input at each DAM I [Gbps]	50,67	470,00
Buffering capacity at DAM I [MB]	12,97	
Output from every DAM	50,67	10,00
Total throughput	1.368,14	270,00

- Sensors DCR: 3-300 kHz (increasing with radiation damage ➔ with experiment lifetime).
- Full detector throughput (FE): 14-1400Gbps
- **A reduction is needed** to match 30 channels aggregated bandwidth (and safety margin)
- EIC beams bunch spacing: 10 ns ➔ bunch crossing rate of 100 MHz
- For the low interaction cross-section (DIS) ➔ one interaction every ~100 bunches ➔ interaction rate of ~1MHz.
- **A system tagging the (DIS) interacting bunches** could solve the issue reducing down to ~1/100 the data throughput

Two complementary approaches are possible:

1. **Develop a dedicated sub-detector tagging relevant interactions.**
2. **This proposal.**

dRICH: Data Reduction

Online Signal/ Noise discrimination using ML

- **Signal (i.e. Merged Phys Signal + Bkg)**:
  - **Physics Signal:**
    - e.g DIS
  - **Physics Background:**
    - e/p with beam pipe
    - Synchrotron radiation (currently not including it)

- **SiPM Noise:**
  - Dark current rate (DCR) modelled in the reconstruction stage (*recon.rb* eic-shell method)

ML task:

Discriminate between **Noise Only** and **Signal + Noise** events

dRICH: Dataset for training, classes

Phys Signal+Phys Background+Noise

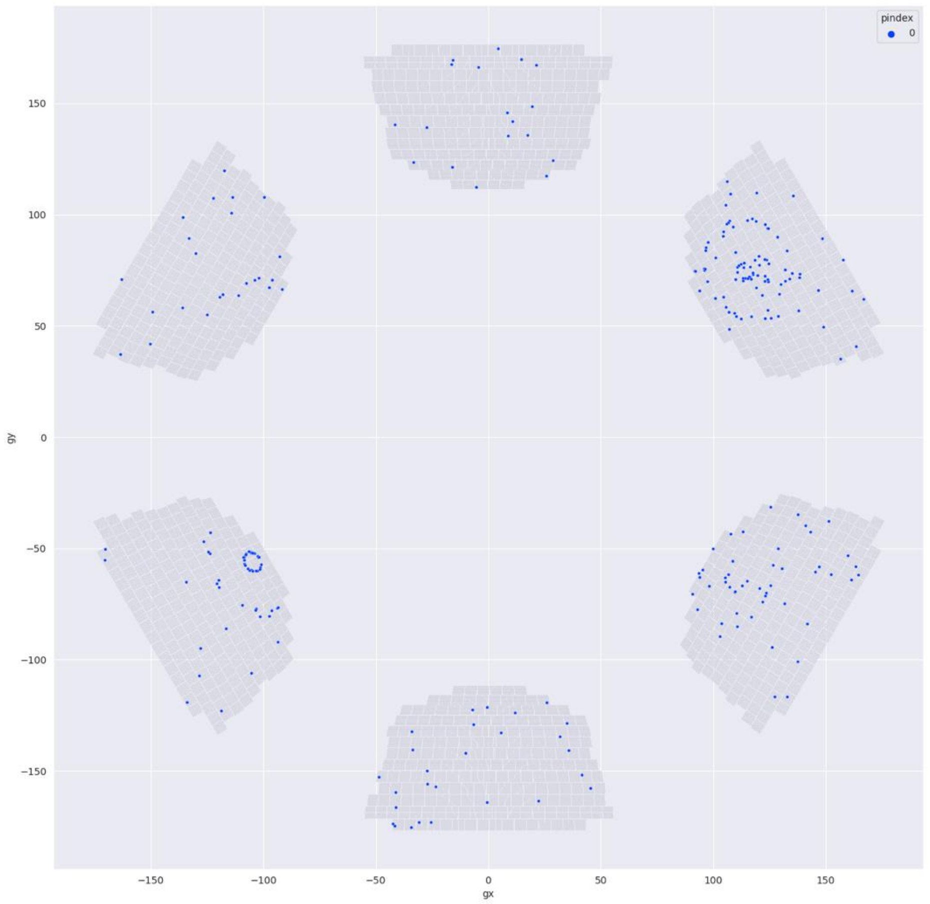

Noise Only

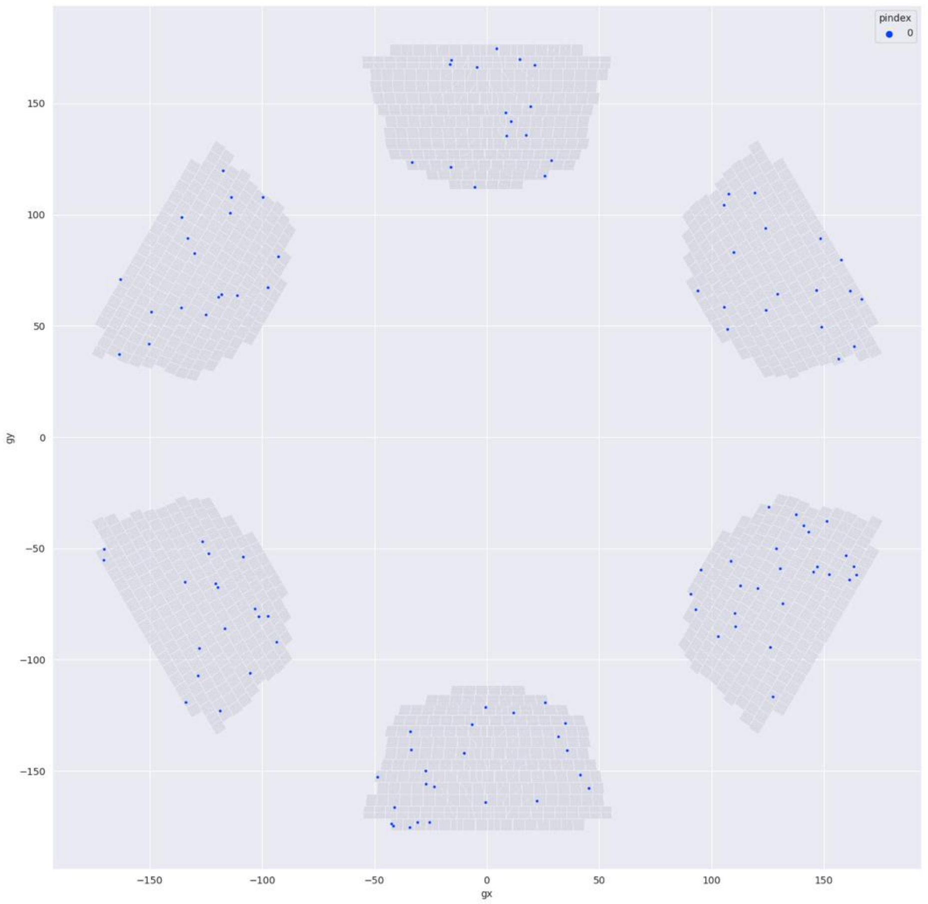

dRICH Data Reduction Stage on FPGA

- Online «Noise only» classifier using ML
  - Study of Inference Models
    - Restricting our study to inference models that can be deployed on FPGA with reasonable effort (using a High-Level Synthesis workflow)
      - MLP, CNN, GNN Models (HLS4ML)
    - Inference throughput (98.5 MHz) is the main challenge.
    - HDL optimized implementation is an option.
    - Not necessarily ML-based…
- Deployment on multiple Felix DAMs and on an additional FPGA (TP – Trigger Processor) directly interconnected with the APE communication IP.
- Possibly integrate with the **dRICH Interaction Tagger** to boost performance

dRICH DAQ

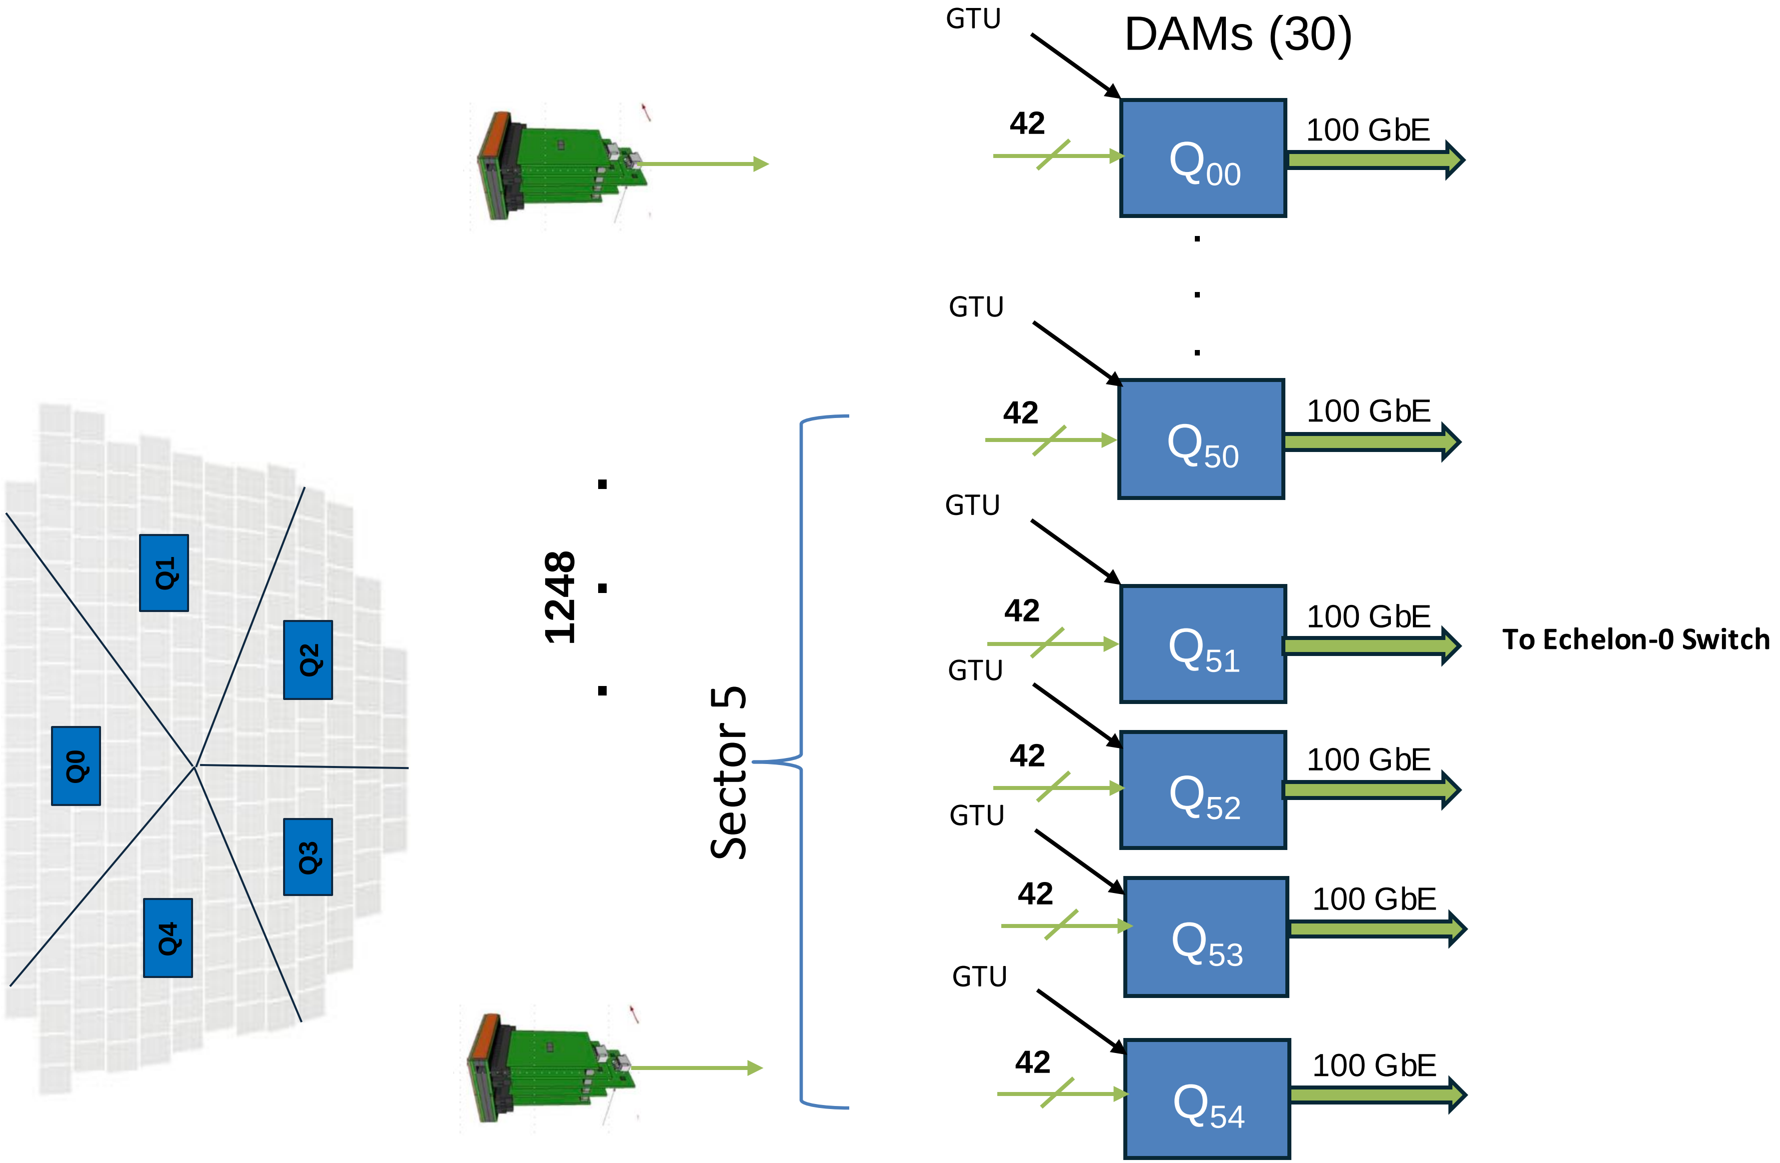
DAMs (30)
GTU
42
Q00
100 GbE
GTU
42
Q50
100 GbE
GTU
42
Q51
100 GbE
To Echelon-0 Switch
GTU
42
Q52
100 GbE
GTU
42
Q53
100 GbE
GTU
42
Q54
100 GbE
1248
Sector 5
Q0
Q1
Q2
Q3
Q4

dRICH DAQ & Data Reduction

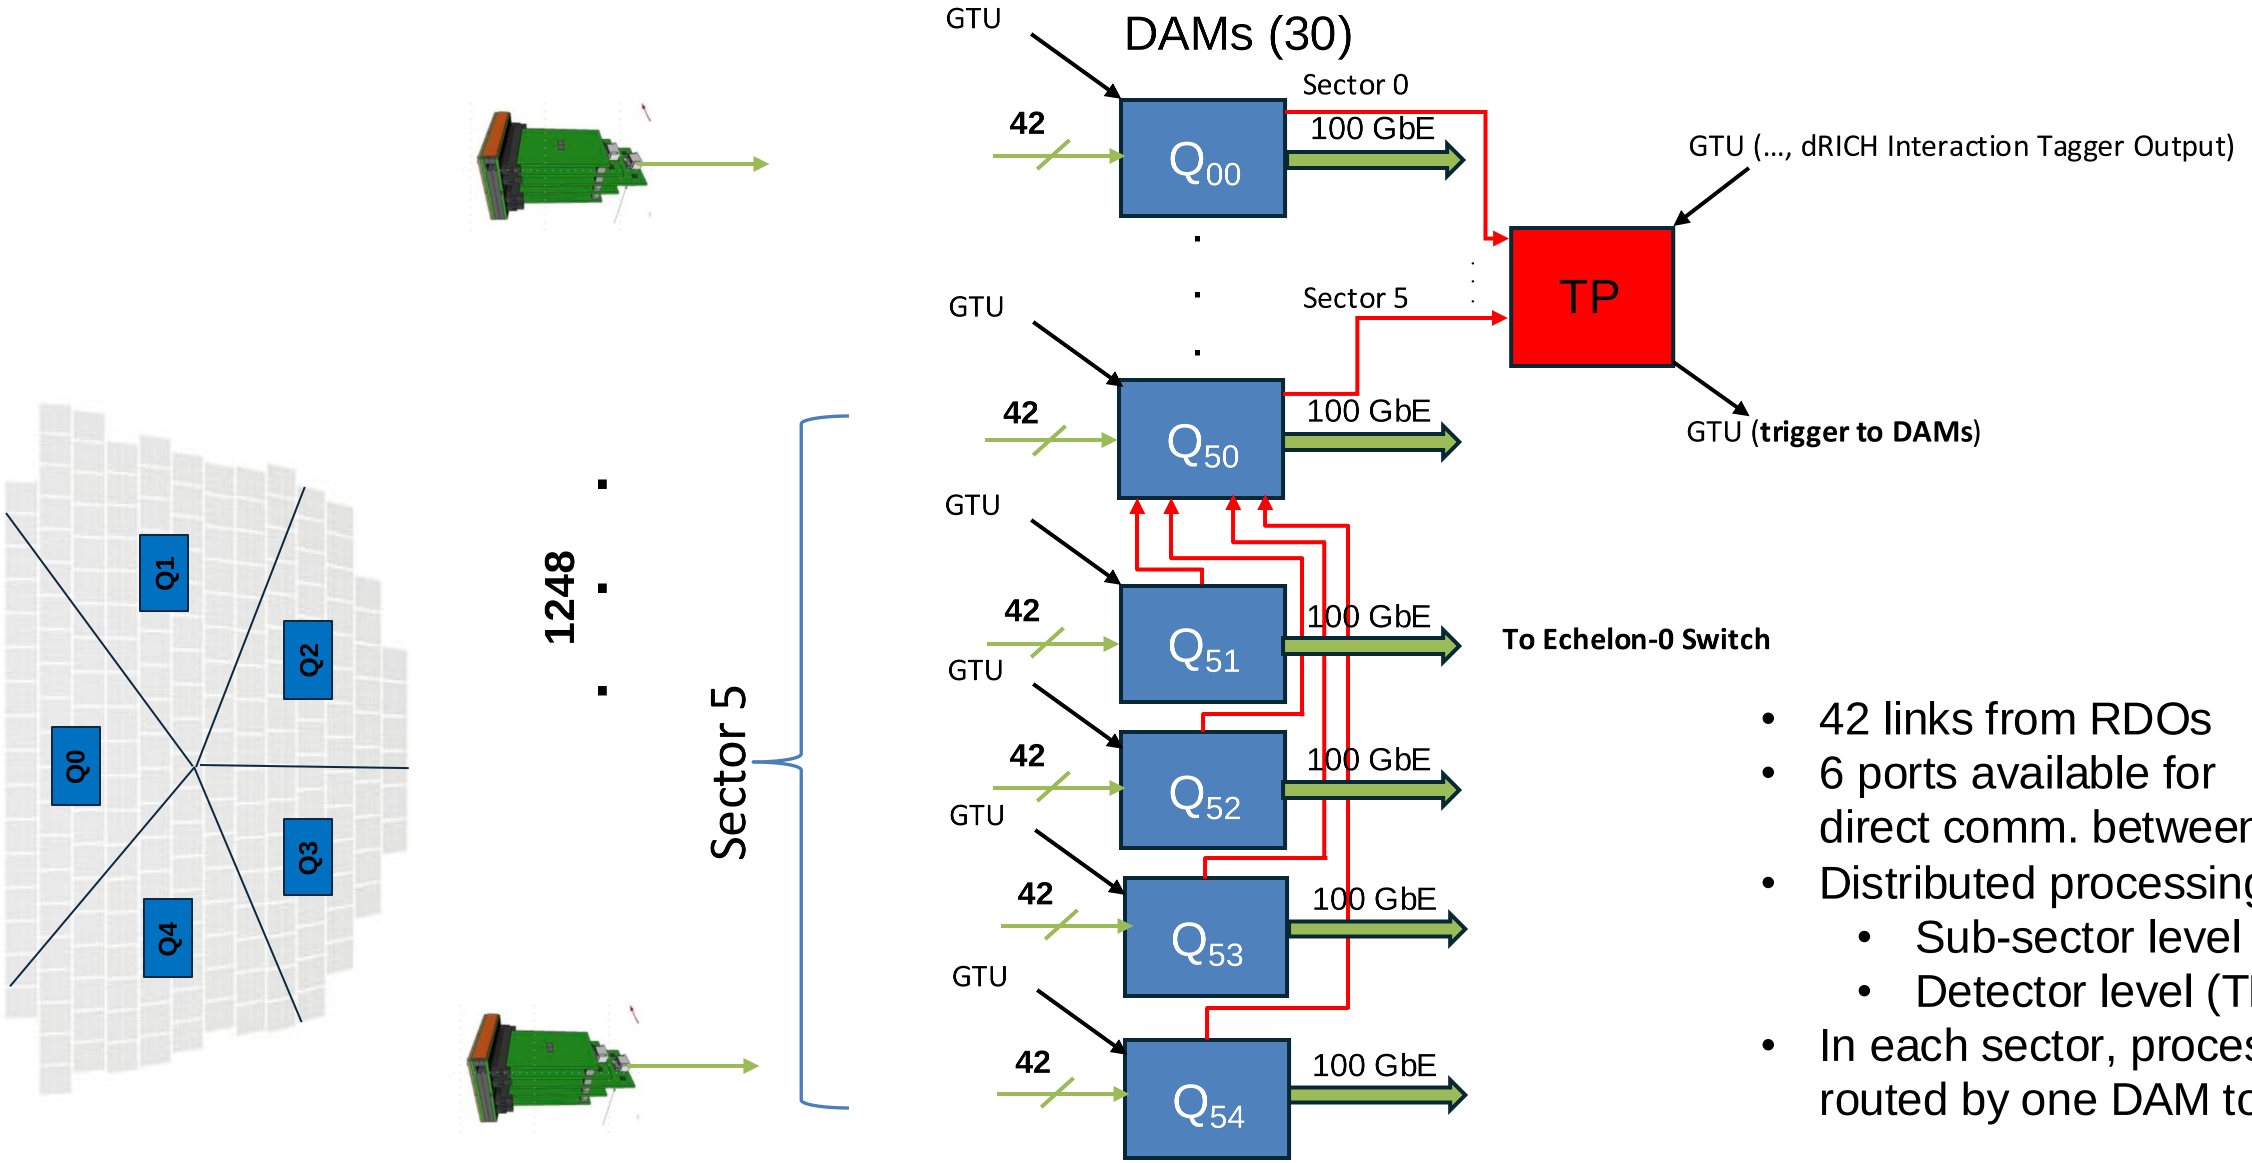

- 42 links from RDOs
- 6 ports available for direct comm. between DAMs
- Distributed processing on FPGAs
  - Sub-sector level (DAM)
  - Detector level (TP)
- In each sector, processed data routed by one DAM to TP

Some Background Activities

- INFN APE Lab @ Roma1/2: design and development of 4 generations of parallel computing architectures (mainly) dedicated to LQCD (1986-2010) **https://apegate.roma1.infn.it**

- Two recent research activities are relevant for this presentation:
  - APEIRON: a framework offering hardware and software support for the execution of real-time dataflow applications on a system composed by interconnected FPGAs. [https://doi.org/10.1051/epjconf/202429511002]
  - FPGA-RICH: online ring counting system based on FPGA for the RICH detector of the NA62 experiment at CERN. In publication.

- Other research activities of possible interest
  - APENet: a high-throughput network interface card based on FPGA used in hybrid, GPU-accelerated clusters with a 3D toroidal mesh topology. [http://doi.org/10.1088/1742-6596/898/8/082035]
  - NaNet: a family of FPGA-based PCIe Network Interface Cards (with GPUDirect/RDMA capability) for High Energy Physics to bridge the front-end electronics and the software trigger computing nodes. [https://doi.org/10.1088/1742-6596/1085/3/032022]

dRICH DAQ & Data Reduction

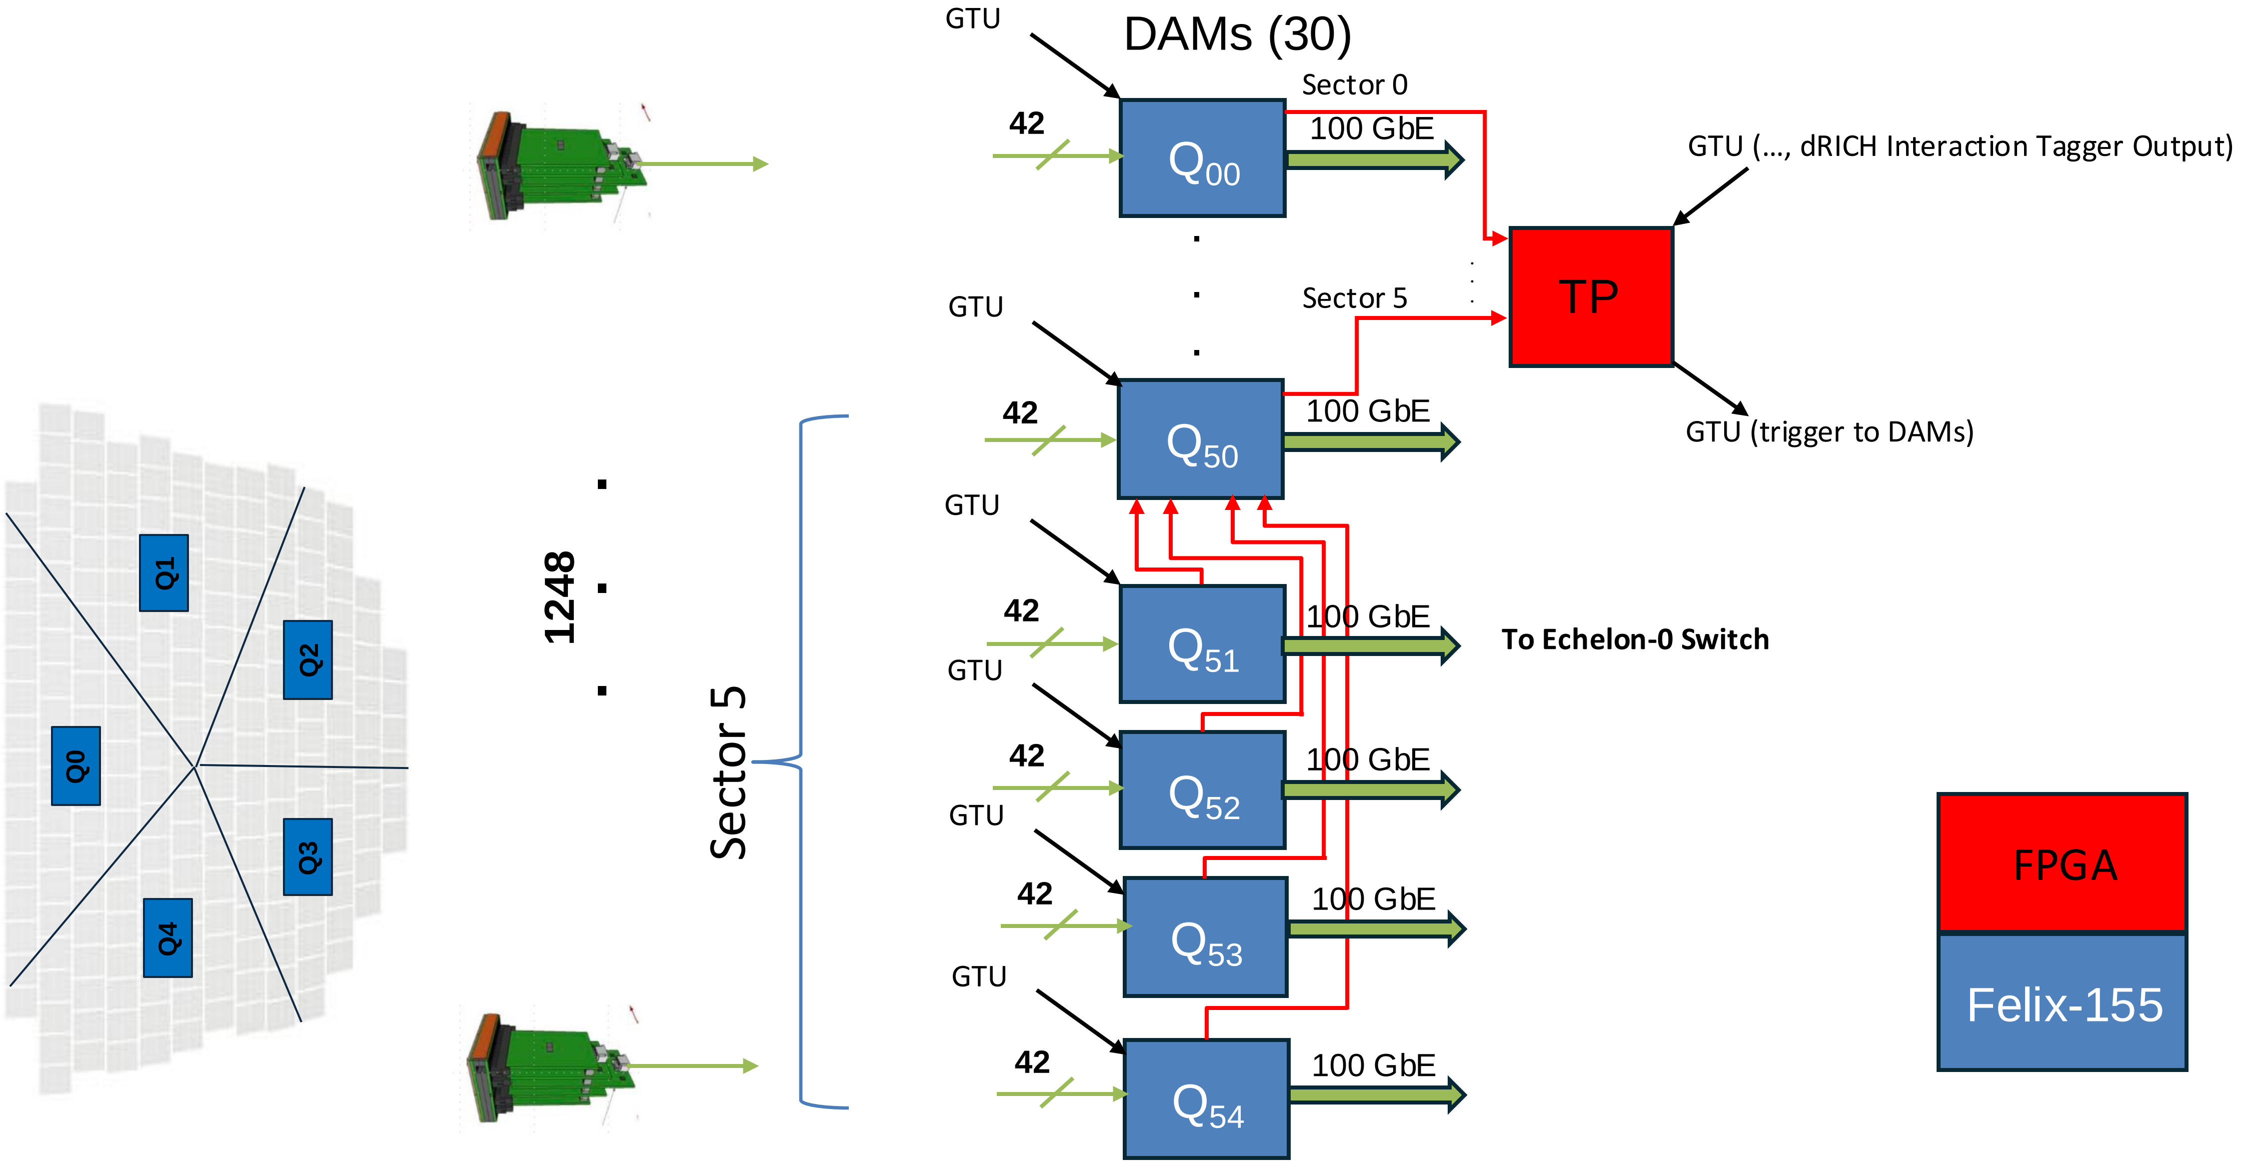

dRICH: Data reduction ➔ Subsectors

each subsector readout information discretized to a **8x8 grid ➔ 64 inputs to NN**

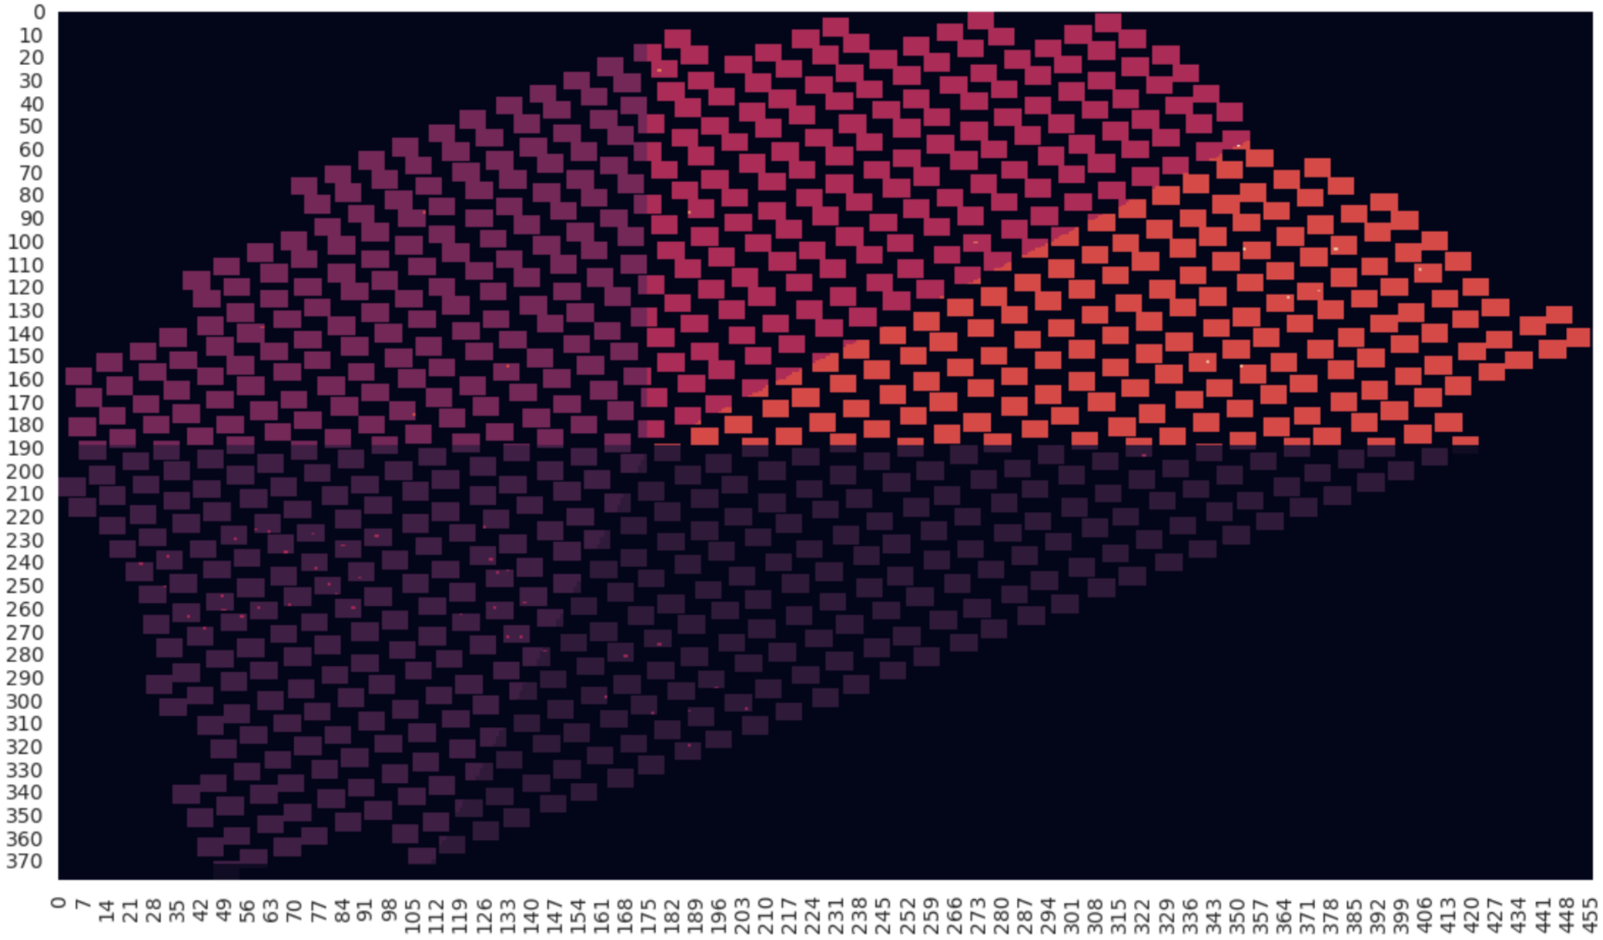

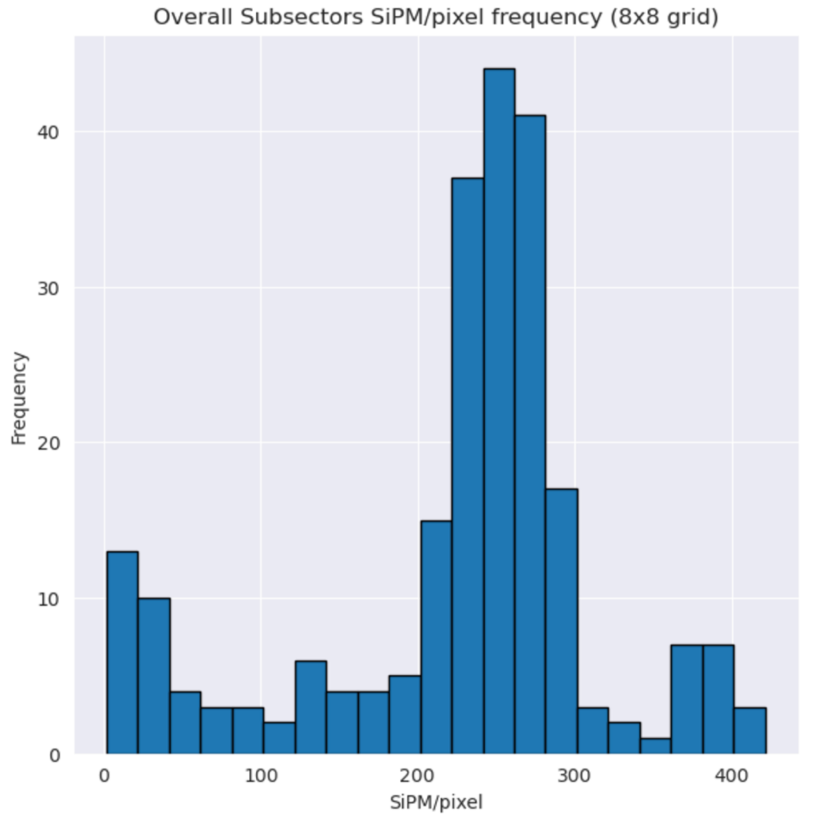

dRICH Data Reduction on FPGA - Deployment

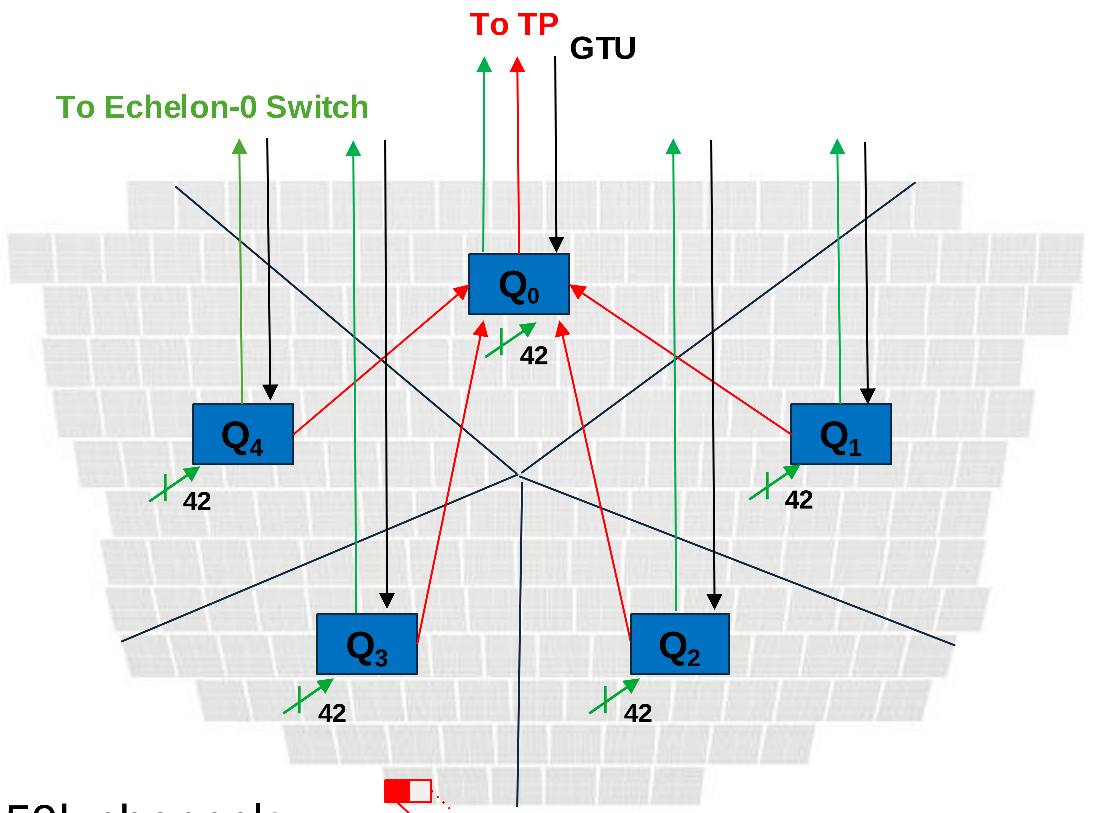

~50k channels
per sector

RDO flow
PROC flow
GTU

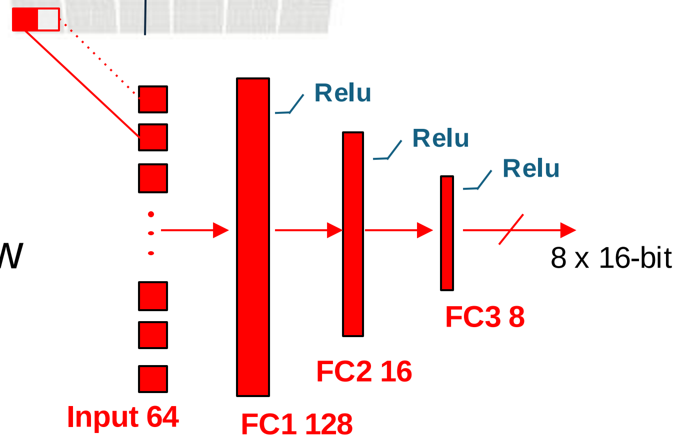

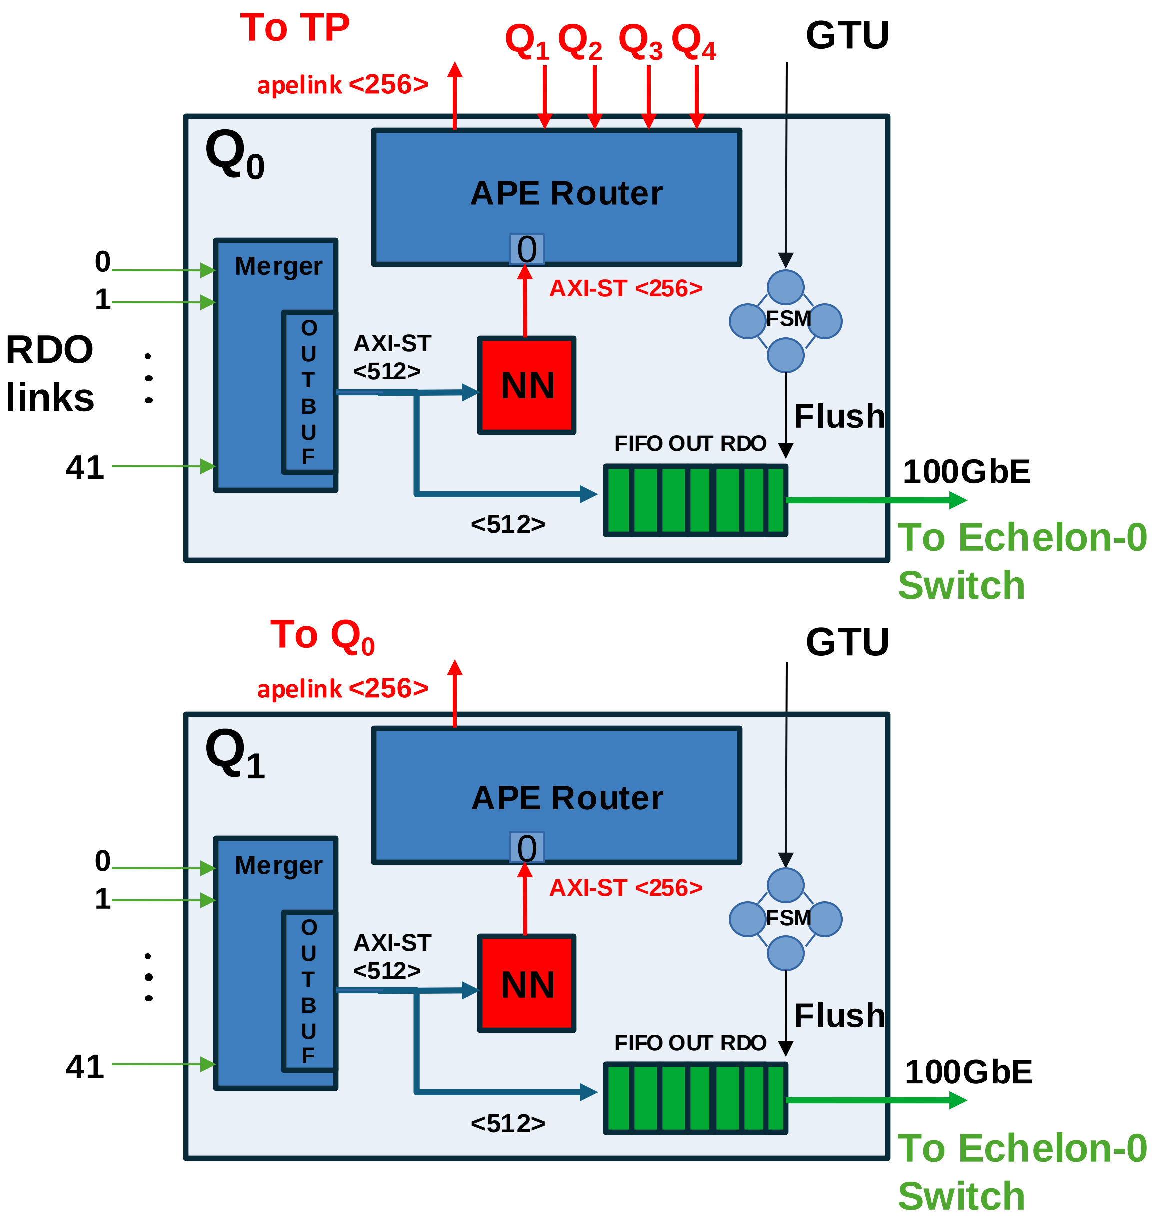

dRICH Data Reduction on FPGA - Deployment

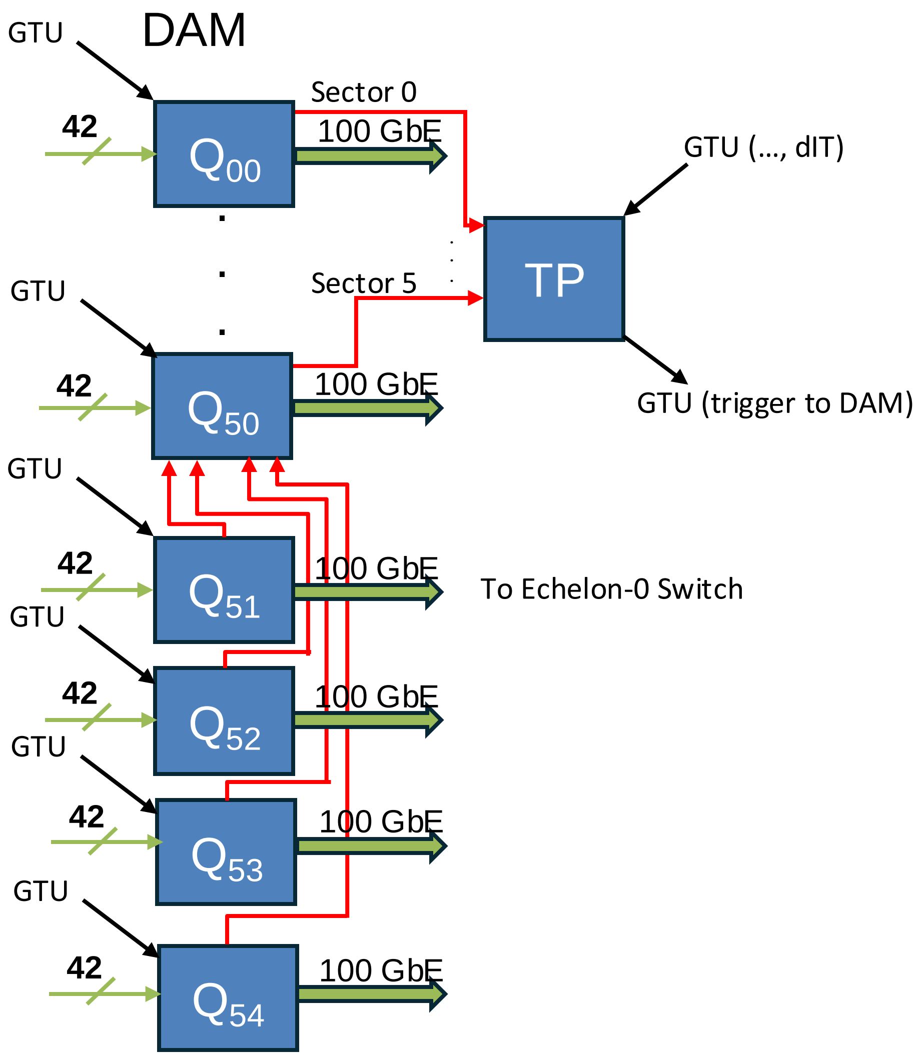

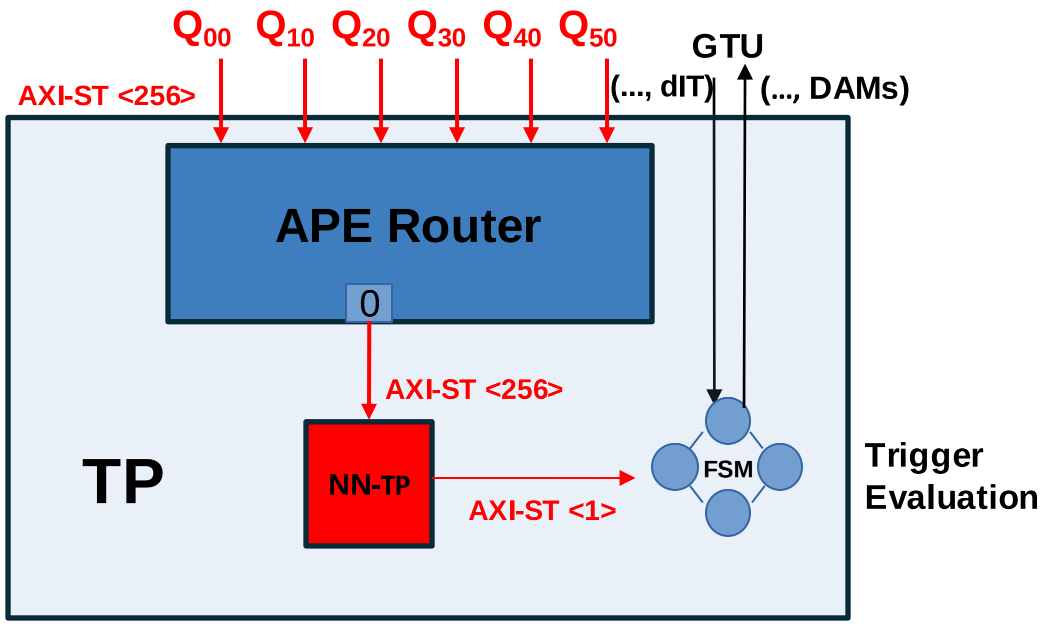

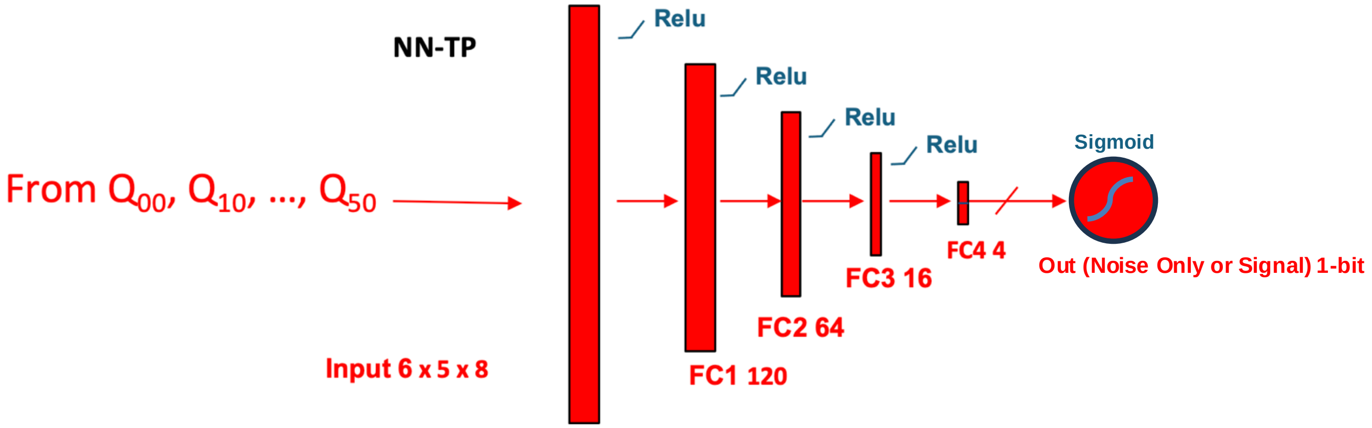

dRICH: Data reduction Dataset

ePIC simulation pipeline:

HEPMC
MC events

PHYS signal e.g. DIS

PHYS bkg e - beamgas
p - beamgas

HEPMC_merger

```
./SignalBackgroundMerger --signalFile dis.hepmc --signalFreq 0 \
        --bg1File hgas.hepmc --bg1Freq 31.9 \
        --bg2File egas.hepmc --bg2Freq 3177.25 \
```

FULL
Geant4

Geant4 hits on detectors sensitive volume

root://dtn-eic.jlab.org//work/eic2/EPIC/FULL/24.05.0/epic_craterlake/BACKGROUNDS/MERGED/HEPMC_merger-1.0.2/10x100/RealisticSignalPerFrame/

RECO
EICRecon

In dRICH case, applies digitization, quantum efficiency, safety factor. Possibly SiPM noise

By default SiPM noise disabled and not present on root://dtn-eic.jlab.org

Options:

- ➢ Start from Merged FULL root files available on server and enable noise at RECO stage using *drich-dev/recon.rb* with configs (but only ~ 7k events present on dtn-eic)
- ➢ **Run the entire simulation pipeline ourselves, starting from HEPMC files.**
  - ○ Up to now we have produced 600k events to train and test our ML models.

dRICH Data reduction:
Input Data (Features Definition)

- **Gaussian** dark current SiPM **noise hits distribution,** obtained by modifying EICRecon source:
- avg = noiseRate*noiseTimeWindow
- sigma = 0.1*avg
- noiseTimeWindow = 2 ns

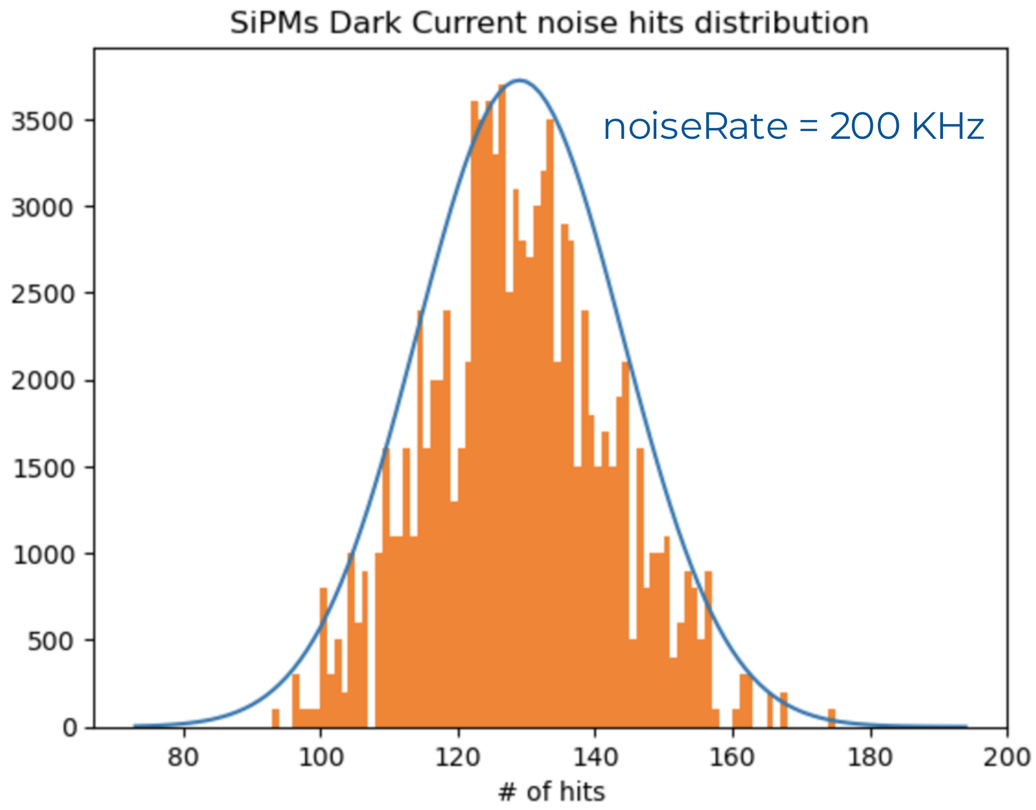

Distribution of Events Particles Momenta

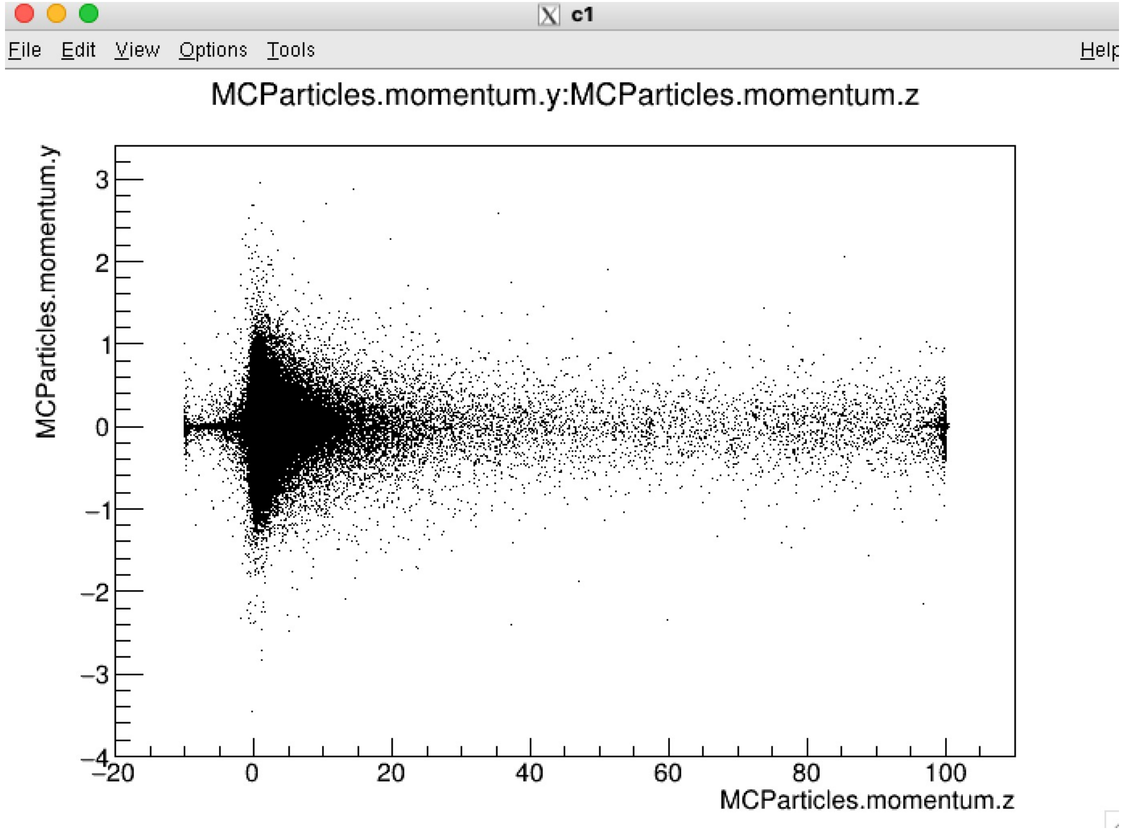

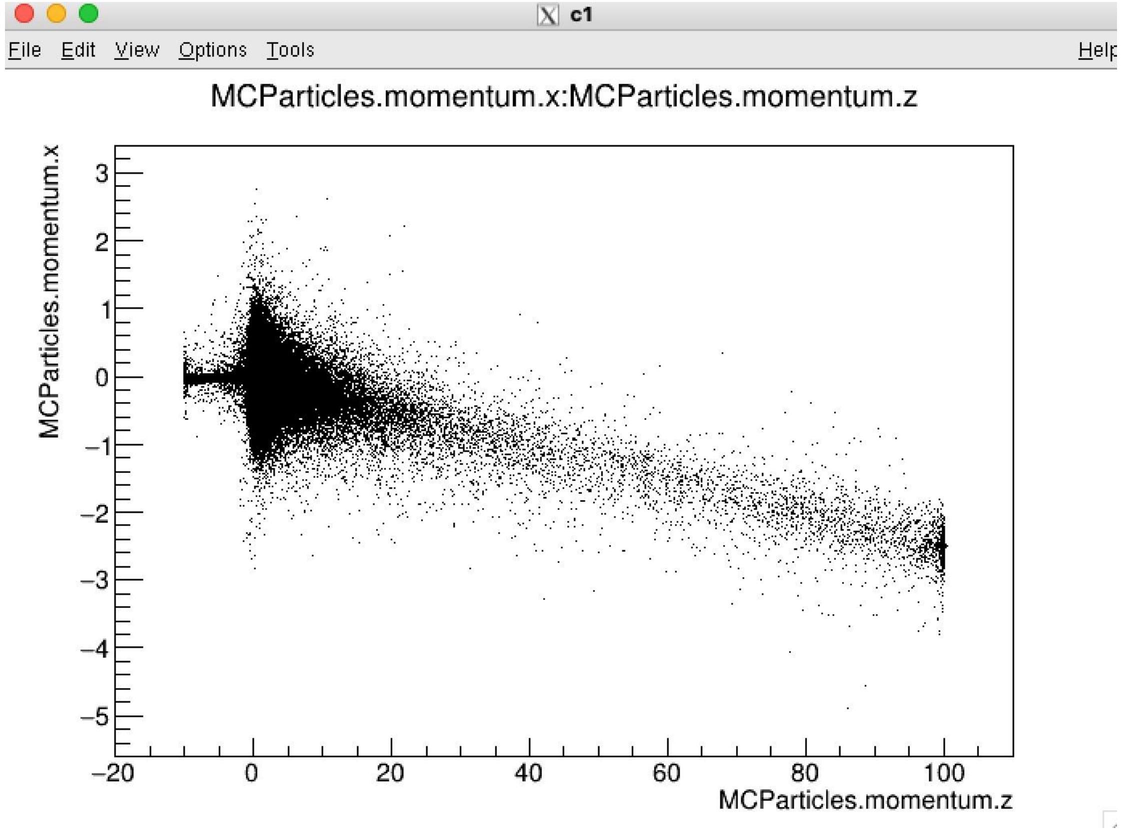

dRICH Data reduction:
Input Data (Features Definition)

- **Signal+Background+Noise**

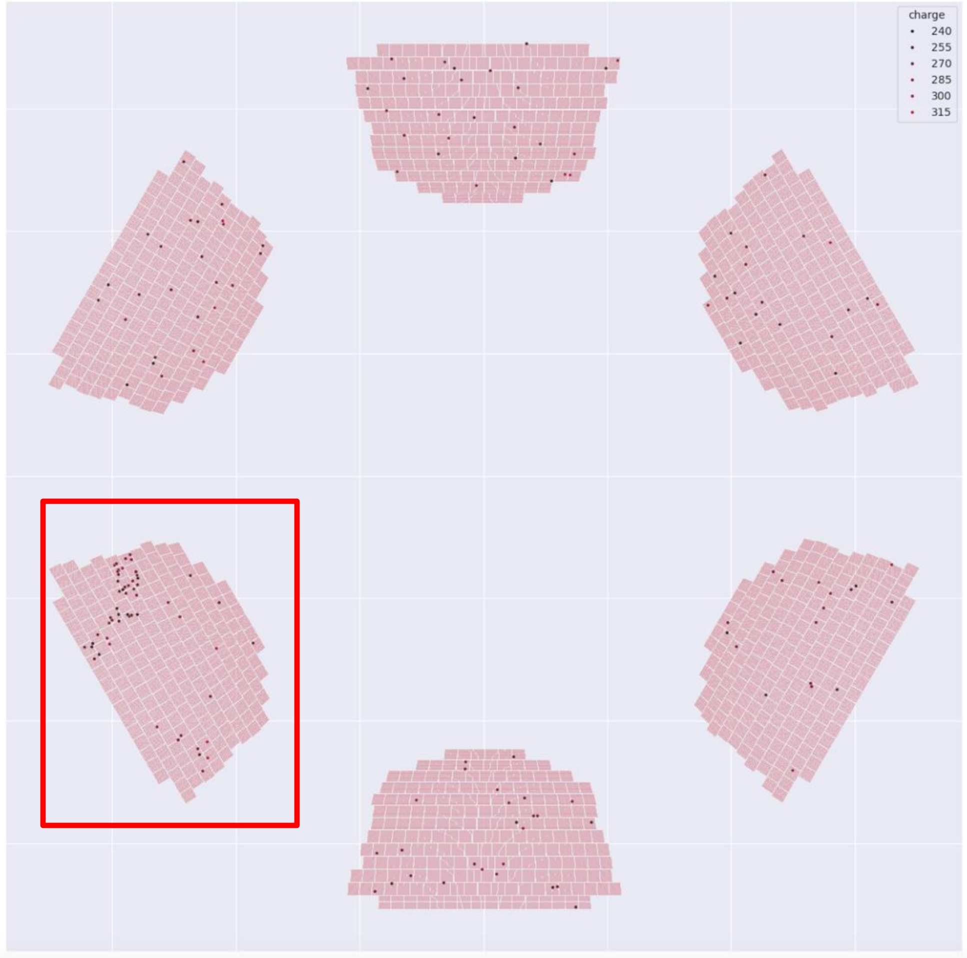

- **Noise Only**

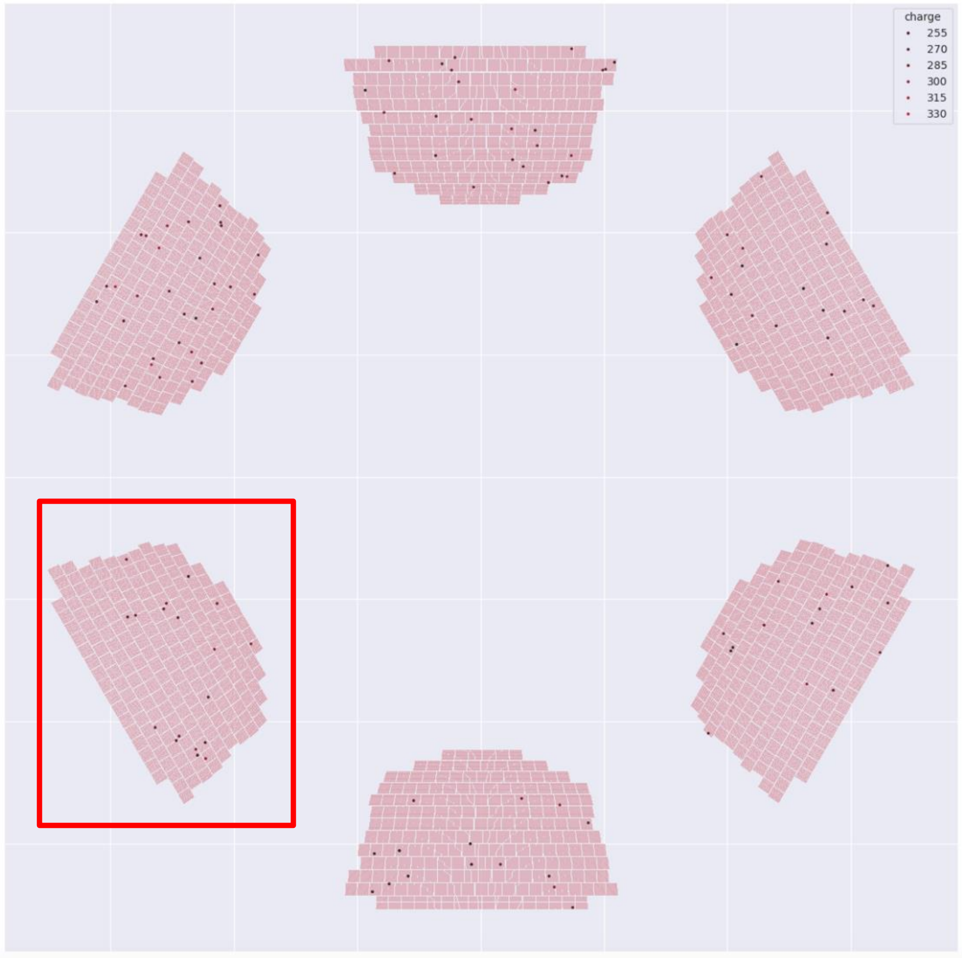

dRICH Data Reduction: Input Data

- **Signal+Background+Noise**

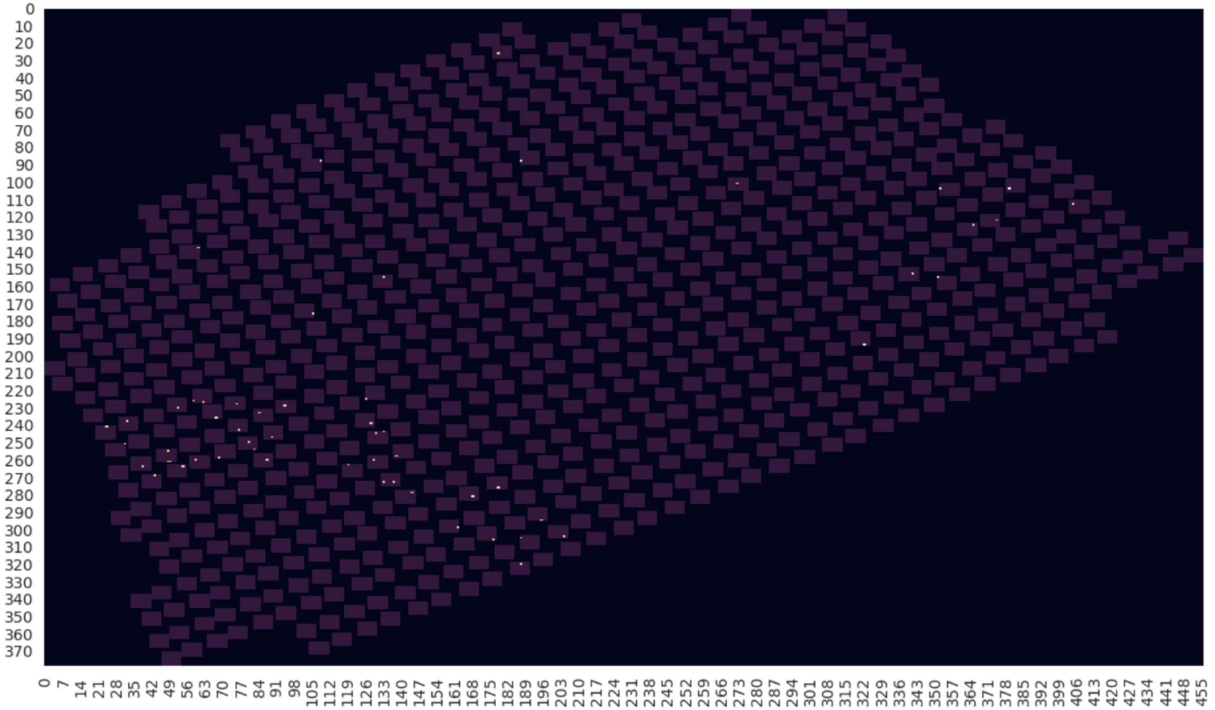

- **Noise**

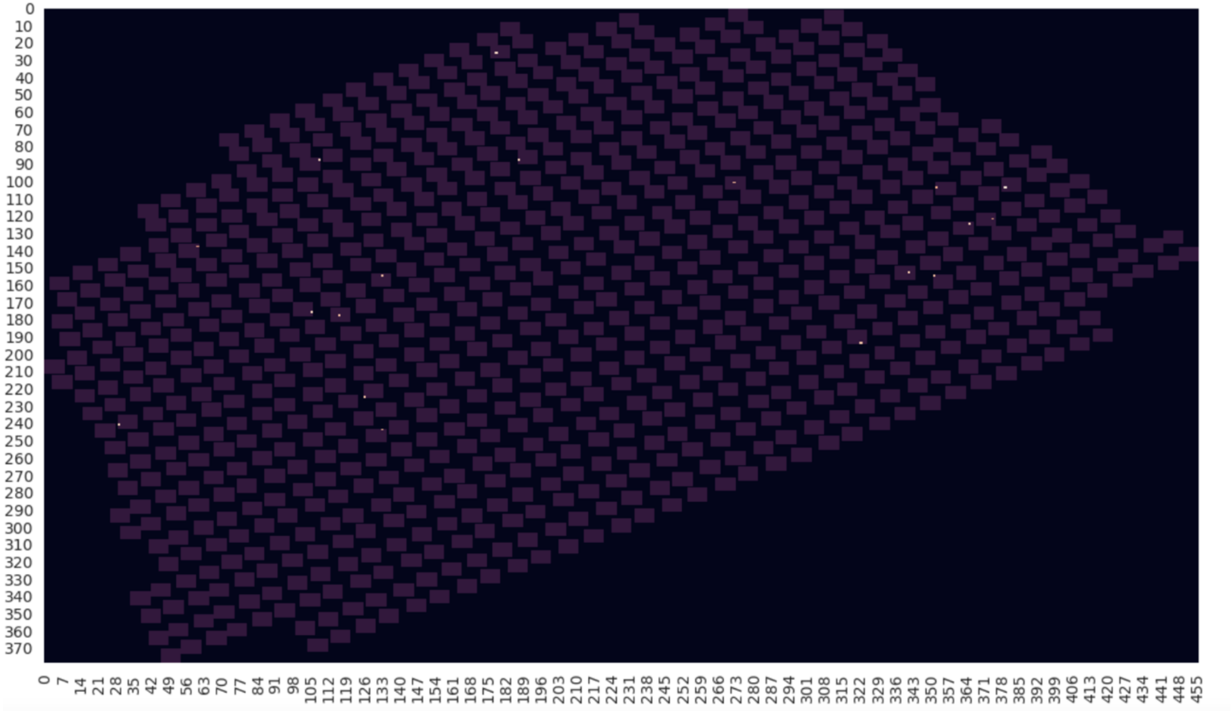

dRICH Data Reduction: 8x8 Grid

- ➢ **Signal+Background+Noise**

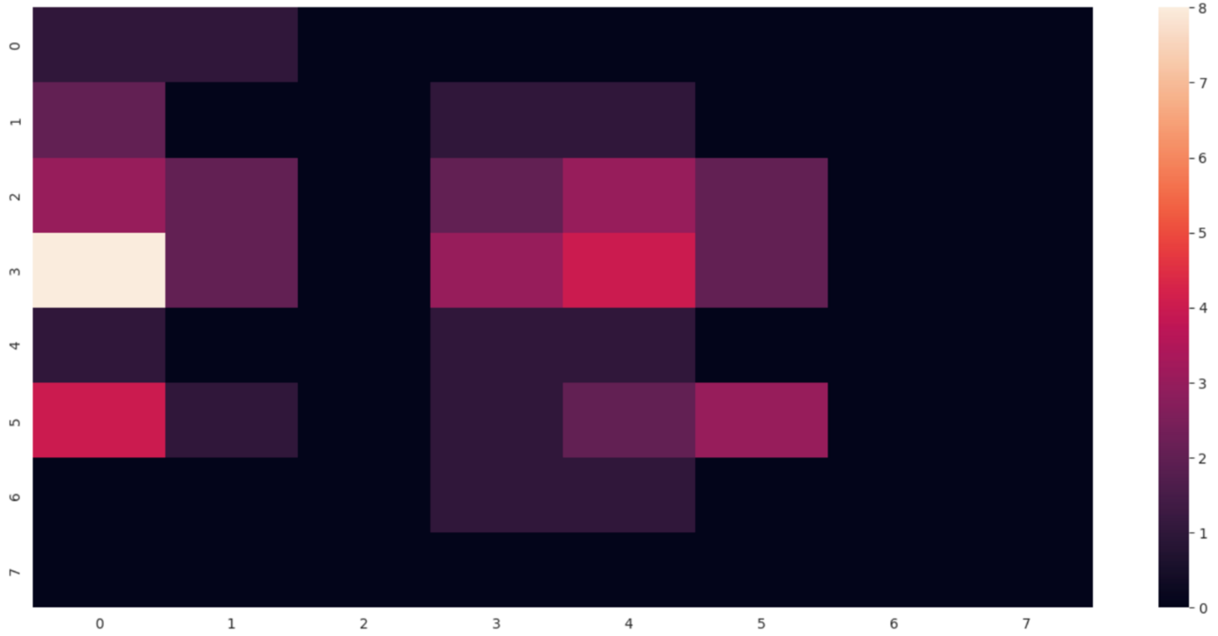

- ➢ **Noise**

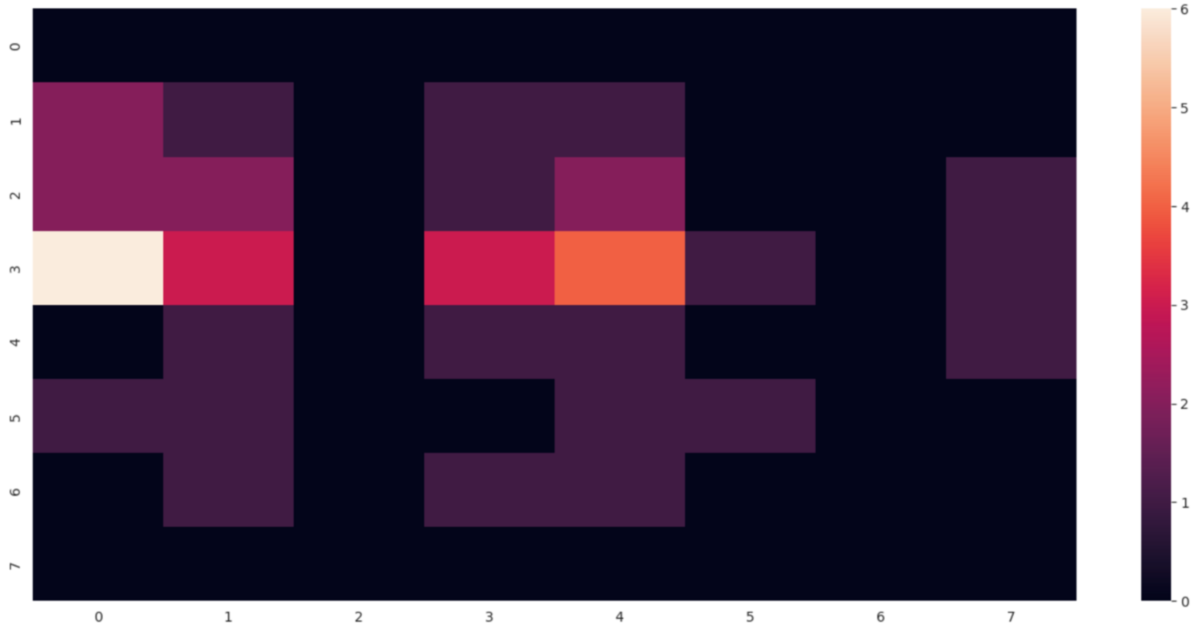

8x8 Grid ➔ 64 input NN neurons

dRICH Data reduction: Tensorflow-Keras Model definition

- To be coherent with the hardware design composition of the proposed system, we trained **30** (# of subsectors x #number of sectors) **concatenated MLP networks** into a single MLP model to be deployed on 30 DAM FPGAs + 1 TP FGPA

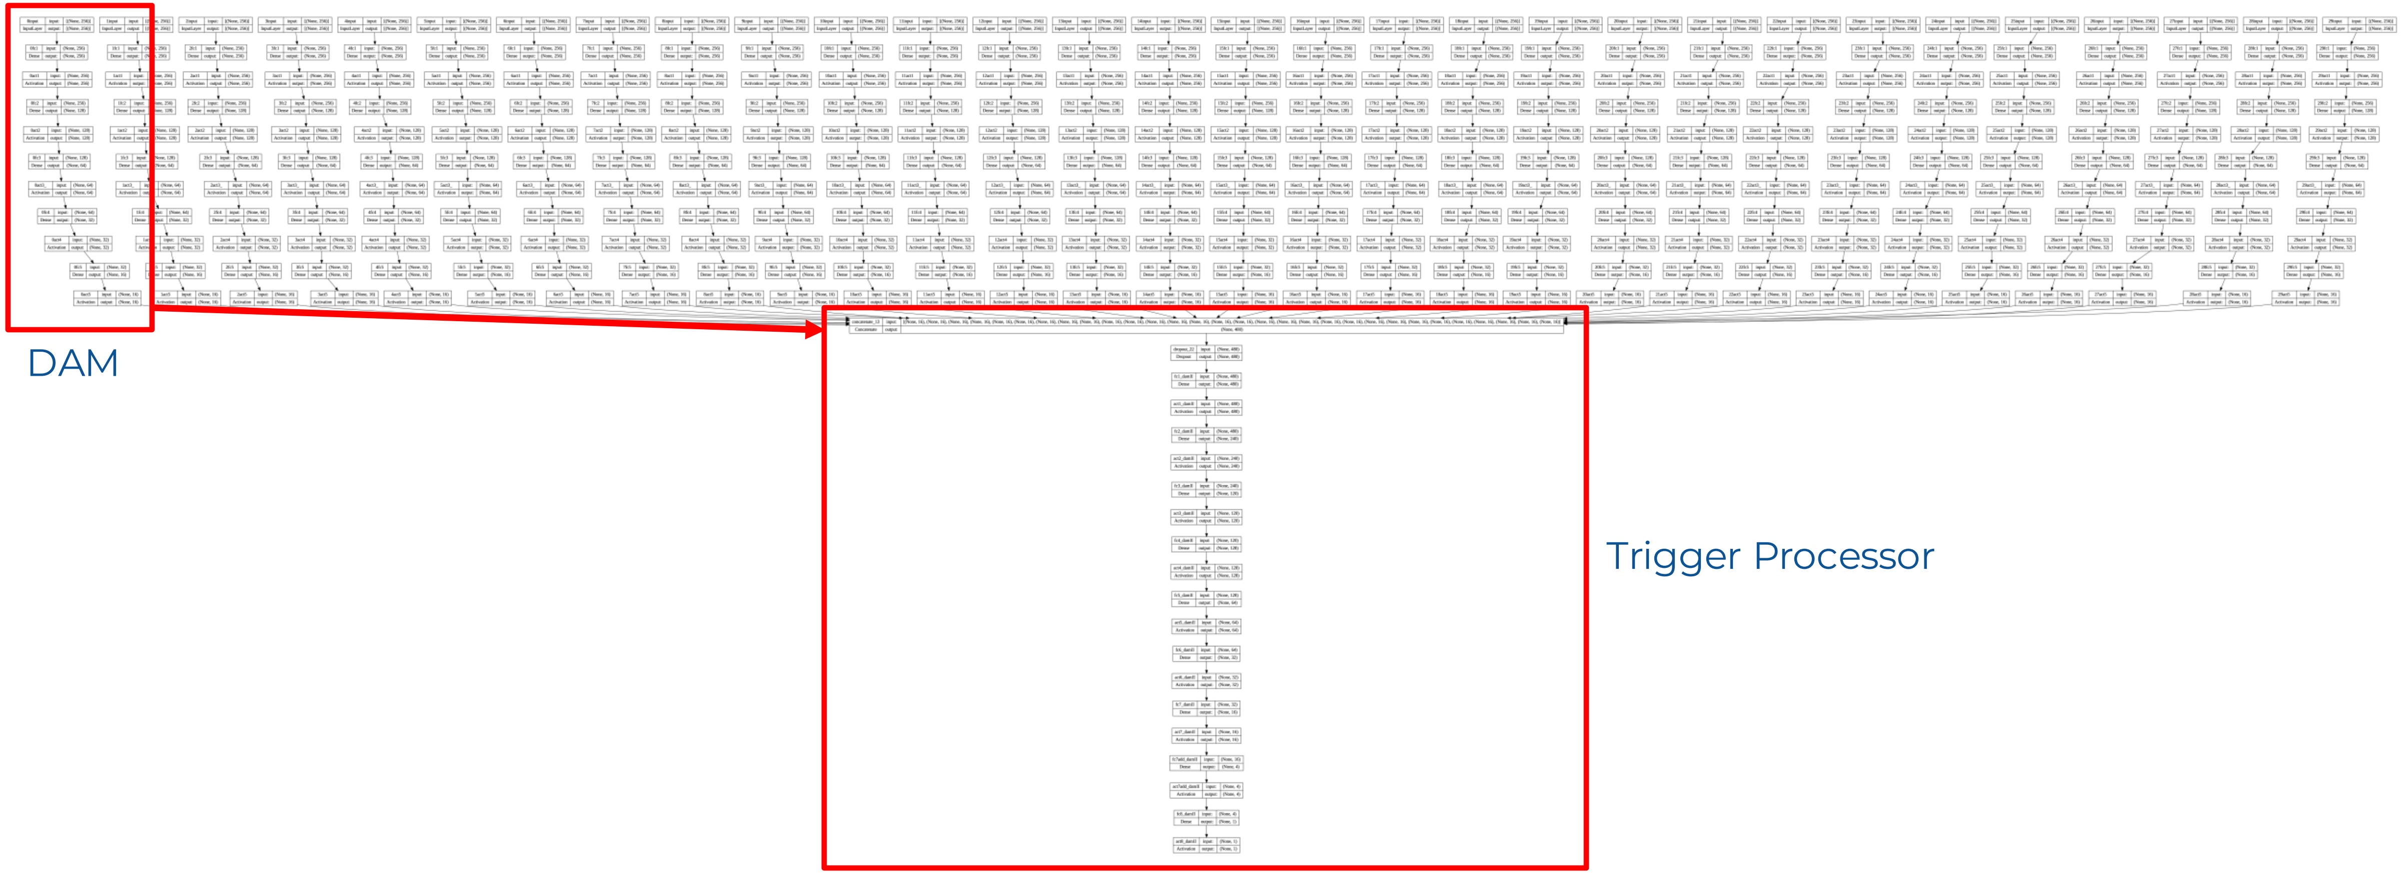

dRICH Data reduction: Tensorflow-Keras Model definition

- To be coherent with the hardware design composition of the proposed system, we trained **30** (# of subsectors x #number of sectors) **concatenated MLP networks** into a single MLP model to be deployed on 30 DAM FPGAs + 1 TP FGPA

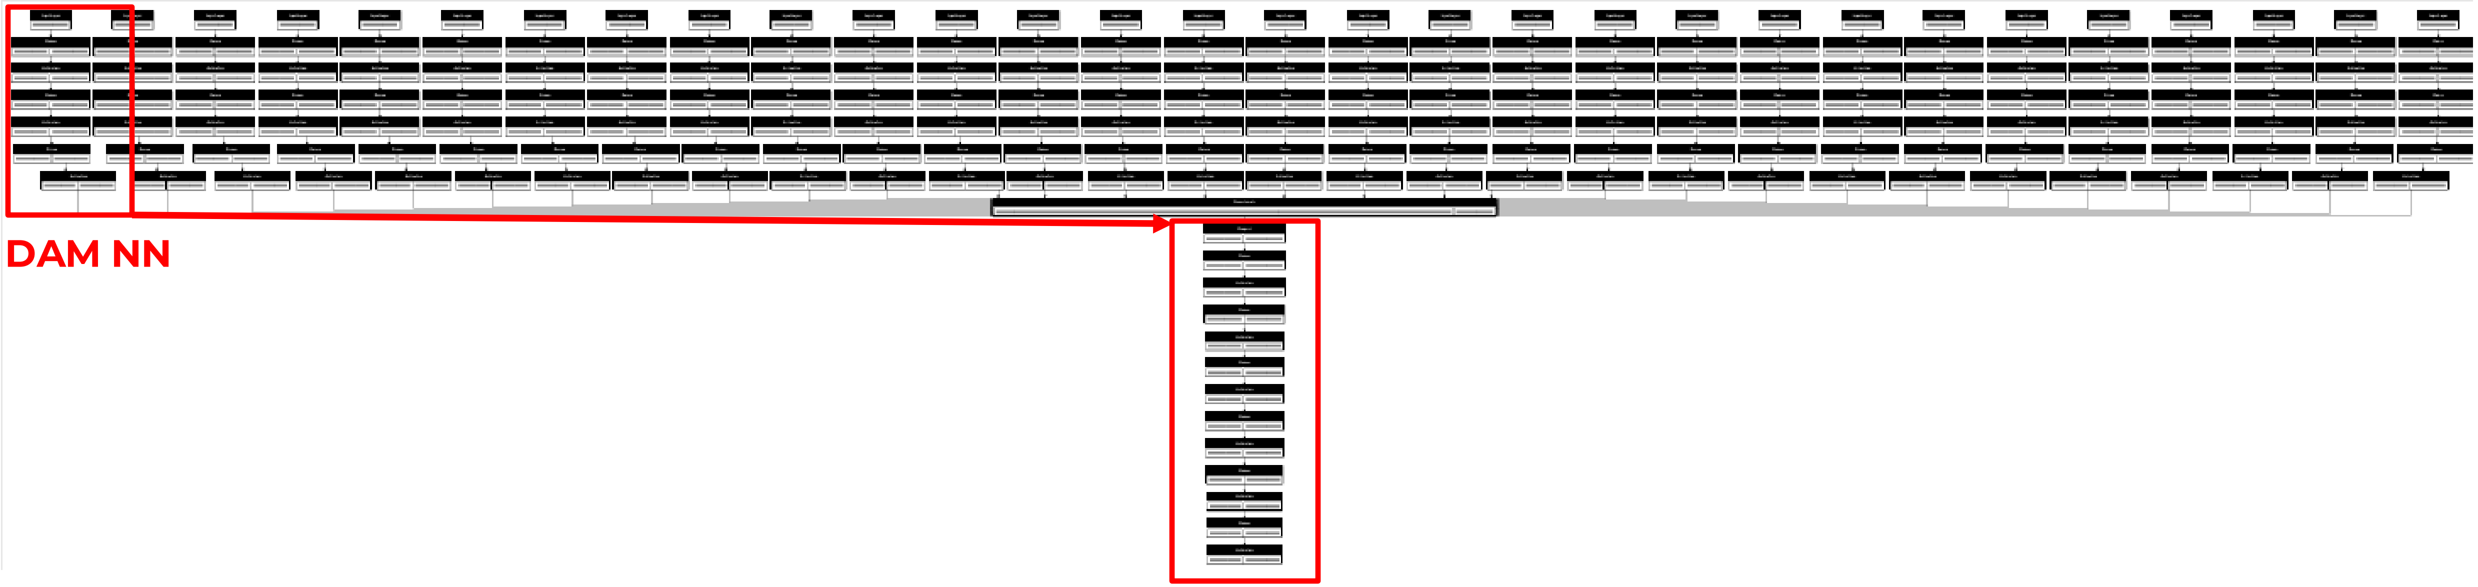

«Distributed MLP Model»

Distributed MLP Tensorflow Model

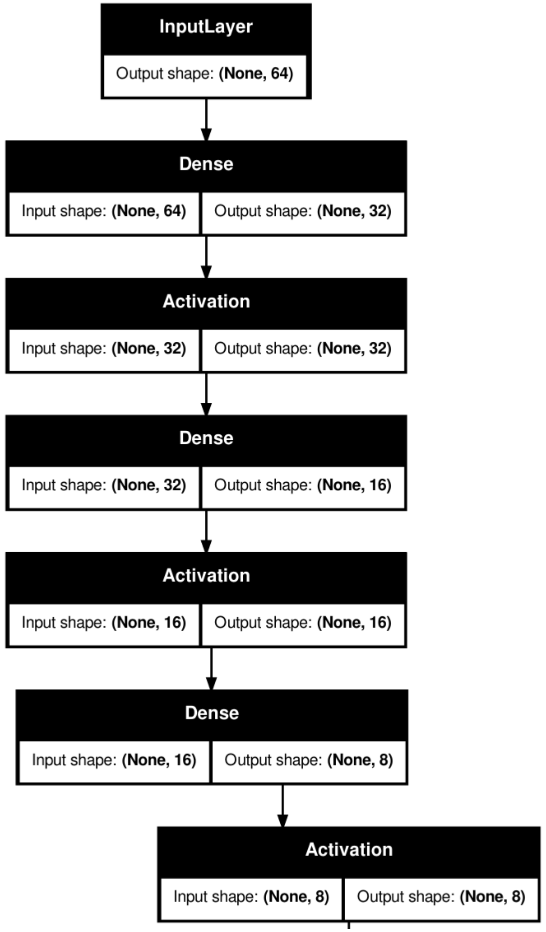

Each MLP DAM output (**embedding**) is concatenated to the others to feed the final stage of the MLP (deployed on TP)

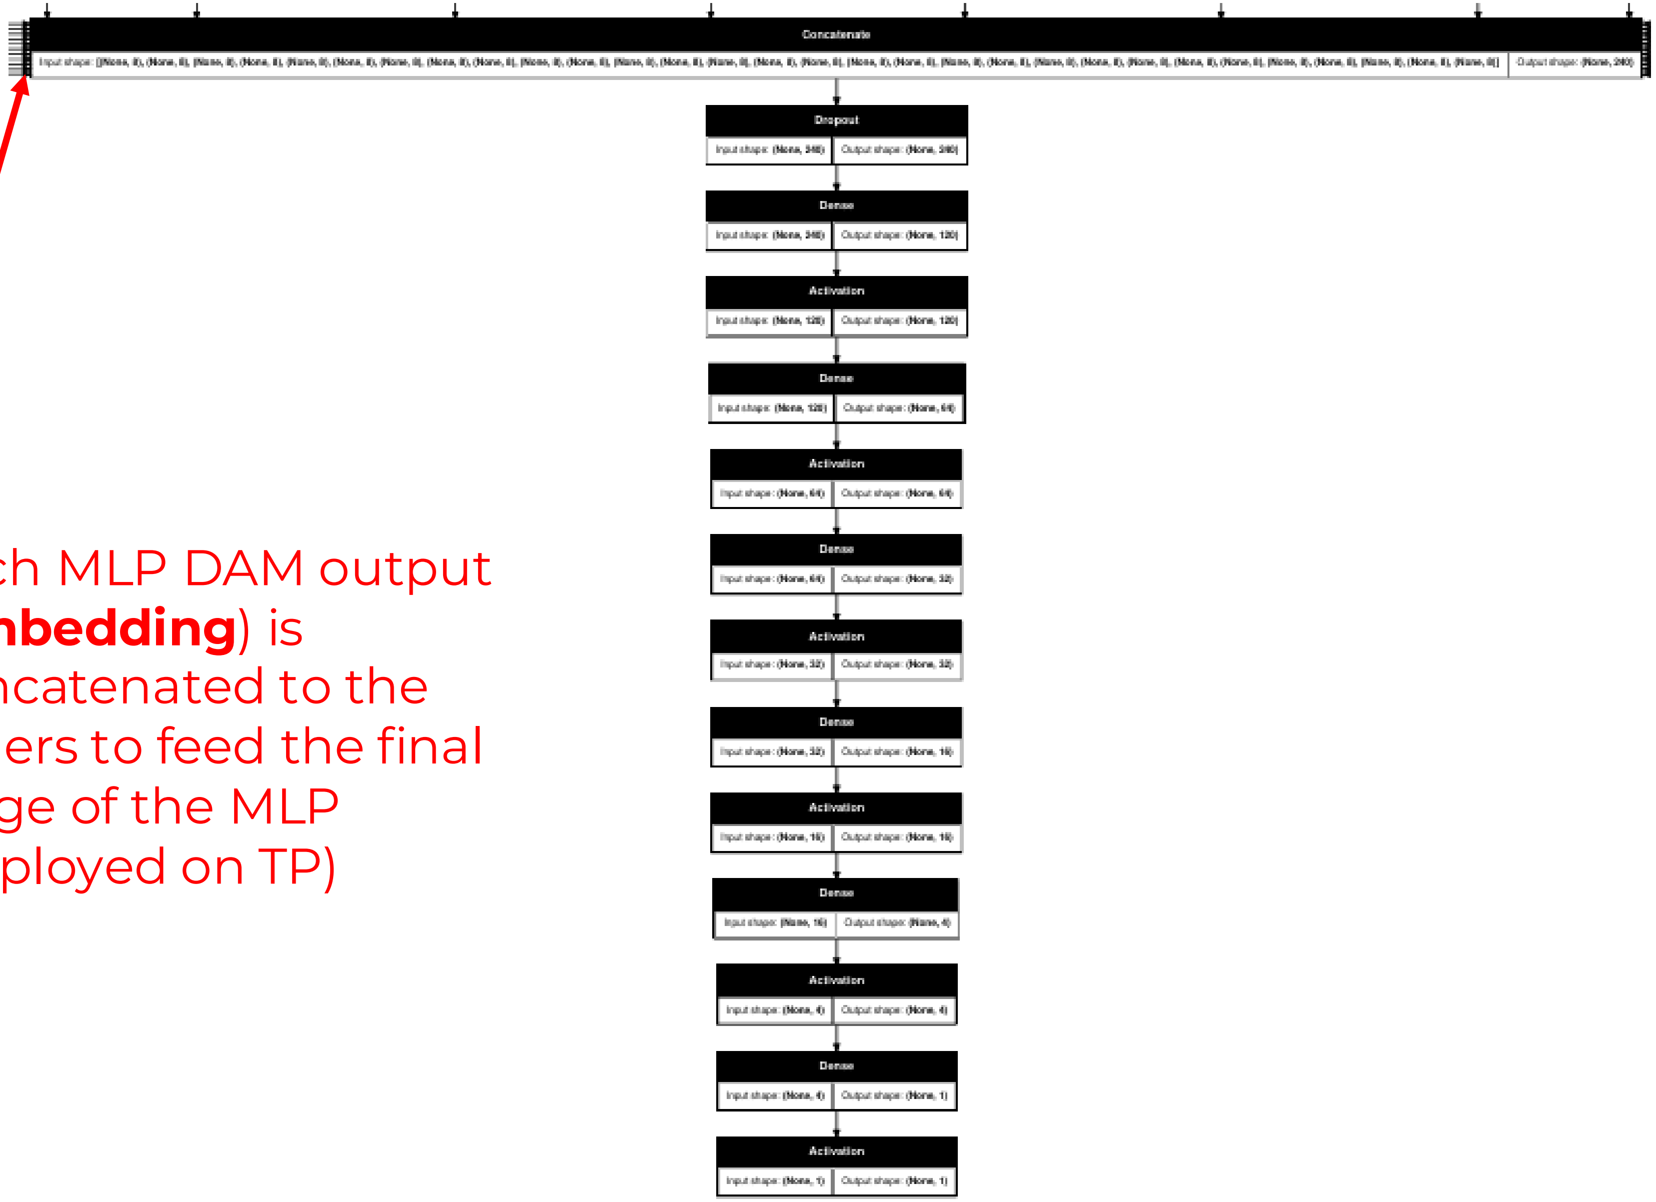

dRICH Data reduction: model training & validation

- We trained the 30 MLP DAM models concatenated to the single MLP TP model by using 100k Signal+Background+Noise and 100k Noise Only event
- **200k balanced dataset** (90% training set, 10% validation set) for any of the considered values of noiseRate (100 KHz, 200 KHz, 400 KHz)
- We minimize a typical Binary CrossEntropy loss function in 1000 epochs, **backpropagating** the result to all the input models ➔ in this way, trained 30 MLP DAM models result are **uncorrelated**, coherently with the target design in which each susbector NN is oblivious to the others subsector NNs
- Training and validation has been repeated after quantization

Model training & validation: Loss

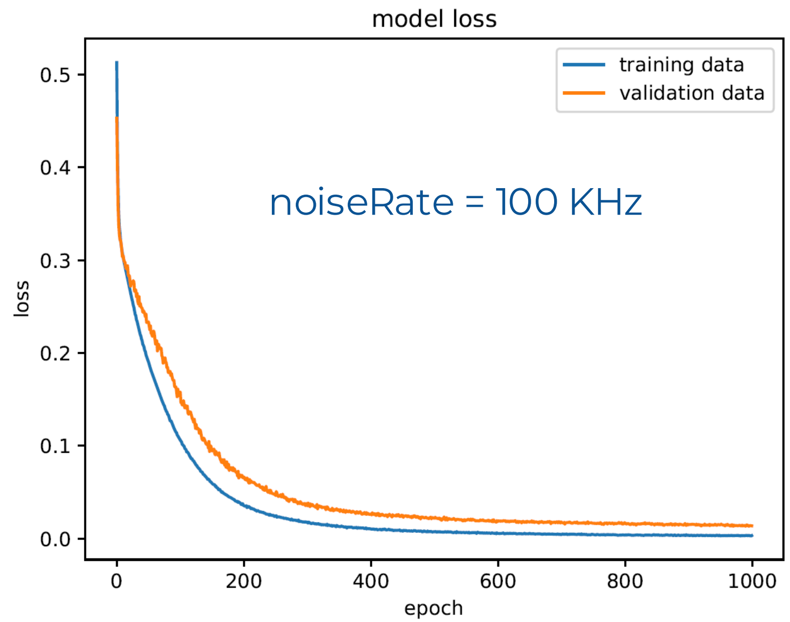

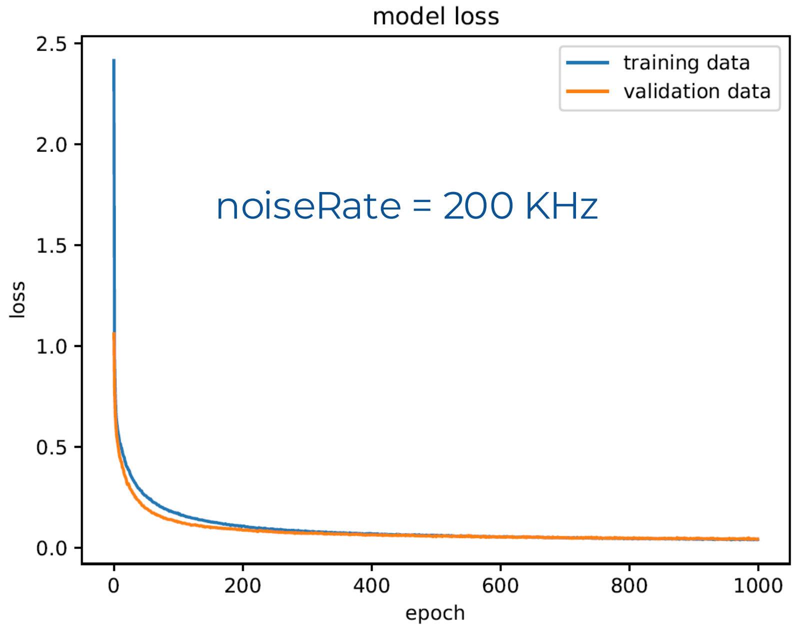

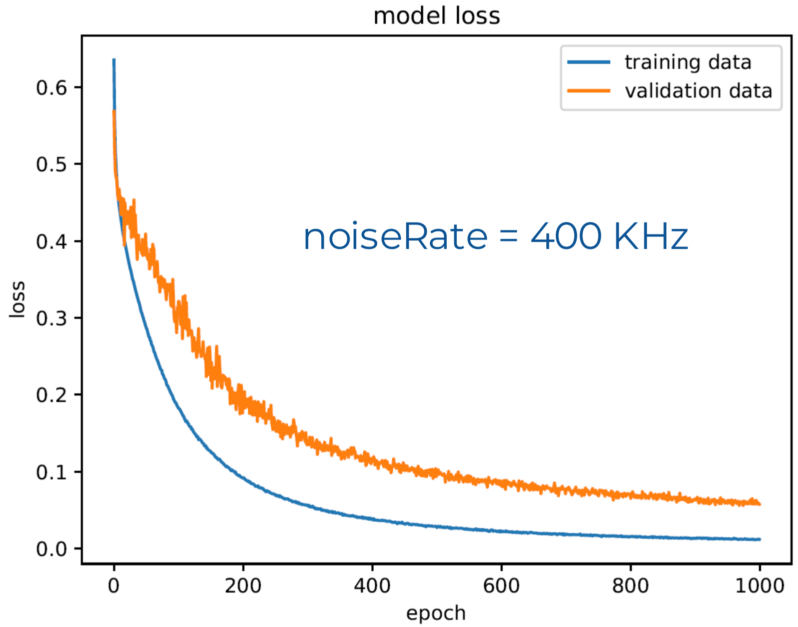

Model training & validation: Accuracy

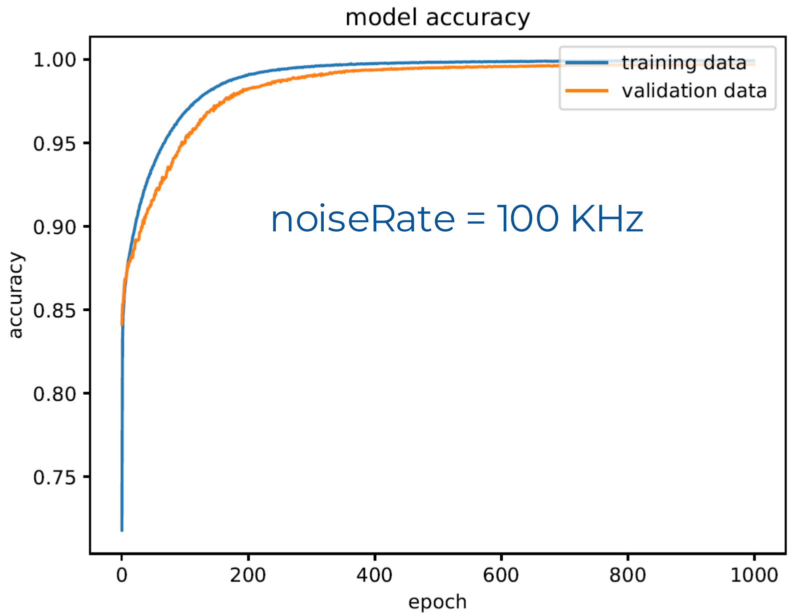

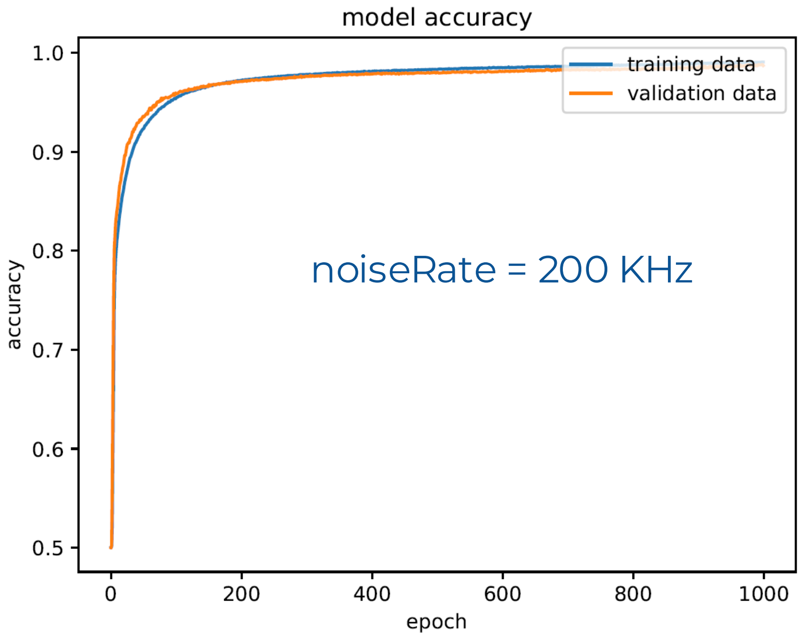

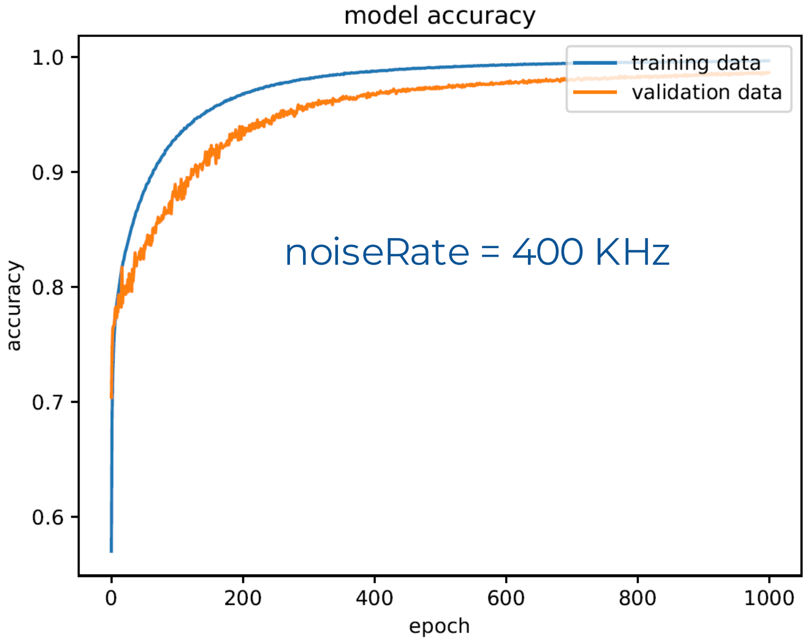

Model performance @ noiseRate = 100 KHz

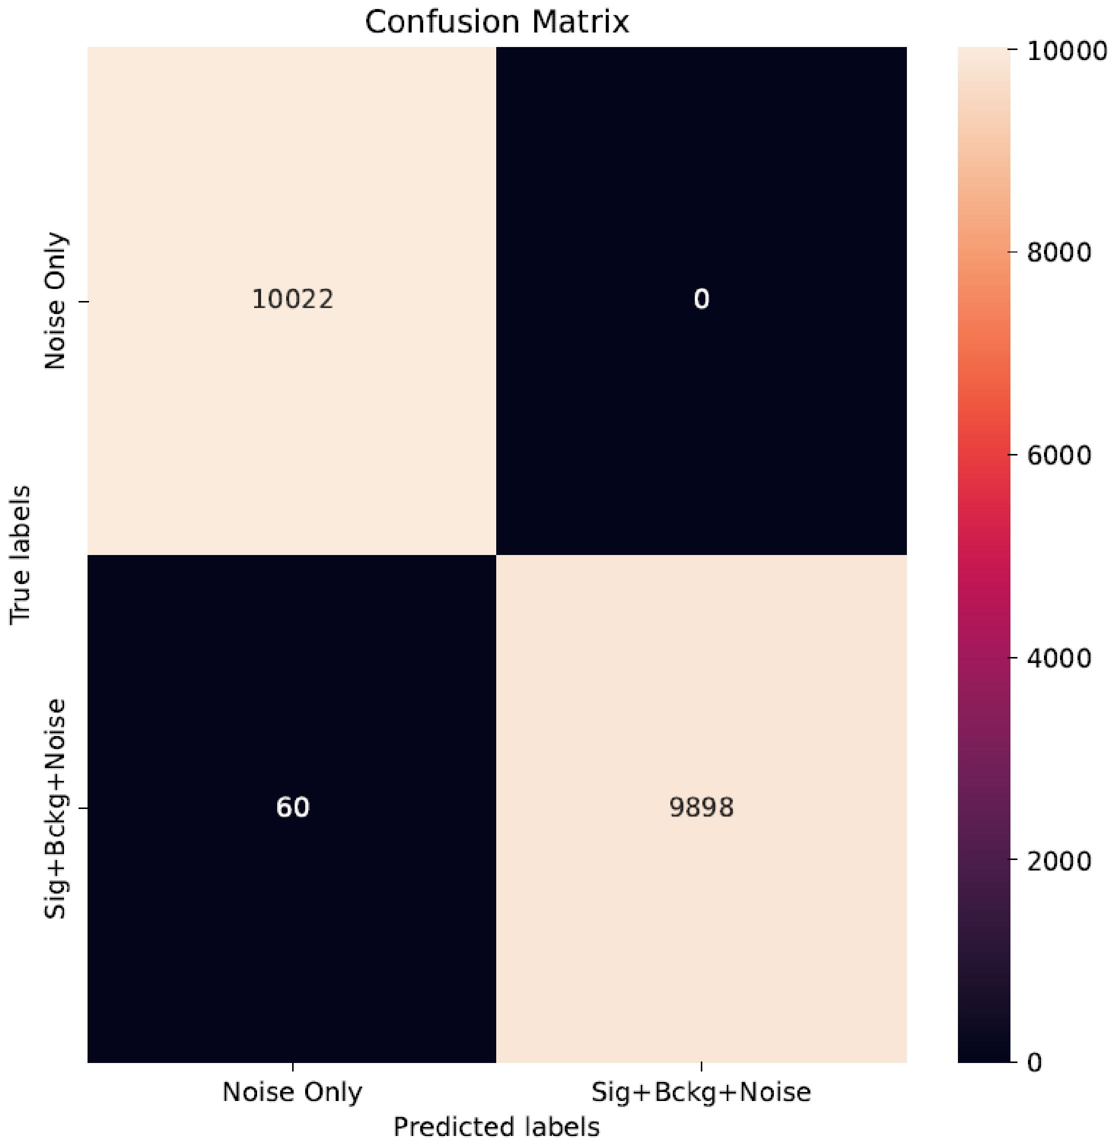

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.997
- Precision = TP/(TP+FP) = 0.994
- Recall = TP/(TP+FN) = 1.000

Model performance @ noiseRate = 200 KHz

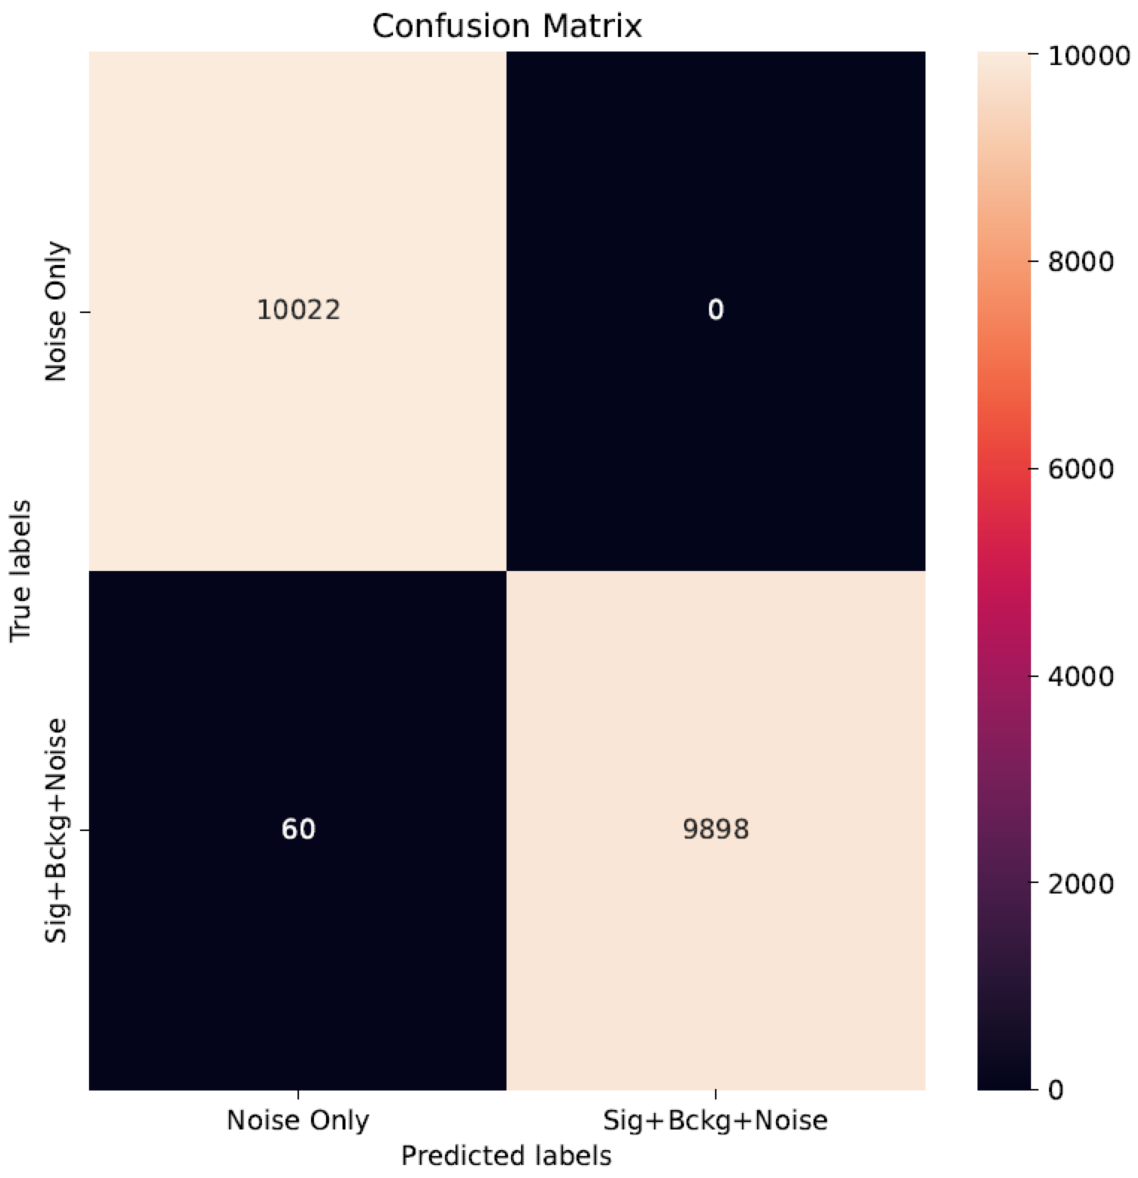

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.997
- Precision = TP/(TP+FP) = 0.994
- Recall = TP/(TP+FN) = 1.000

Model performance @ noiseRate = 400 KHz

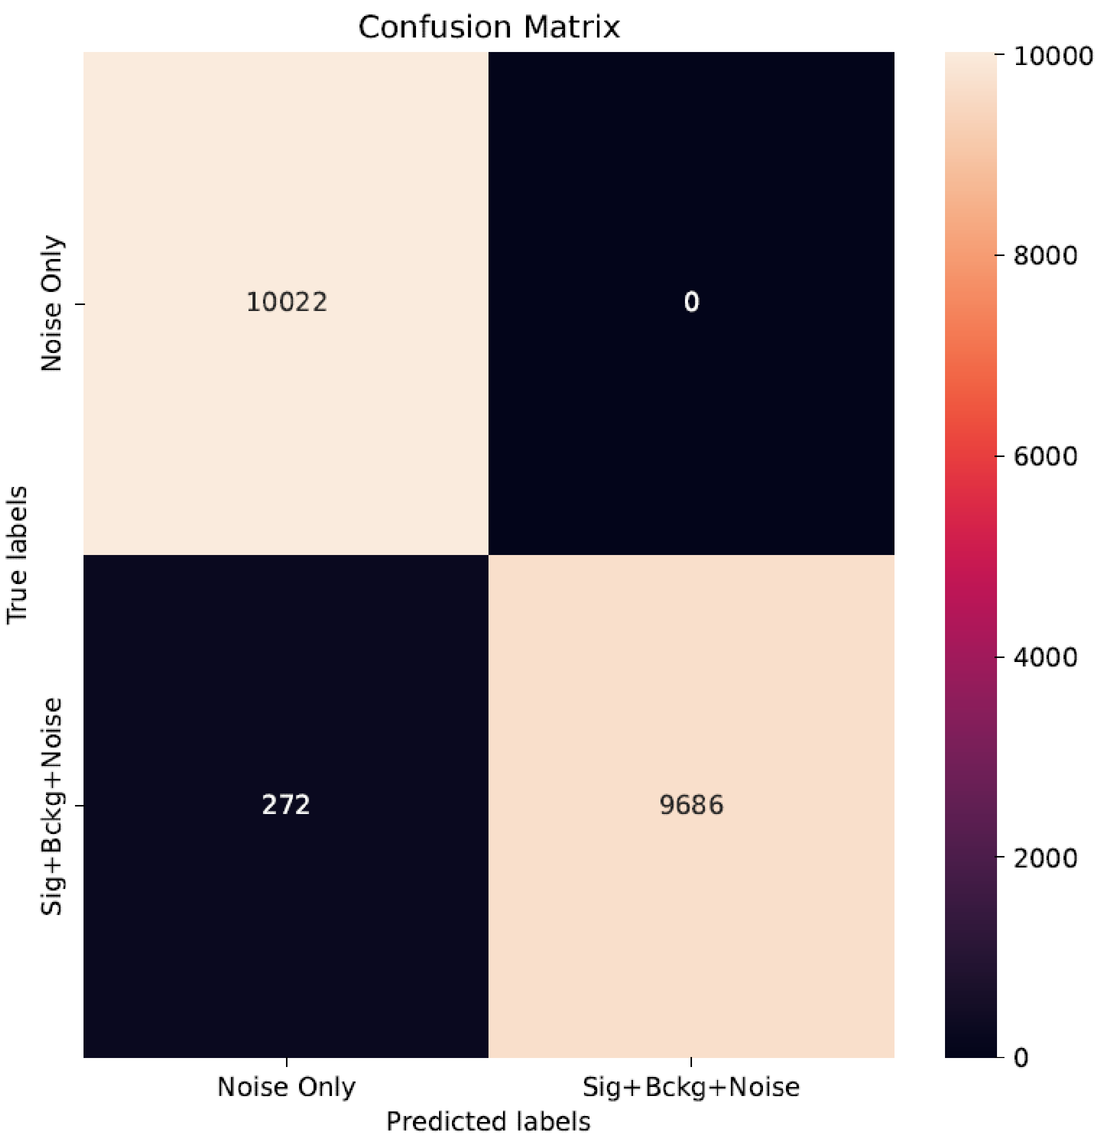

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.986
- Precision = TP/(TP+FP) = 0.974
- Recall = TP/(TP+FN) = 1.000

Quant. model performance @ noiseRate = 100 KHz

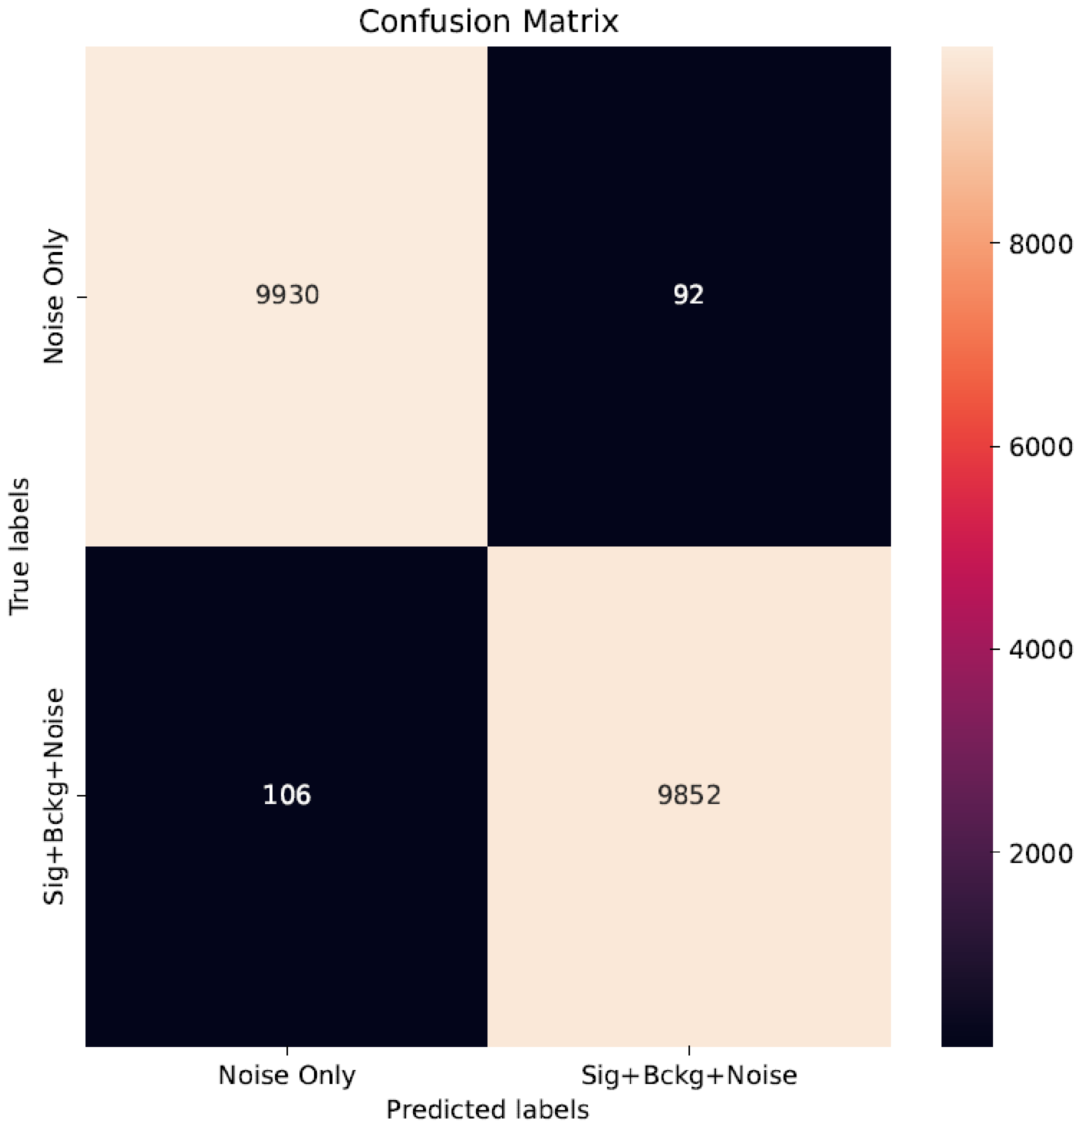

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.990
- Precision = TP/(TP+FP) = 0.989
- Recall = TP/(TP+FN) = 0.991

Model Quantization

- Inputs, Activations: fixed point<16,6>
- Weights, Biases: fixed point<8,1>

Quant. Model performance @ noiseRate = 200 KHz

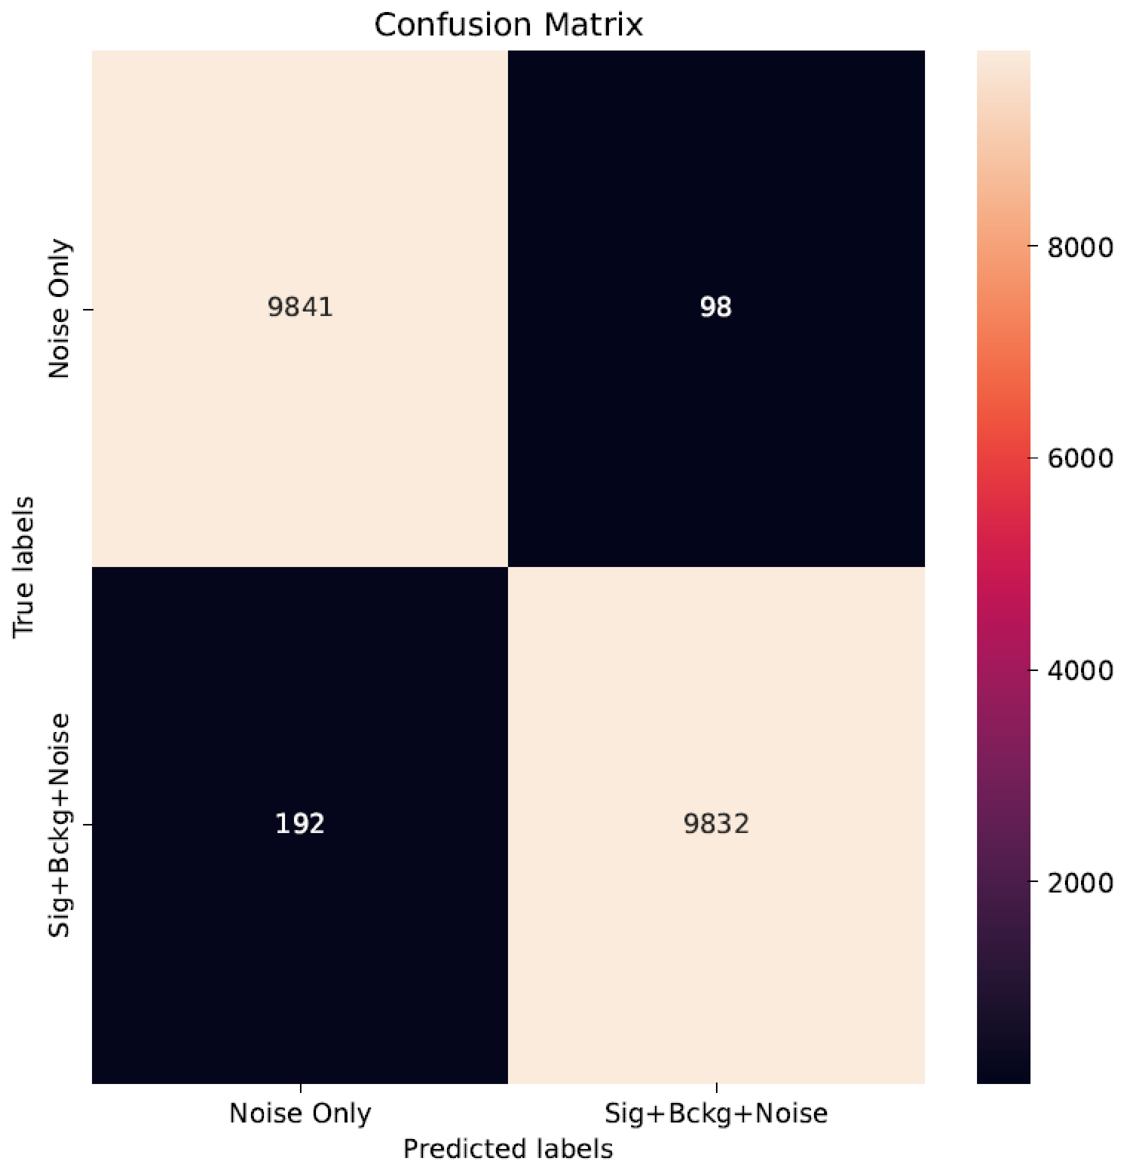

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.985
- Precision = TP/(TP+FP) = 0.981
- Recall = TP/(TP+FN) = 0.990

Model Quantization

- Inputs, Activations: fixed point<16,6>
- Weights, Biases: fixed point<8,1>

Quant. model performance @ noiseRate = 400 KHz

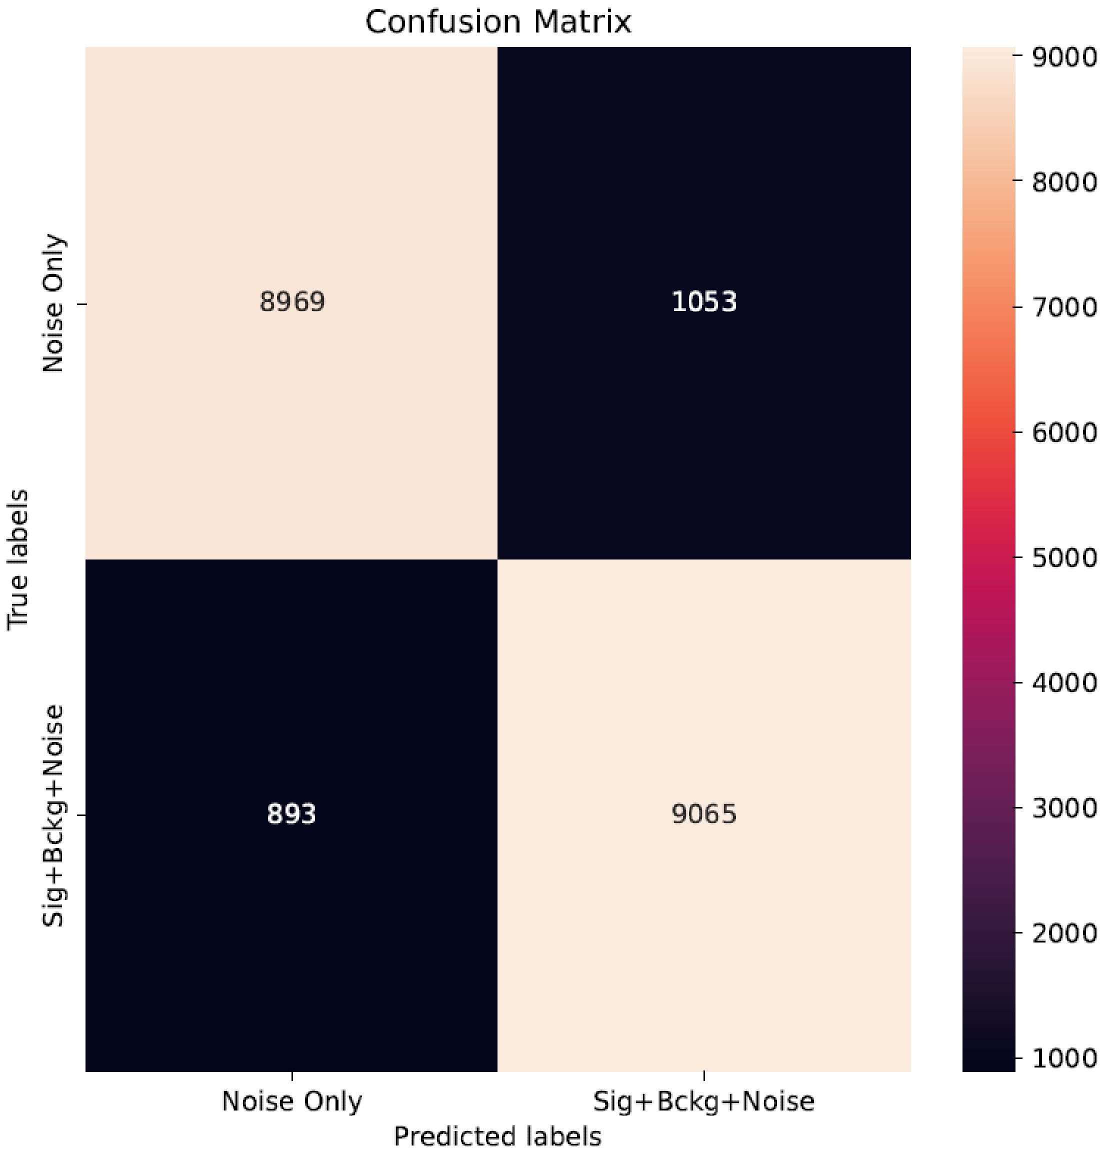

- Accuracy = (TP+TN)/(TP+TN+FP+FN) = 0.903
- Precision = TP/(TP+FP) = 0.909
- Recall = TP/(TP+FN) = 0.895

Model Quantization

- Inputs, Activations: fixed point<16,6>
- Weights, Biases: fixed point<8,1>

Summary of Distributed MLP Performance

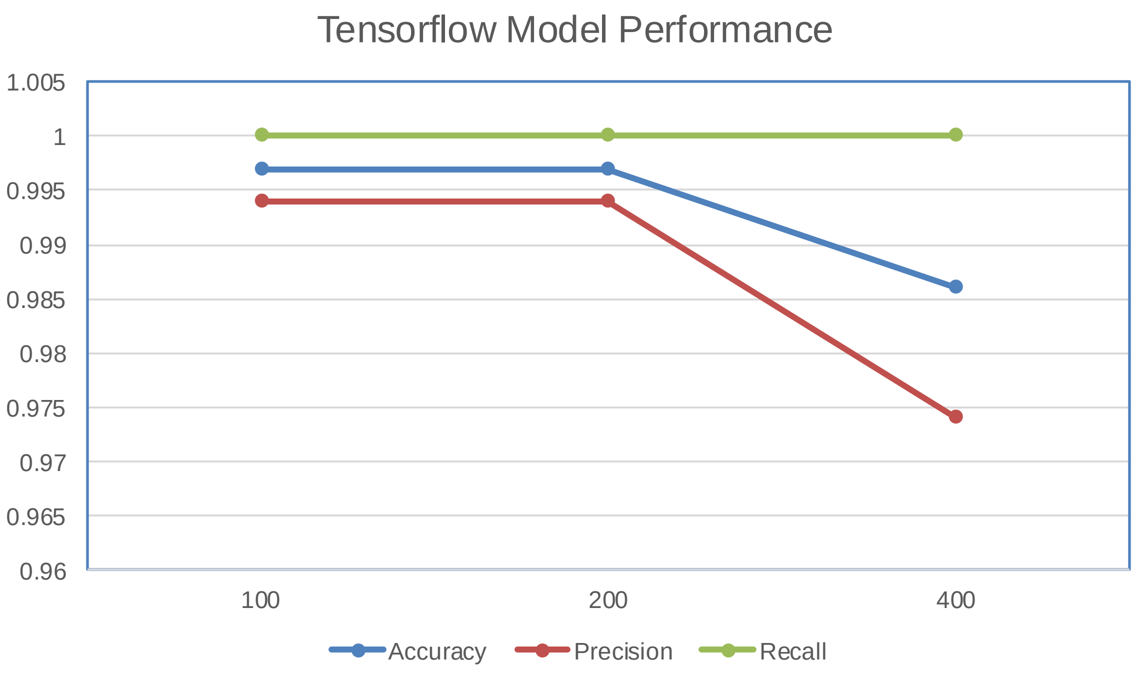

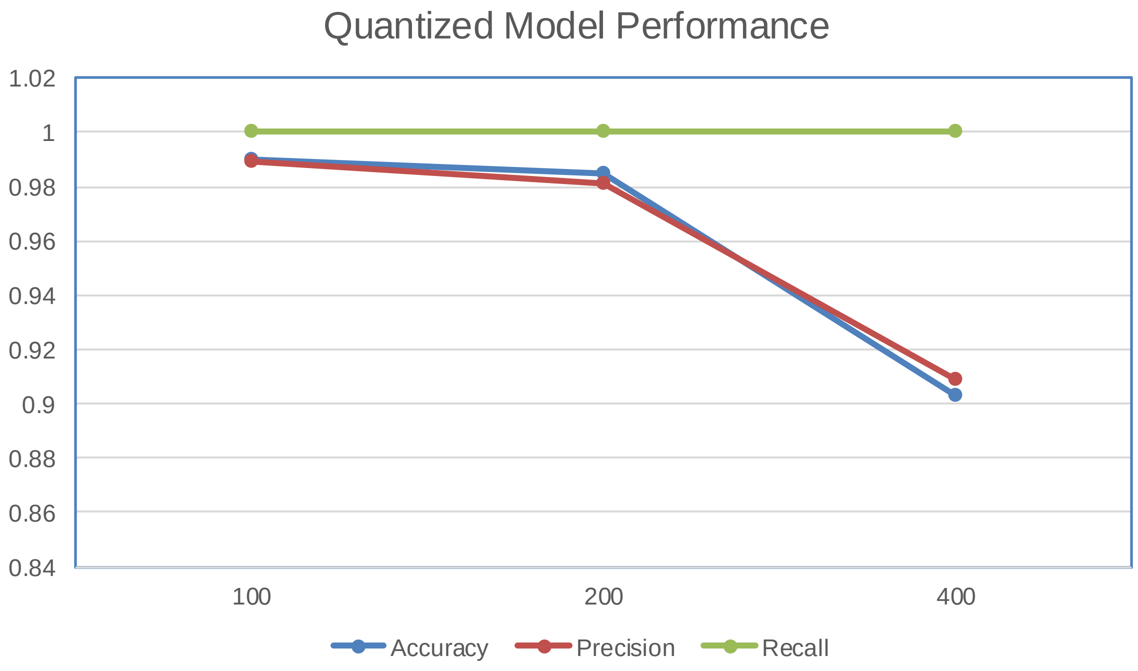

Close-up → False Positive Events

Training and validating with datasets of 100 kHz dark count rates, we obtain a **99% accurate model**.

BUT WHAT ABOUT THESE FALSE POSITIVE EVENTS?

WHAT THEY LOOK LIKE?

ARE THEY **TRULY** SHOWING SIGNAL+BACKGROUND FEATURES?

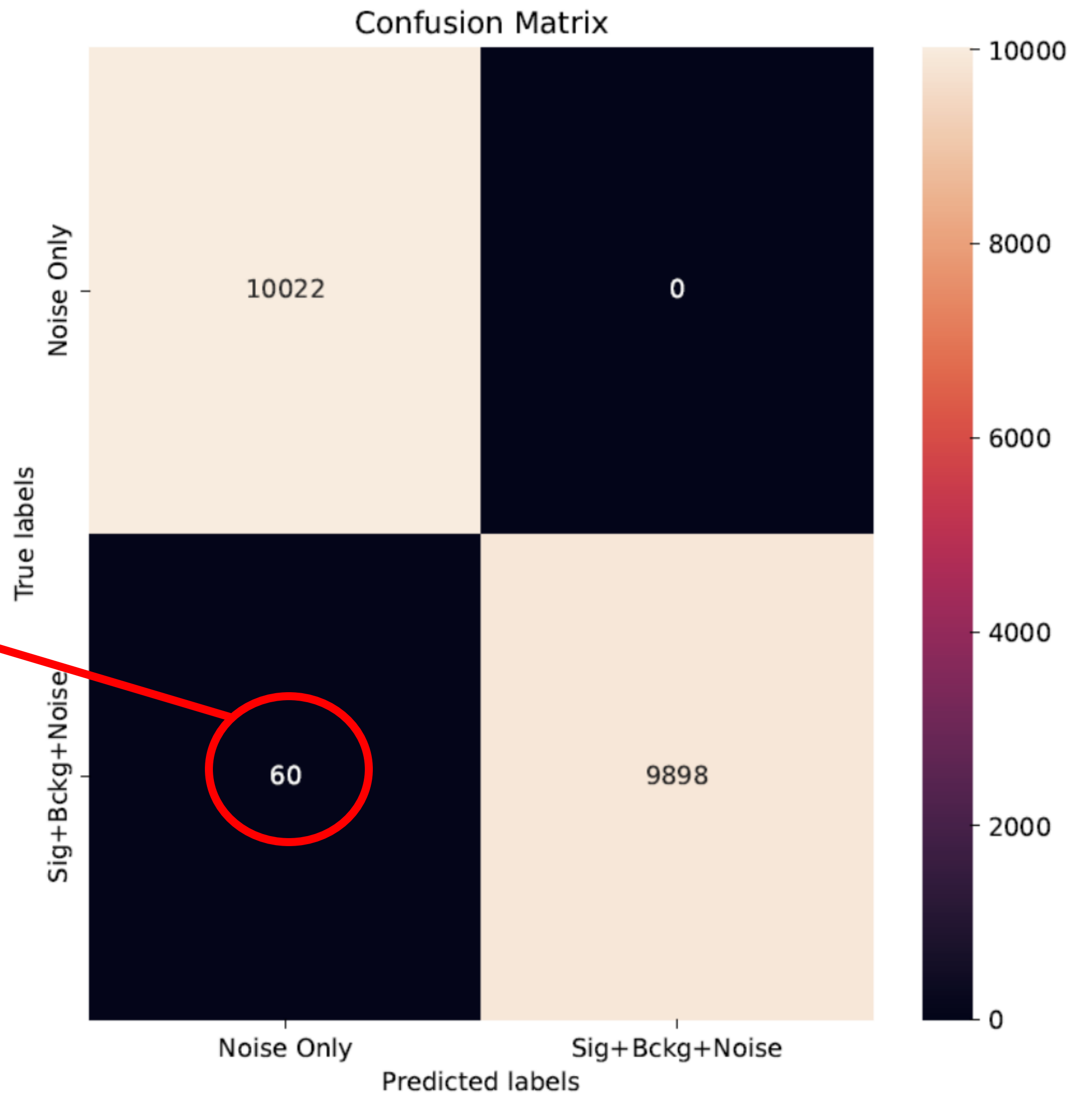

Close-up ➔ False Positive Events

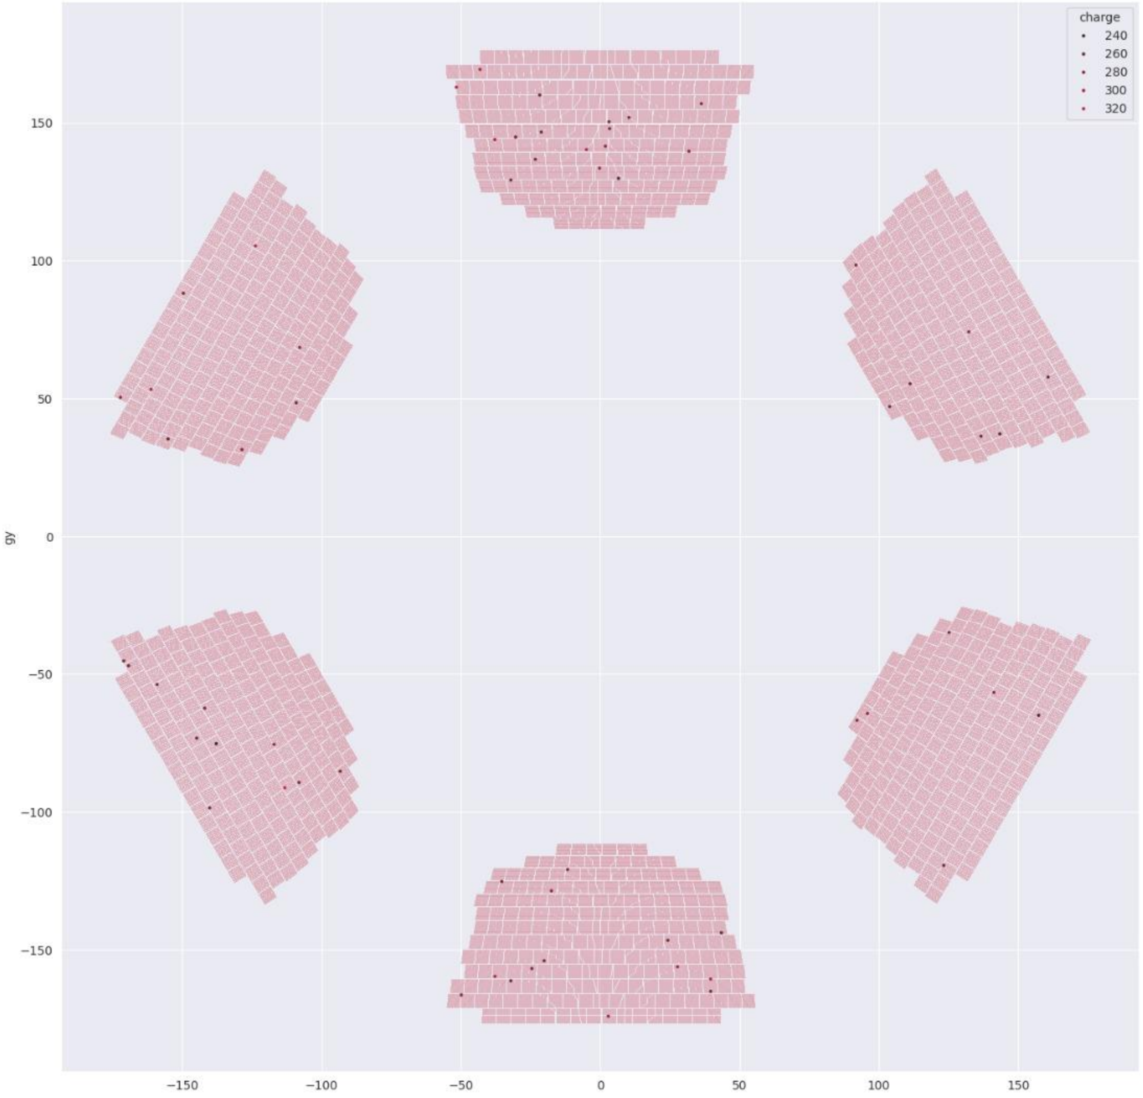

Example of a **False Positive event** (signal+background+noise, but **classified as noise**):

- Low number of dRICH hits
- **No Cherenkov rings detected**
- No evident dRICH hits clusters
- Homogenous dRICH hits distribution **➔ comparable with a noise hits distribution**

Close-up ➔ False Positive ROOT TTree

MCParticles.PDG = 22, 11, 2212, 9900330, 2212, -311, 313, 2212, 11, 130, 311, 111, 310, 22, 22, 111, 111, 22, 22, 22
MCParticles.generatorStatus = 21, 21, 21, 21, 21, 2, 2, 1, 1, 1, 2, 2, 2, 1, 1, 2, 2, 1, 1, 1
[…]
MCParticles.time = 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.184769, 212.184769, 212.184769, 212.184769, 212.184769
[…]
MCParticles.momentum.x = 0.000092, -0.000105, -2.521645, 0.352699, -2.874251, 0.030792, 0.321907, -2.874251, -0.000105, 0.030793, 0.075927, 0.245985, 0.075927, 0.121866, 0.124117, 0.146451, -0.070528, 0.060760, 0.085690, -0.060012
MCParticles.momentum.y = -0.000563, 0.000807, -0.012031, 0.239004, -0.251596, -0.168178, 0.407180, -0.251596, 0.000807, -0.168186, 0.058748, 0.348438, 0.058748, 0.155661, 0.192775, 0.125229, -0.066484, 0.037693, 0.087535, 0.007046
MCParticles.momentum.z = -1.228703, -8.770502, 99.992050, -0.499420, 99.262772, 0.339107, -0.838527, 99.262772, -8.770502, 0.339123, -0.206120, -0.632420, -0.206120, -0.410238, -0.222180, -0.279635, 0.073525, -0.013916, -0.265717, 0.087298

Close-up ➔ False Positive ROOT TTree

γ

MCParticles.PDG = **22**, 11, 2212, 9900330, 2212, -311, 313, 2212, 11, 130, 311, 111, 310, 22, 22, 111, 111, 22, 22, 22

MCParticles.generatorStatus = **21**, 21, 21, 21, 21, 2, 2, 1, 1, 1, 2, 2, 2, 1, 1, 2, 2, 1, 1, 1

[…]

MCParticles.time = **212.173264**, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.184769, 212.184769, 212.184769, 212.184769, 212.184769

[…]

MCParticles.momentum.x = **0.000092**, -0.000105, -2.521645, 0.352699, -2.874251, 0.030792, 0.321907, -2.874251, -0.000105, 0.030793, 0.075927, 0.245985, 0.075927, 0.121866, 0.124117, 0.146451, -0.070528, 0.060760, 0.085690, -0.060012

MCParticles.momentum.y = **-0.000563**, 0.000807, -0.012031, 0.239004, -0.251596, -0.168178, 0.407180, -0.251596, 0.000807, -0.168186, 0.058748, 0.348438, 0.058748, 0.155661, 0.192775, 0.125229, -0.066484, 0.037693, 0.087535, 0.007046

MCParticles.momentum.z = **-1.228703**, -8.770502, 99.992050, -0.499420, 99.262772, 0.339107, -0.838527, 99.262772, -8.770502, 0.339123, -0.206120, -0.632420, -0.206120, -0.410238, -0.222180, -0.279635, 0.073525, -0.013916, -0.265717, 0.087298

Close-up ➔ False Positive ROOT TTree

e^-

MCParticles.PDG = 22, **11**, 2212, 9900330, 2212, -311, 313, 2212, 11, 130, 311, 111, 310, 22, 22, 111, 111, 22, 22, 22
MCParticles.generatorStatus = 21, **21**, 21, 21, 21, 2, 2, 1, 1, 1, 2, 2, 2, 1, 1, 2, 2, 1, 1, 1
[…]
MCParticles.time = 212.173264, **212.173264**, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.184769, 212.184769, 212.184769, 212.184769, 212.184769
[…]
MCParticles.momentum.x = 0.000092, **-0.000105**, -2.521645, 0.352699, -2.874251, 0.030792, 0.321907, -2.874251, -0.000105, 0.030793, 0.075927, 0.245985, 0.075927, 0.121866, 0.124117, 0.146451, -0.070528, 0.060760, 0.085690, -0.060012
MCParticles.momentum.y = -0.000563, **0.000807**, -0.012031, 0.239004, -0.251596, -0.168178, 0.407180, -0.251596, 0.000807, -0.168186, 0.058748, 0.348438, 0.058748, 0.155661, 0.192775, 0.125229, -0.066484, 0.037693, 0.087535, 0.007046
MCParticles.momentum.z = -1.228703, **-8.770502**, 99.992050, -0.499420, 99.262772, 0.339107, -0.838527, 99.262772, -8.770502, 0.339123, -0.206120, -0.632420, -0.206120, -0.410238, -0.222180, -0.279635, 0.073525, -0.013916, -0.265717, 0.087298

Close-up ➔ False Positive ROOT TTree

p

```
MCParticles.PDG = 22, 11, 2212, 9900330, 2212, -311, 313, 2212, 11, 130, 311,
111, 310, 22, 22, 111, 111, 22, 22, 22
MCParticles.generatorStatus = 21, 21, 21, 21, 21, 2, 2, 1, 1, 1, 2, 2, 2, 1, 1,
2, 2, 1, 1, 1
[…]
MCParticles.time = 212.173264, 212.173264, 212.173264, 212.173264, 212.173264,
212.173264, 212.173264, 212.173264, 212.173264, 212.173264, 212.173264,
212.173264, 212.173264, 212.173264, 212.173264, 212.184769, 212.184769,
212.184769, 212.184769, 212.184769
[…]
MCParticles.momentum.x = 0.000092, -0.000105, -2.521645, 0.352699, -2.874251,
0.030792, 0.321907, -2.874251, -0.000105, 0.030793, 0.075927, 0.245985,
0.075927, 0.121866, 0.124117, 0.146451, -0.070528, 0.060760, 0.085690, -
0.060012
MCParticles.momentum.y = -0.000563, 0.000807, -0.012031, 0.239004, -0.251596, -
0.168178, 0.407180, -0.251596, 0.000807, -0.168186, 0.058748, 0.348438,
0.058748, 0.155661, 0.192775, 0.125229, -0.066484, 0.037693, 0.087535, 0.007046
MCParticles.momentum.z = -1.228703, -8.770502, 99.992050, -0.499420, 99.262772,
0.339107, -0.838527, 99.262772, -8.770502, 0.339123, -0.206120, -0.632420, -
0.206120, -0.410238, -0.222180, -0.279635, 0.073525, -0.013916, -0.265717,
0.087298
```

Close-up ➔ False Positive ROOT TTree

Example of a **False Positive simulate event TTree ROOT File entries:**

- High Momentum Z-component charged particles (e-, p) ➔ pseudorapidity not in the dRICH acceptance (1.5 – 3.5)
- Neutral secondary products ➔ no Cherenkov rings
- Low momentum secondary products

dRICH Data Reduction:
HLS4ML ➔ HW Synthesis for 8x8 Grid DAM NN

➔ To correctly synthetize the model at 200 MHz of operational clock, we used a **REUSE FACTOR = 1**, obtaining an instantiation interval **II = 5 clock cycles**

➔ **Throughput = 40MHz (< 100 MHz)**

+ Timing:
  * Summary:

Clock	Target	Estimated	Uncertainty
ap_clk	5.00 ns	4.374 ns	0.62 ns

+ Latency:
  * Summary:

Latency (cycles) min	Latency (cycles) max	Latency (absolute) min	Latency (absolute) max	Interval min	Interval max	Pipeline Type
14	14	70.000 ns	70.000 ns	5	5	dataflow

dRICH Data Reduction:
HLS4ML ➔ HW Synthesis for 8x8 Grid DAM NN

➔ The possible **overhead** in the full II pipepline **introduced by the communication between DAMs and TP** will be considered **in further developments**

STILL LOW, BUT PROMISING! (can be improved via modifying part of HLS4ML code)

➔ To correctly synthetize the model at 200 MHz of operational clock, we used a **REUSE FACTOR = 1**, obtaining an instantiation interval **II = 5 clock cycles**

➔ **Throughput = 40MHz (< 100 MHz)**

\+ Timing:
* Summary:

Clock	Target	Estimated	Uncertainty
ap_clk	5.00 ns	4.374 ns	0.62 ns

\+ Latency:
* Summary:

Latency (cycles)		Latency (absolute)		Interval		Pipeline
min	max	min	max	min	max	Type
14	14	70.000 ns	70.000 ns	5	5	dataflow

Realistic Noise Model for EICRECON

```
/**
    this function returns the probability for a channel to fire due to dark noise
    as a function of the radial position, the selection time window and the integrated luminosity
**/

const float baseline_dcr = 3.e3; // [Hz] new sensors at T = -30 C and Vover = 4V
const float dcr_increase = 300.e3 / 1.e9; // [Hz/neq]
float neq_radius_params[6] = { -3.27029e+09, 1.26055e+08, -1.88568e+06, 13929.1, -50.9931, 0.0741068 };

float neq_radius(float radius /* cm */)
{
  float neq = 0.;
  for (int ipar = 0; ipar < 6; ++ipar)
    neq += neq_radius_params[ipar] * std::pow(radius, ipar);
  return neq;
}

float
noise_probability(float radius = 150. /* cm */, float window = 10. /* ns */, float luminosity = 100. /* fb-1 */)
{
  float neq = neq_radius(radius) * luminosity;
  float dcr = baseline_dcr + dcr_increase * neq;
  float pro = dcr * window * 1.e-9;
  return pro;
}
```

Conclusions

- We sketched a data reduction system designed based on DAM's FPGAs as a risk-mitigation action to the possible problem of an excessive data bandwidth requirement from the dRICH to Echelon-0 due to SiPMs DCR.
- We showed results of the initial activities we made to proof the design concept.
- The design is based on a **distributed Dense MLP NN** model, that can reach **near-optimal performance (using simulated data), and promising performance in terms of throughput of the first part of the pipeline (need to improve by a x2.5 factor).** **<u>These results need to be confirmed with a more realistic noise model (started…)</u>**
- Next steps:
  - Deploy the distributed NN on two FPGAs already available in our lab (Xilinx Alveo U200) representing a DAM and the TP, integrating the communication in the pipeline and assessing its impact on pipeline throughput (and latency).
  - In addition different NN models (CNNs, GNN,…) and data reduction tasks/ideas (Cherenkov ring detection…) can be explored
  - Become familiar with the FELIX board HW and FW (we received a FLX-182 on loan from JLab) to start devising the integration of our design in its FW.
  - A initial «parasitic mode» deployment would allow the tuning and assessment of performace of the system, with periodic re-training of the NN with real data.

Backup Slides

APEIRON: the Node

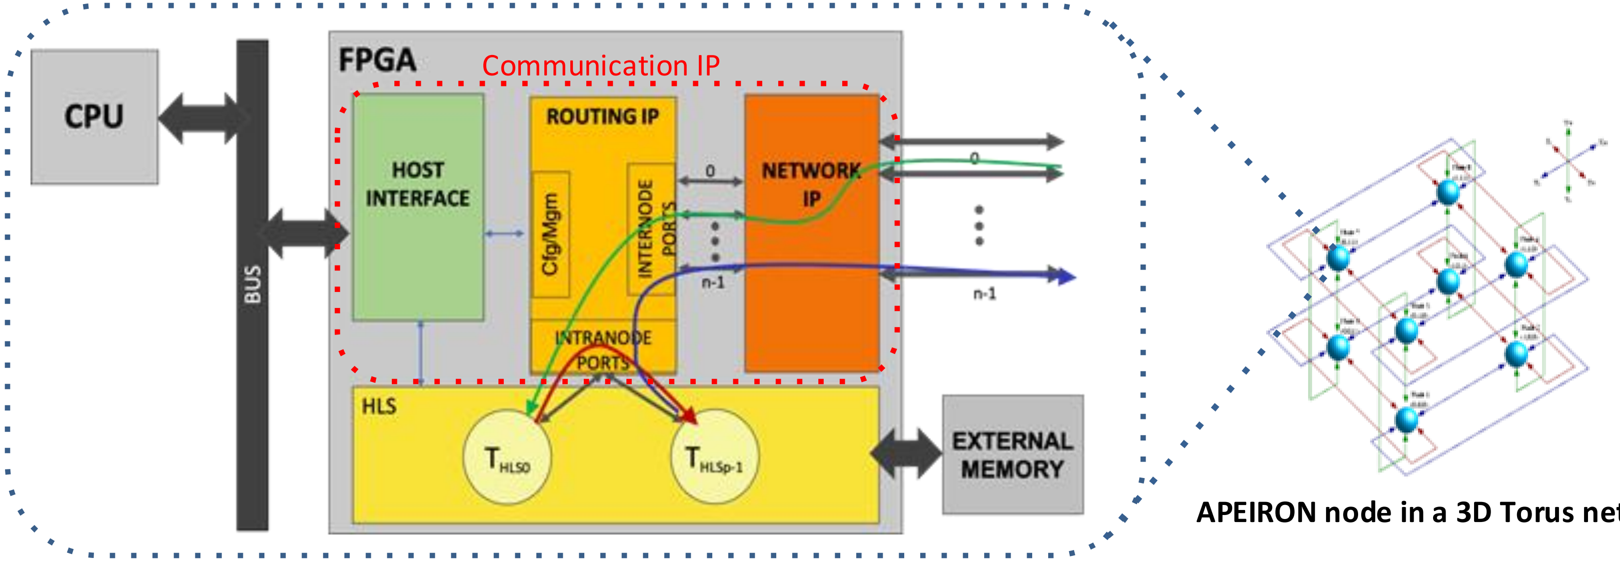

APEIRON node in a 3D Torus network topology

- **Host Interface IP**: Interface the FPGA logic with the host through the system bus.
  - Xilinx XDMA PCIe Gen3
- **Routing IP**: Routing of intra-node and inter-node messages between processing tasks on FPGA.
- **Network IP**: Network channels and Application-dependent I/O
  - APElink 40 Gbps
  - UDP/IP over 10 GbE
- **Processing Tasks**: user defined processing tasks (Xilinx Vitis HLS Kernels)

APEIRON: Communication Latency

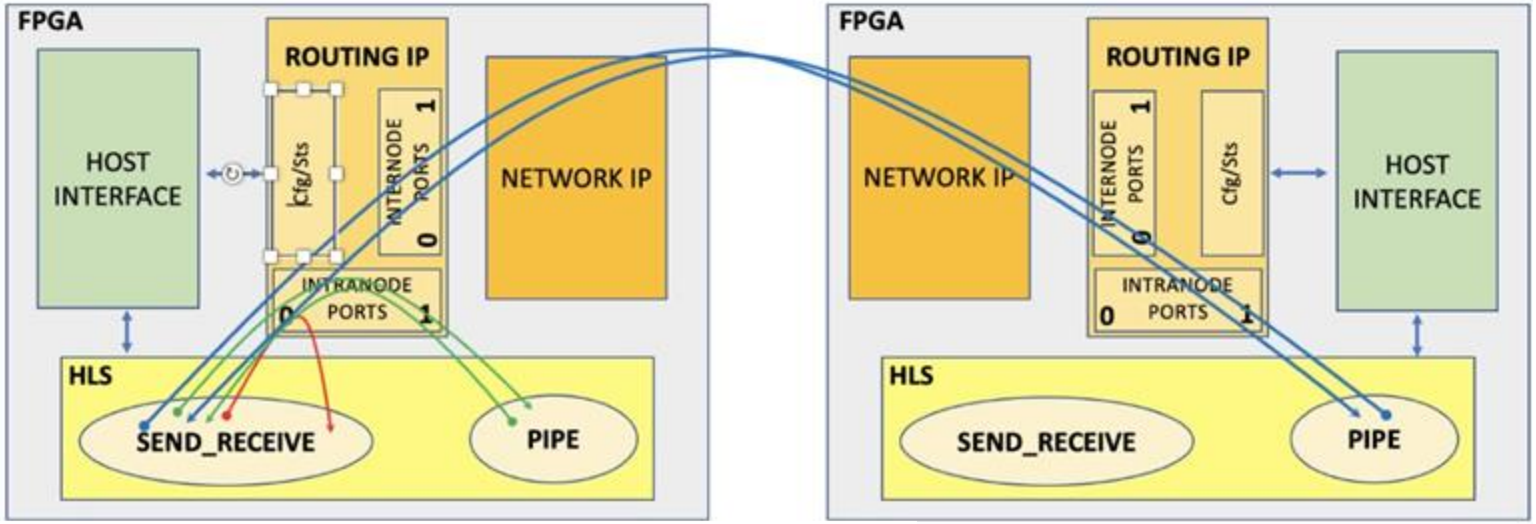

Test modes
- Local-loop (red arrow)
- Local-trip (green arrows)
- Round-trip (blue arrows)

Test Configuration
- IP logic clock @ 200 MHz
- 4 intranode ports
- 2 internode ports
- 256-bit datapath width
- 4 lanes inter-node channels

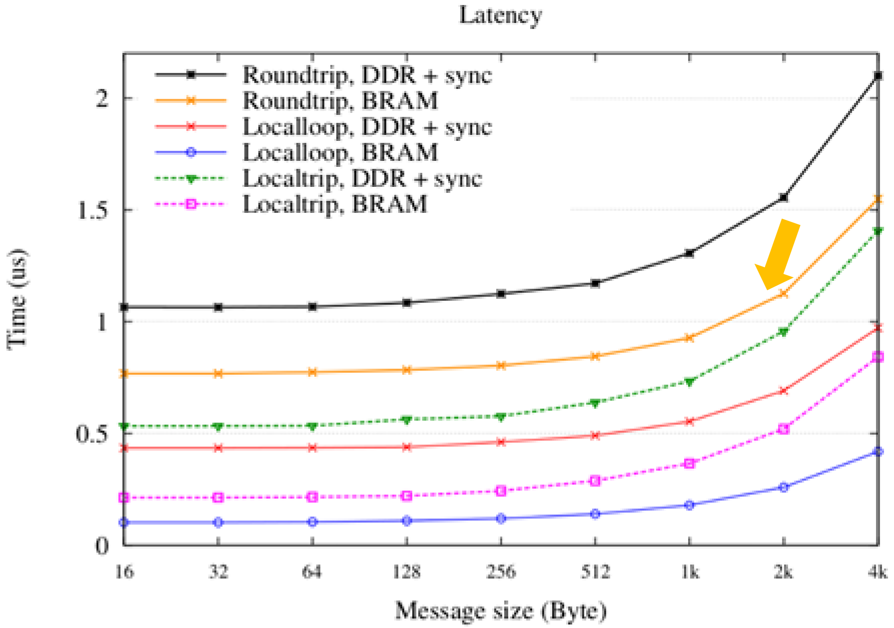

Inter-node LATENCY (orange line) < 1us for packet sizes up to 1kB (source and destination buffers in BRAM)

FELIX Hardware Development at BNL

FLX-711
@2017

PCIe Gen3 x16
48x 10 Gb/s
DDR3

FLX-712
@2018

PCIe Gen3 x16
48x 10 Gb/s

FLX Demo
@2019

PCIe Gen3 x16
25 Gb/s with various optical modules

FLX-181
@2021

PCIe Gen4 x16
FMC+ 16x 25 Gb/s
12x 16 Gb/s
SoC / DDR4

FLX-182
@2022

FLX-182B
@2023

PCIe Gen4 x16
28x 10/25 Gb/s
100 GbE
SoC / DDR4

FLX-155
@2024

PCIe Gen5 x16
48x 10/25 Gb/s
100/400 GbE
SoC / DDR4

FLX-182B Hardware

Assembled FLX-182B

- FPGA: Xilinx Versal Prime XCVM1802
- PCIe Gen4 x16, 256 GT/s
- 24 FireFly links with 3 possible configurations
  - 24 links up to 25 Gb/s
  - 24 links up to 10 Gb/s (CERN-B FireFly)
  - 12 links up to 25 Gb/s + 12 links up to 10 Gb/s
- 4 FireFly links with 2 possible configurations with 14 or 25 Gb/s FireFly TRx
  - LTI interface
  - 100 GbE
- Built-in self test, online configuration and monitoring
- White Rabbit
- DDR4 Mini-UDIMM
- GbE/SD3.0/PetaLinux

FLX-155 Hardware

Assembled FLX-155

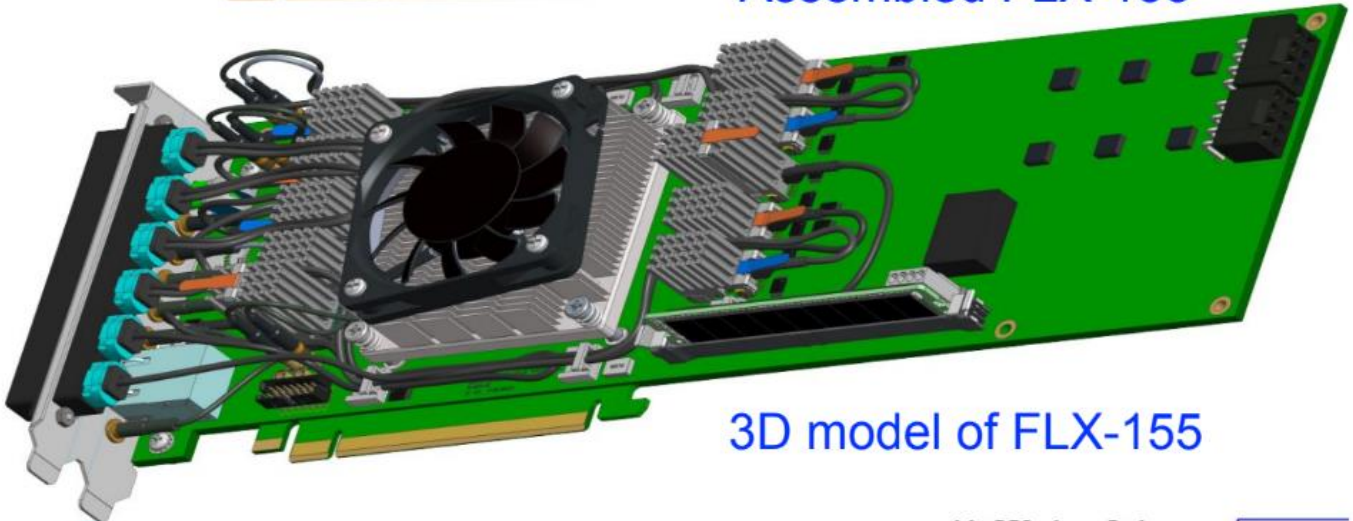

3D model of FLX-155

- AMD/Xilinx Versal Premium FPGA: XCVP1552-2MSEVSVA3340
- 2 x PCIe Gen5 x8 512 GT/s
- 56 FireFly optical links
  - Compatible with various options
  - Default configuration for ATLAS
    - 48 data links up to 25 Gb/s
    - 4 links for LTI
  - Optional 4 links for 100 GbE
- Electrical IOs
- Built-in self test, online configuration and monitoring
- 1 16GB DDR4 Mini-UDIMM
- USB-JTAG/USB-UART
- GbE/SD3.0/PetaLinux
- Optional White Rabbit

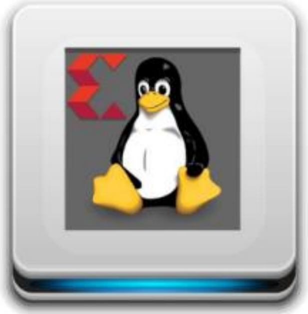

	VP1002	VP1052	VP1102	VP1202	VP1402	VP1502	VP2502	VP1552	VP1702	VP1802	VP2802	VP1902
System Logic Cells	833,000	1,185,800	1,574,720	1,969,240	2,233,280	3,763,480	3,737,720	3,836,840	5,557,720	7,351,960	7,326,200	18,506,880
CLB Flip-Flops	761,600	1,084,160	1,439,744	1,800,448	2,041,856	3,440,896	3,417,344	3,507,968	5,081,344	6,721,792	6,698,240	16,920,576
LUTs	380,800	542,080	719,872	900,224	1,020,928	1,720,448	1,708,672	1,753,984	2,540,672	3,360,896	3,349,120	8,460,288
Distributed RAM (Mb)	12	17	22	27	31	53	52	54	78	103	102	258
Block RAM Blocks	535	751	1,405	1,341	1,981	2,541	2,541	2,541	3,741	4,941	4,941	6,808
Block RAM (Mb)	19	26	49	47	70	89	89	89	132	174	174	239
UltraRAM Blocks	345	489	453	677	645	1,301	1,301	1,301	1,925	2,549	2,549	2,200
UltraRAM (Mb)	97	138	127	190	181	366	366	366	541	717	717	619
Multiport RAM (Mb)	80	80	-	-	-	-	-	-	-	-	-	-
DSP Engines	1,140	1,572	1,904	3,984	2,672	7,440	7,392	7,392	10,896	14,352	14,304	6,864
AI Engines (AIE)	-	-	-	-	-	-	472	-	-	-	472	-
AIE Data Memory (Mb)	-	-	-	-	-	-	118	-	-	-	118	-
APU	Dual-core Arm Cortex-A72; 48 KB/32 KB L1 Cache w/ parity & ECC; 1 MB L2 Cache w/ ECC											
RPU	Dual-core Arm Cortex-R5F; 32 KB/32 KB L1 Cache; TCM w/ECC											
Memory	256 KB On-Chip Memory w/ECC											
Connectivity	Ethernet (x2); UART (x2); CAN-FD (x2); USB 2.0 (x1); SPI (x2); I2C (x2)											
NoC to PL Master / Slave Ports	22	22	30	28	42	52	52	52	76	100	100	192
DDR Bus Width	128	128	192	256	192	256	256	256	256	256	256	896
DDR Memory Controllers (DDRMC)	2	2	3	4	3	4	4	4	4	4	4	14
PCIe w/DMA (CPM4)	2 x Gen4x4	2 x Gen4x4	-	-	-	-	-	-	-	-	-	-
PCIe w/DMA (CPM5)	-	-	-	2 x Gen5x8	-	2 x Gen5x8	2 x Gen5x8	2 x Gen5x8	2 x Gen5x8	2 x Gen5x8	2 x Gen5x8	-
PCIe (PL PCIE4)	1 x Gen4x8	1 x Gen4x8	-	-	-	-	-	-	-	-	-	-
PCIe (PL PCIE5)	-	-	2 x Gen5x4	2 x Gen5x4	2 x Gen5x4	2 x Gen5x4	2 x Gen5x4	8 x Gen5x4	2 x Gen5x4	2 x Gen5x4	2 x Gen5x4	16 x Gen5x4
100G Multirate Ethernet MAC	3	5	6	2	6	4	4	4	6	8	8	12
600G Ethernet MAC	2	3	7	1	11	3	3	1	5	7	7	4
600G Interlaken	1	2	-	-	-	1	1	-	2	3	3	0
High-Speed Crypto Engines	1	1	3	1	4	2	2	2	3	4	4	0
GTY Transceivers[1]	8	8	-	-	-	-	-	-	-	-	-	-
GTYP Transceivers[1]	-	-	8	28[3]	8	28[3]	28[3]	68[3]	28[3]	28[3]	28[3]	128
GTM Transceivers[1] 58Gb/s (112 Gb/s)	24 (12)	36 (18)	64 (32)	20 (10)	96 (64)[2]	60 (30)	60 (30)	20 (10)	100 (50)	140 (70)	140 (70)	32 (16)