Readout & Learning

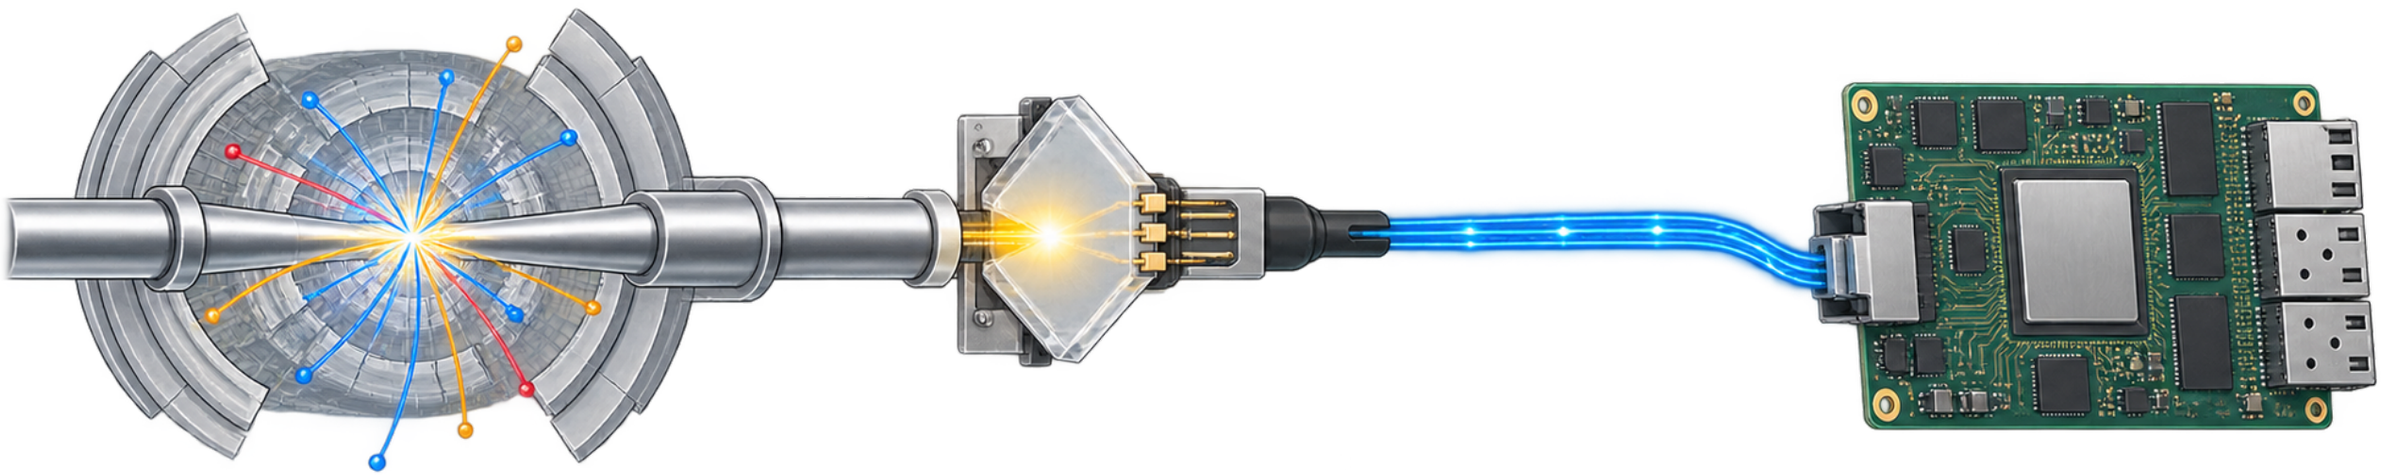
F9 Strategic meeting

September 30th, 2026

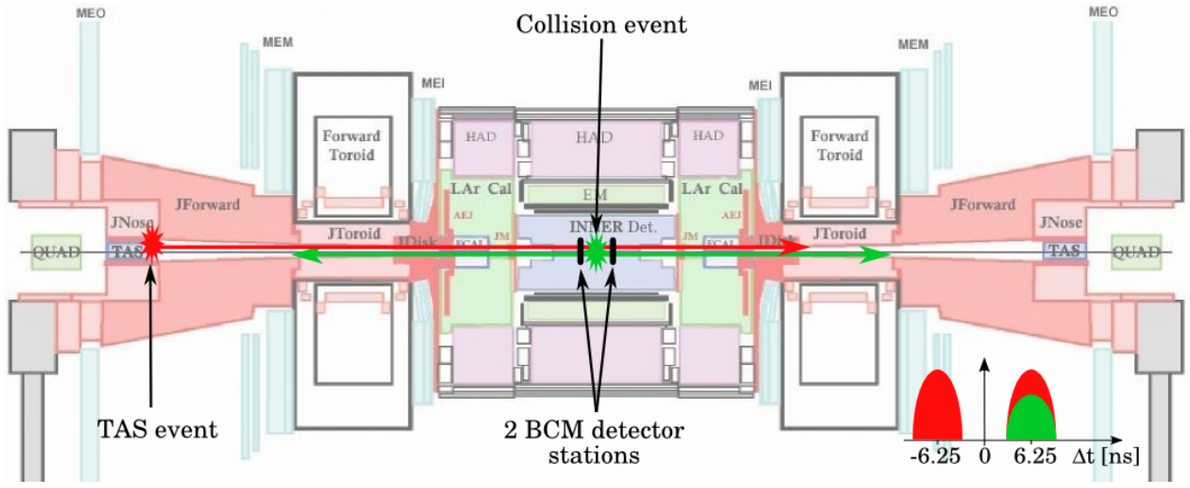
Boštjan Maček

Jožef Stefan Institute

BCM’ readout



BCM' system



protection - beam abort

- beam abort capability to protect the inner ATLAS
- crucial aspects: robustness, reliability
- must run independently of TDAQ infrastructure

BCM' TOF concept

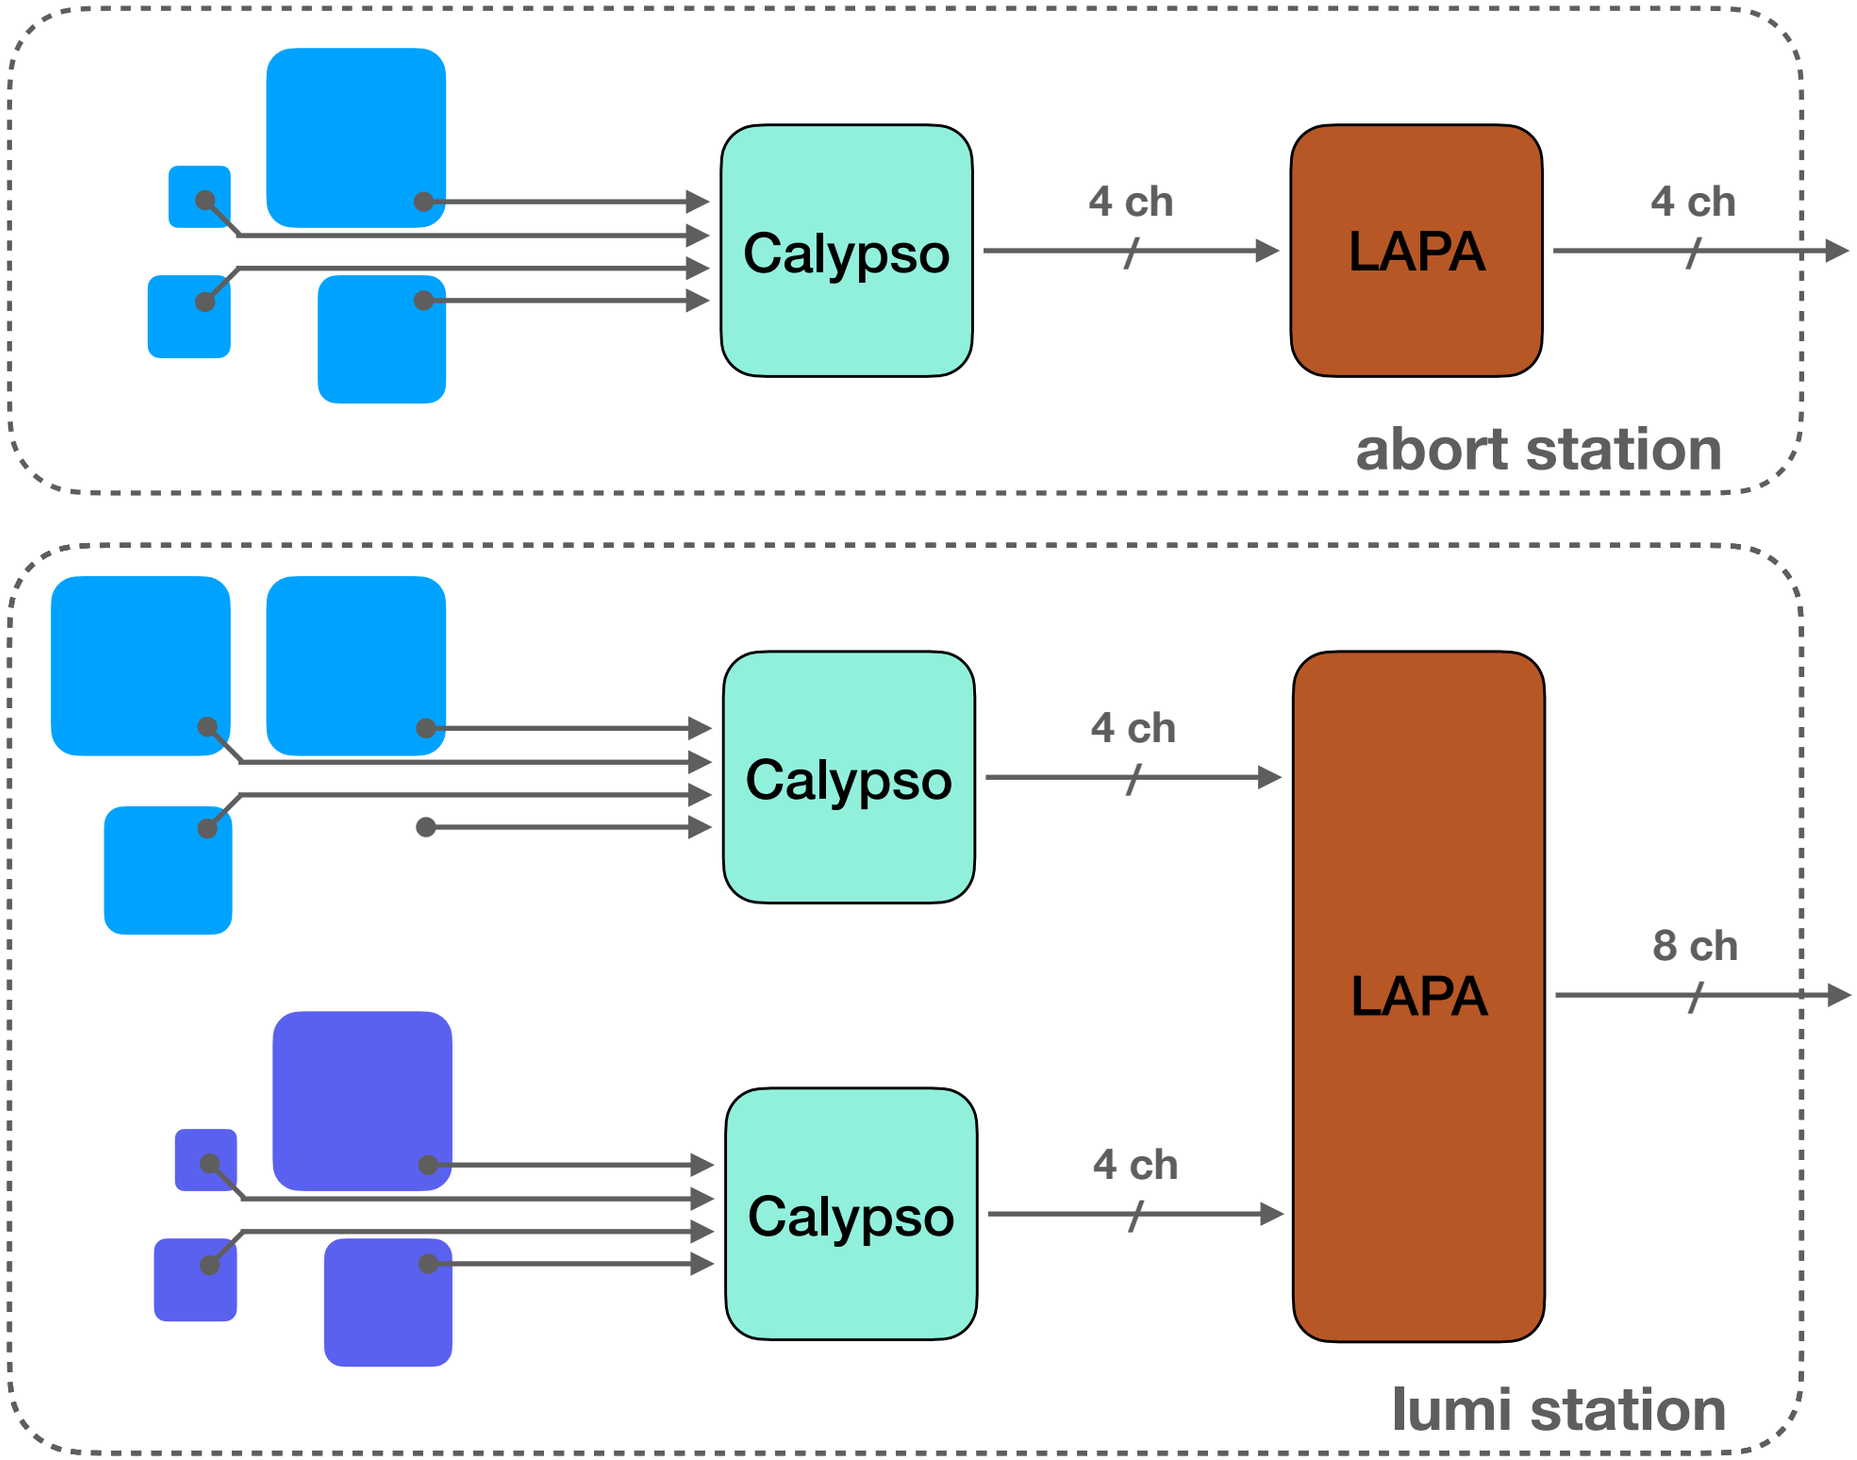
- **Collisions: in-time**
- **Background: out-of-time**

- safety sub-system has to serve ATLAS without TDAQ infrastructure
- multiple streams with different demands

luminosity & background measurement

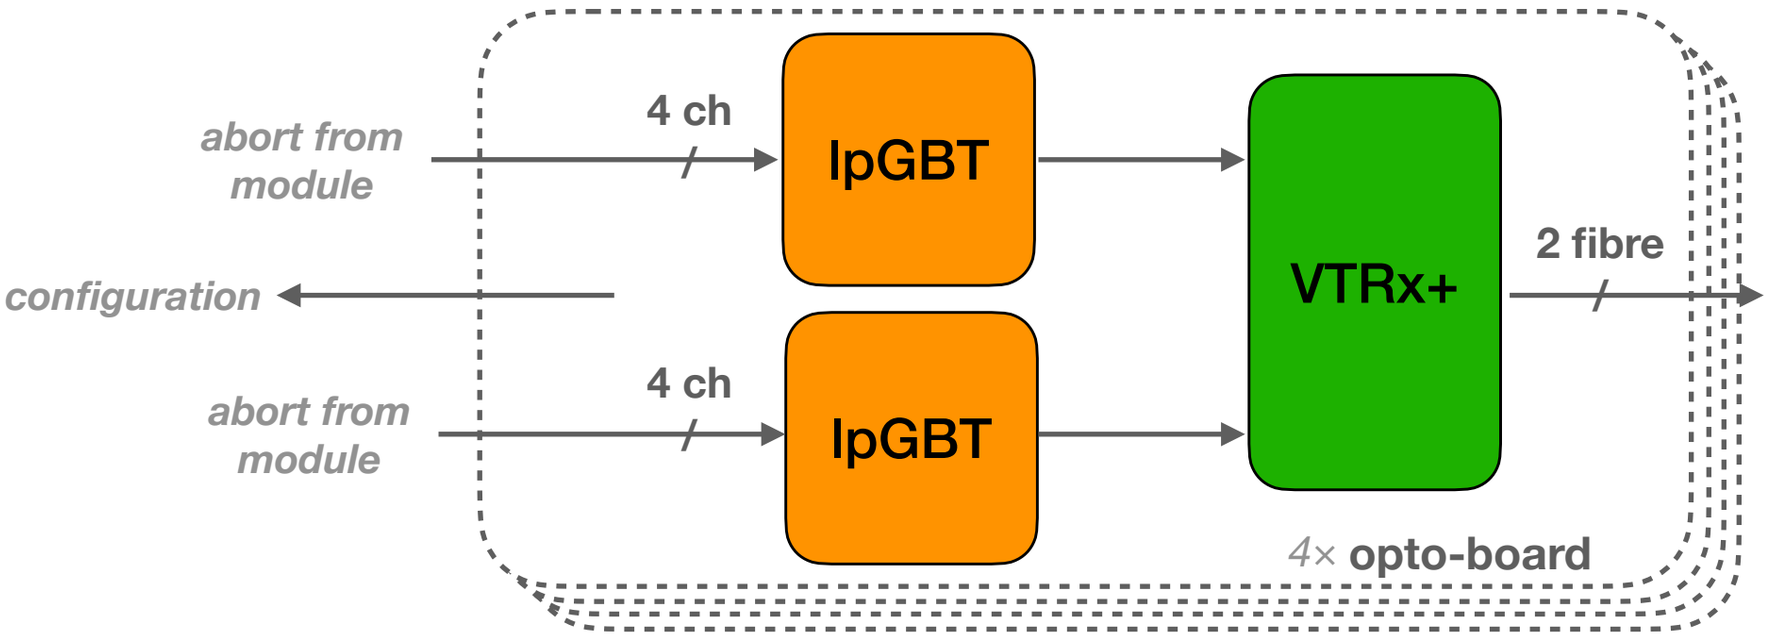
- higher sensitivity
- crucial aspects: signal stability

Module - data source



- 8 identical modules
- each module serves $abort_a$ & $luminosity_l$
- 4 ch_a+8 ch_l / module

Optoboards - data sampling



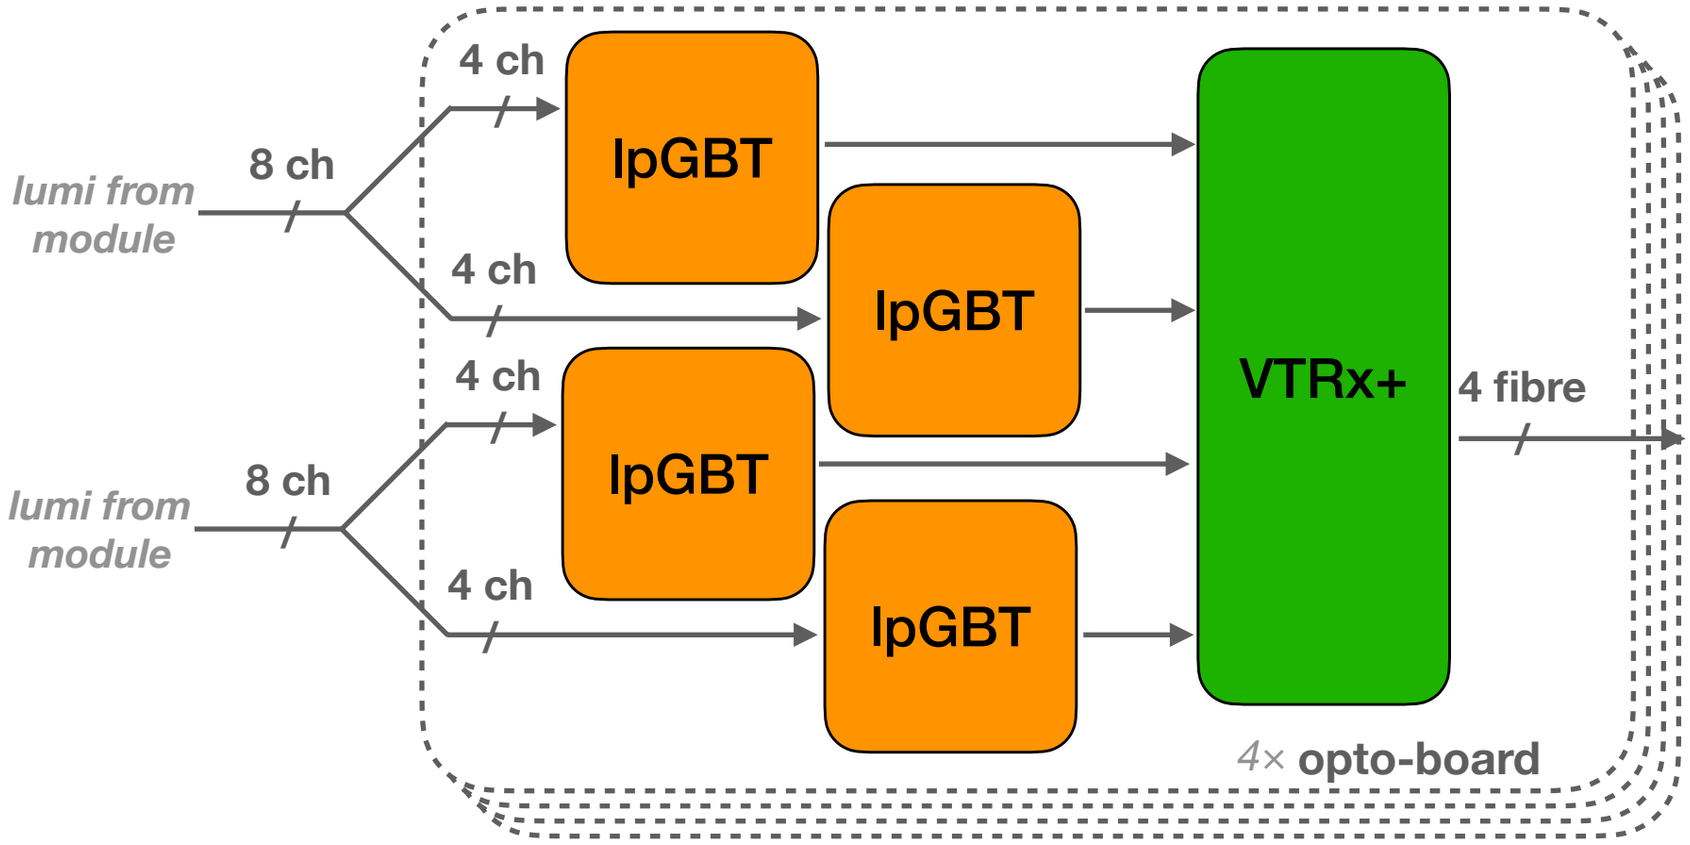
functionalities served separately

each module provides:

- 1 abort fibre
- 2 lumi fibres

system wide:

- 8 abort fibres
- 16 lumi fibres



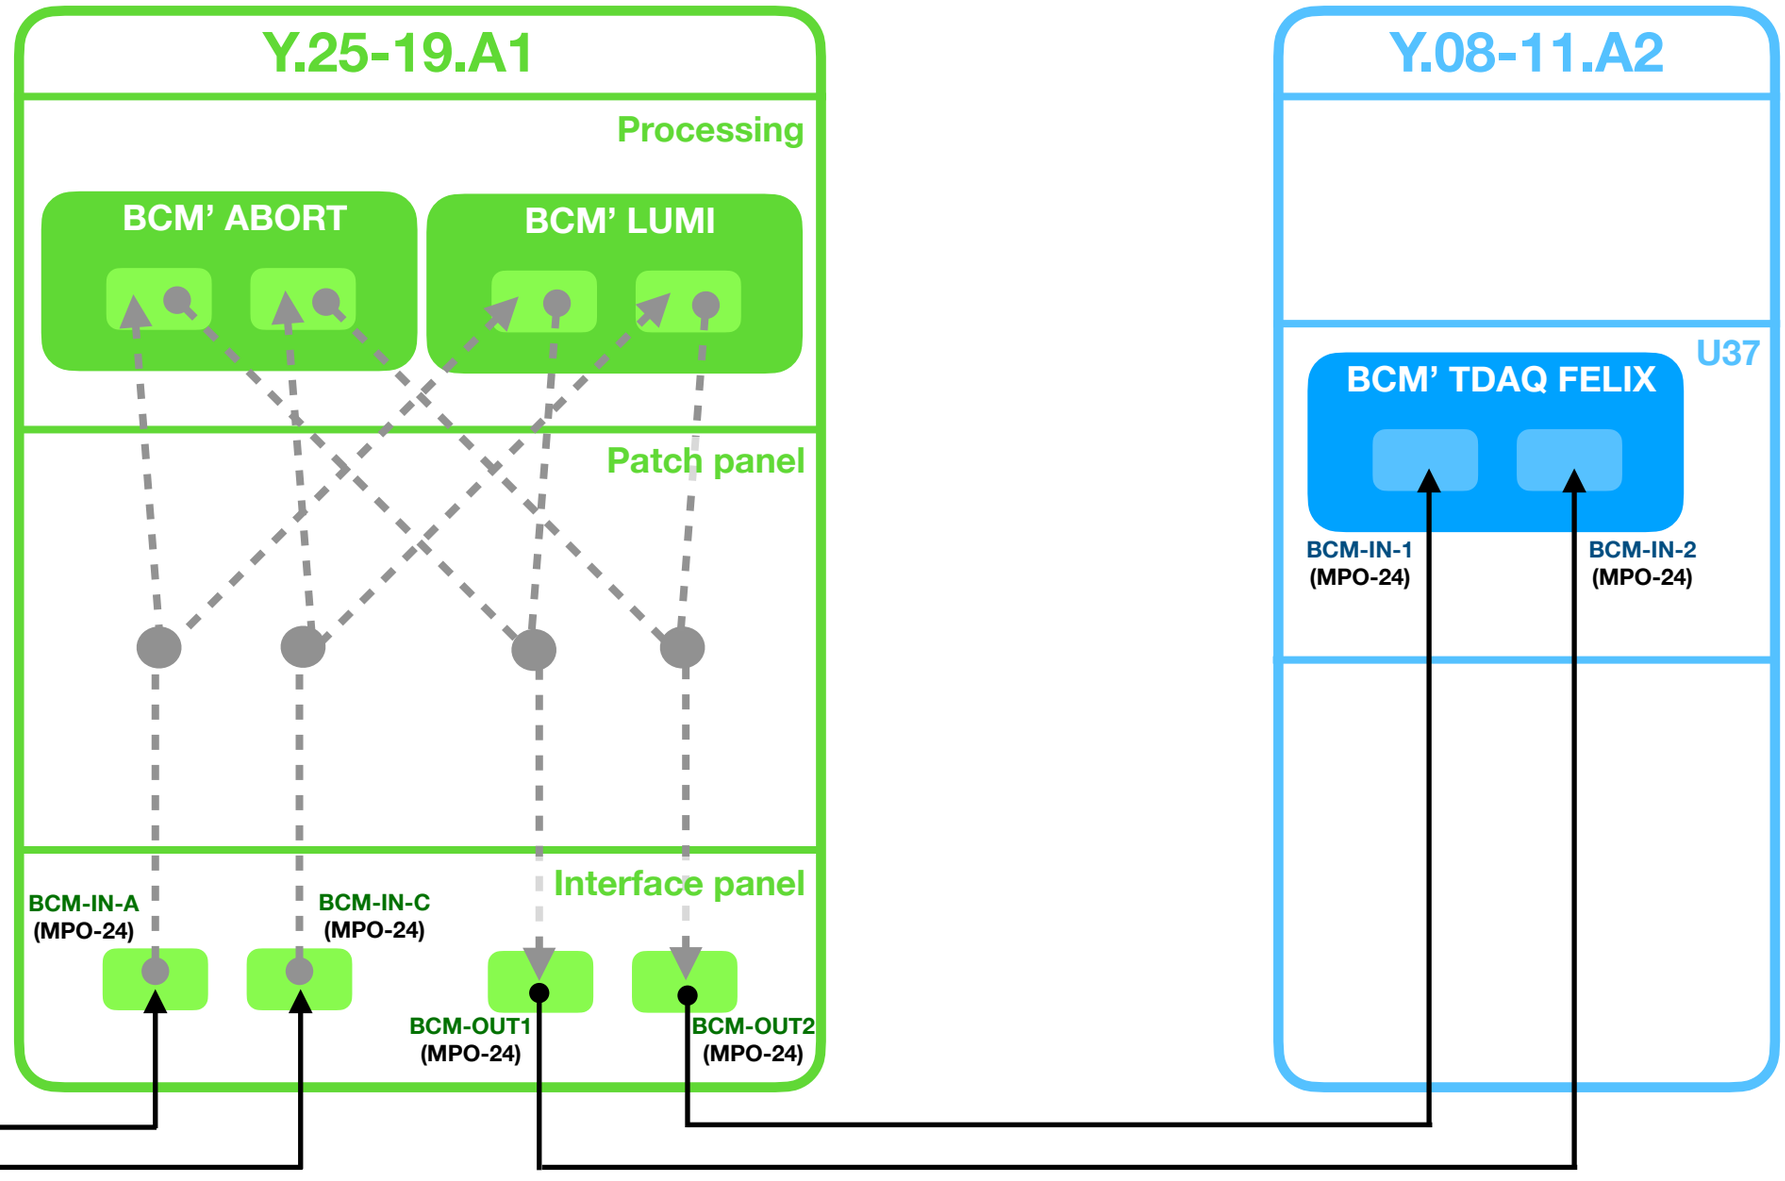
USA15 readout chain

system's speciality:

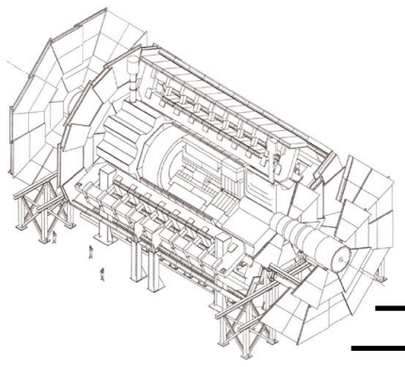
- front-ends do not encode
- non-trigger based measurements

solution :

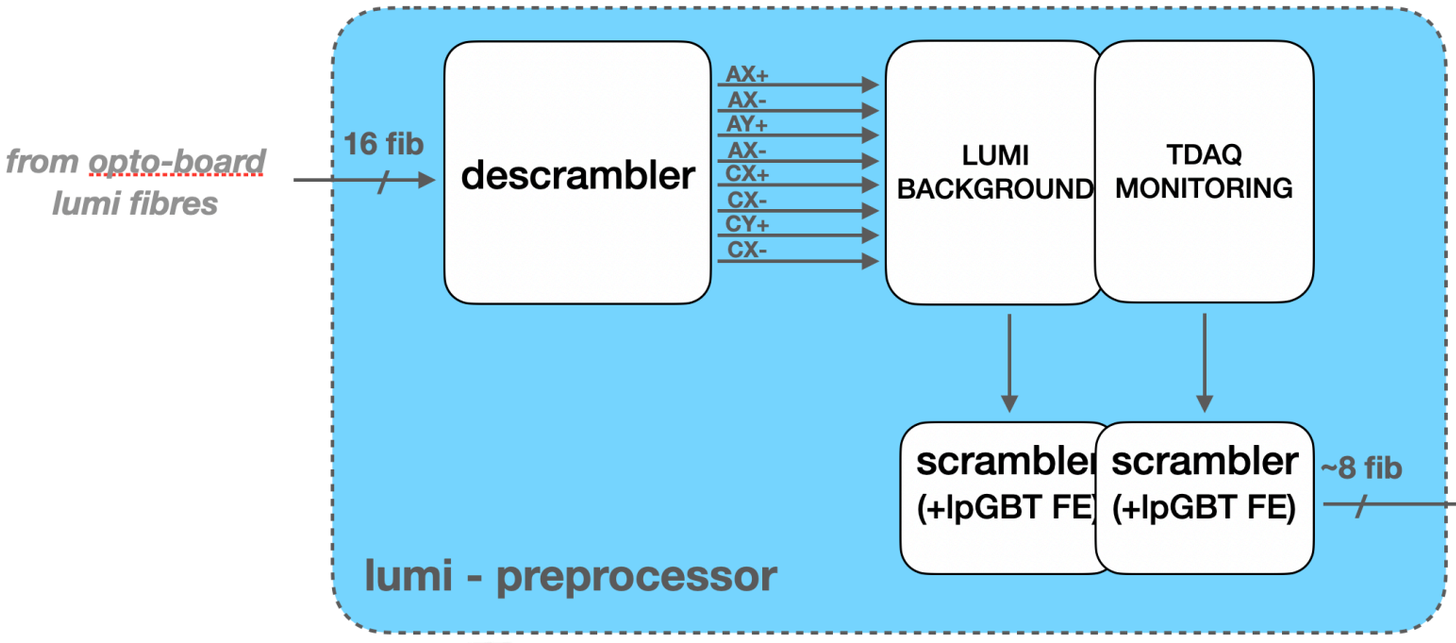
- dedicated real-time processing for abort
- dedicated preprocessing for lumi



DETECTOR/OPTO-PANELS



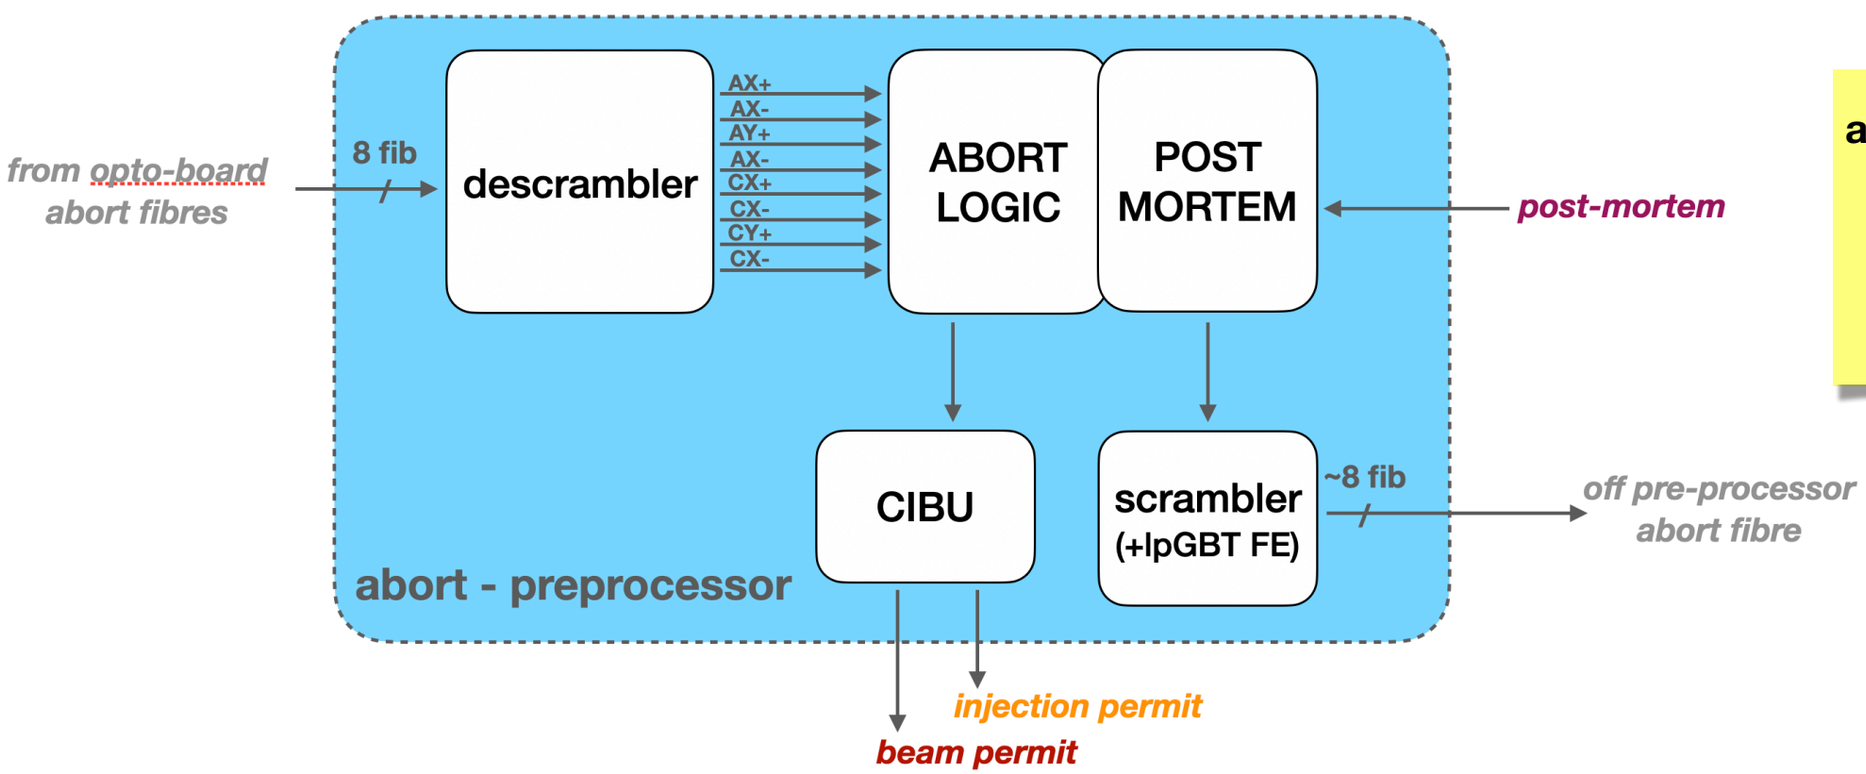
Pre-processing FW



data for ATLAS:

- data digested and reformatted
- all data combined
- from TDAQ side it looks like lpGBTs, compatible with standard TDAQ

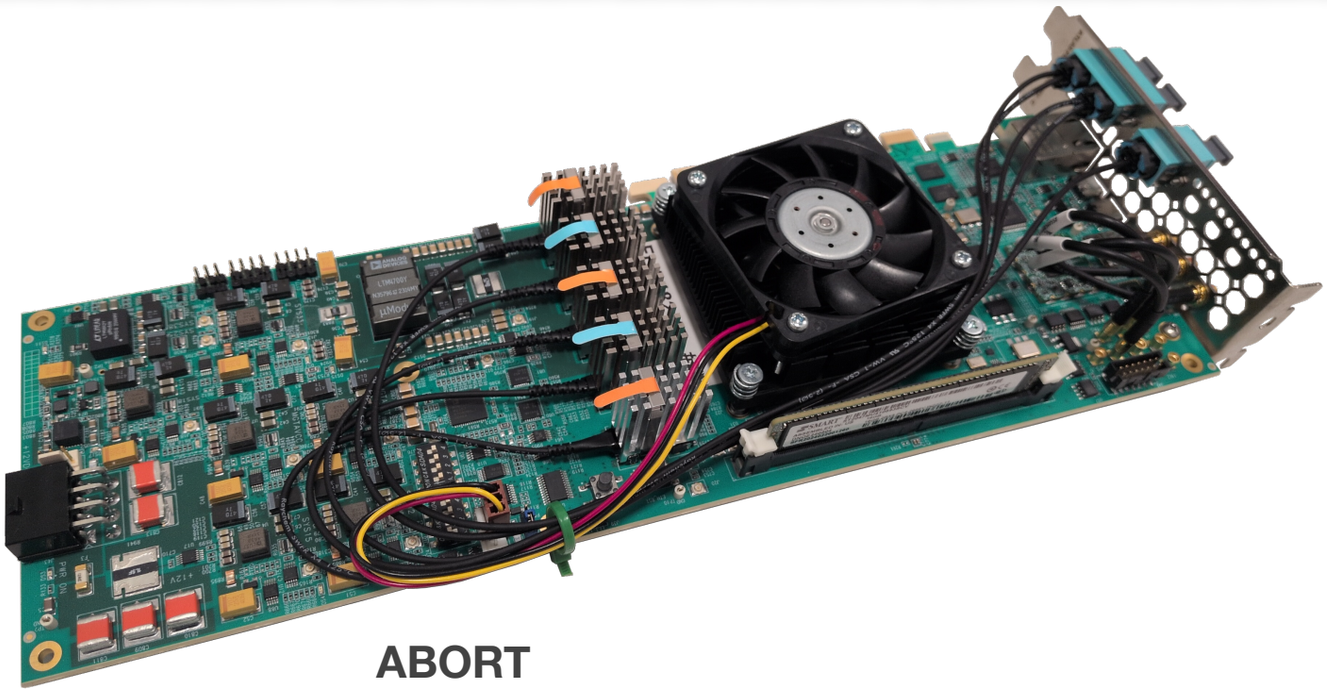
off pre-processor
lumi fibre



abort related signals:

- dedicated signals for CIBU interface
- 2 signals needed

Pre-processing HW



ABORT

- single logical stream serving post-mortem data
  - high volume (all raw data transferred)
  - low processing power needed (only few bits need to be monitored to sense for PM events)
- data stored in a cyclical buffer
- on detected PM flags the actual data processing would be done (heavy, but rare, not extremely time critical)

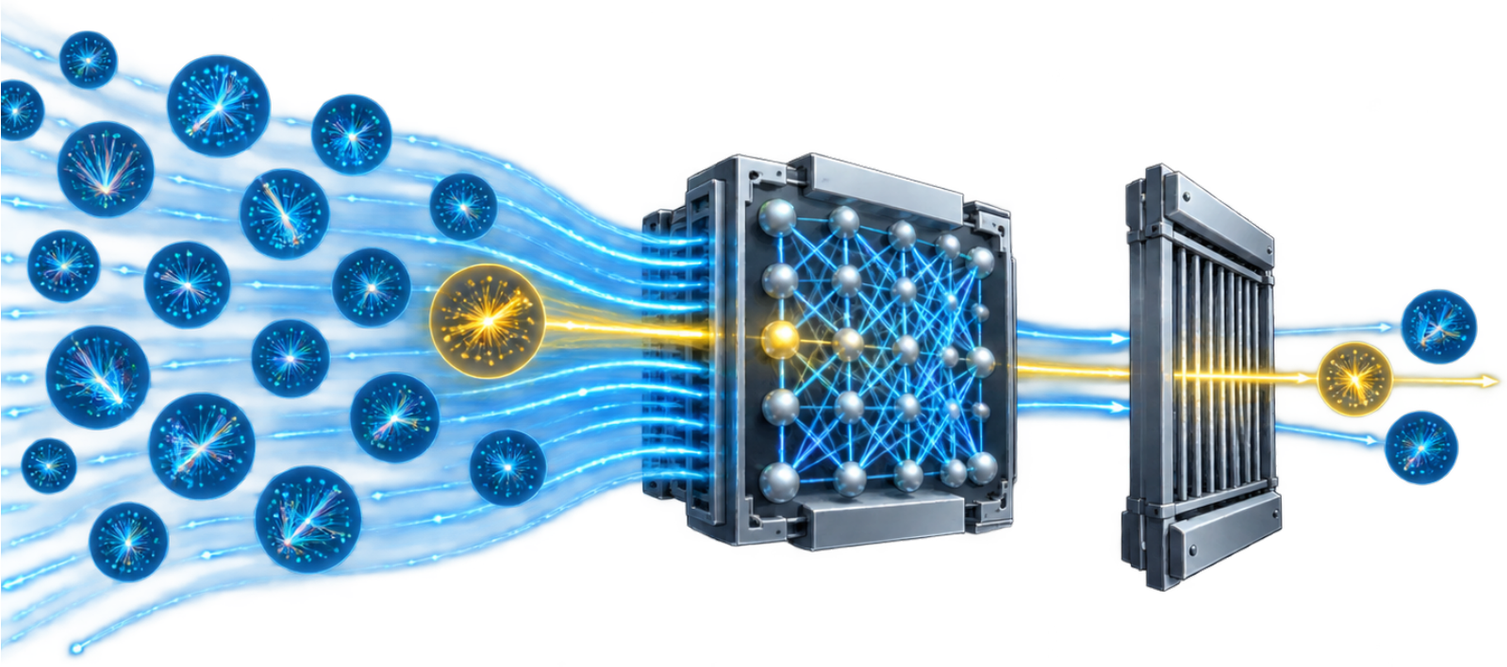
abort preprocessor:

- hosted on FLX-182 hardware - reusing standard ATLAS FELIX
- hosted within dedicated container with no OS
- only power needed

LUMI

- multiple streams served
  - luminosity
    - data already heavily reduced - low volume but constant processing
  - background & monitoring
    - low volume of heavily digested data
  - ATLAS events (standard TDAQ)
    - high volume, all pre-digested data transferred

Learning before filtering



Learning before filtering



Source: https://a3d3.ai

- experiments filter data using static triggers
- stored data is strongly biased
- selection depends on prior theoretical assumptions
- precise detector modelling is required for trigger studies

Can we learn from unbiased pre-trigger data streams?

- identify interesting regions via model-agnostic anomaly detection
- use unsupervised models due to lack of labels
- adapt continuously to changing detector conditions

Learning before filtering



Learning in hardware

Hardware constraints:

- extreme input bandwidth
- limited resources
- limited power for cooling
- minimal data storage & optimal movement

Methodology constraints:

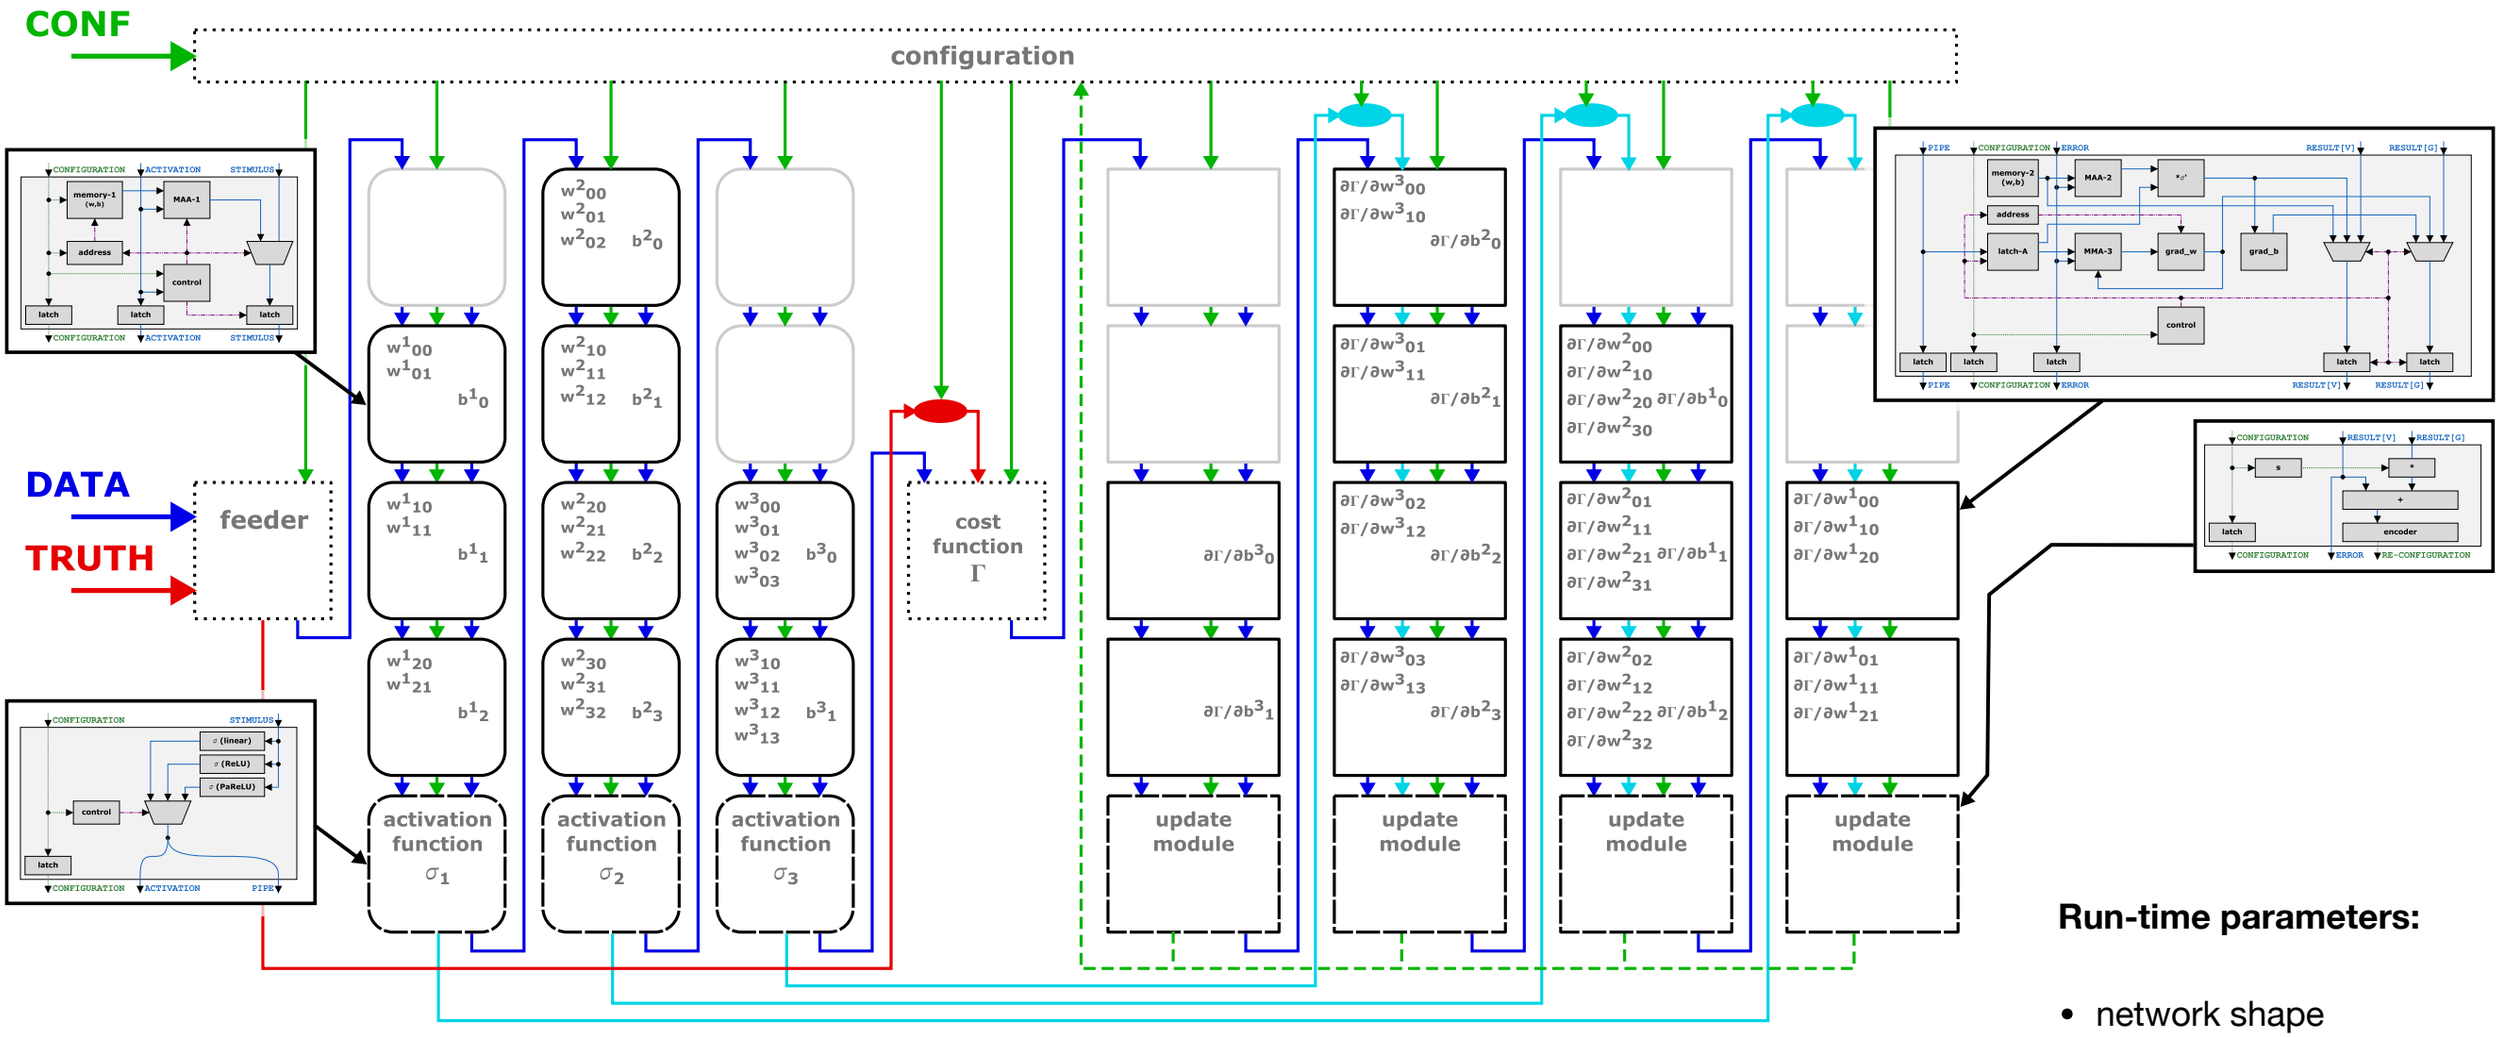
- truth labels do not exist
- detector operating conditions and systematics can change
- incoming data distribution as the only source of knowledge

First milestone - standalone training of neural network

- fully connected neural-network provides the essential training operation
- most demanding part of computing - measuring resource usage and computing throughput
- Basic building block that can later be adapted for auto encoders, normalisation flows and more sophisticated anomaly detection

First milestone is to train simple NN.

Design implementation



Self contained block, that is easily scalable.

Design-time parameters:

- max. layers
- max. neurons per layer

Run-time parameters:

- network shape
- activation functions
- cost functions
- initial parameters

Real life verification



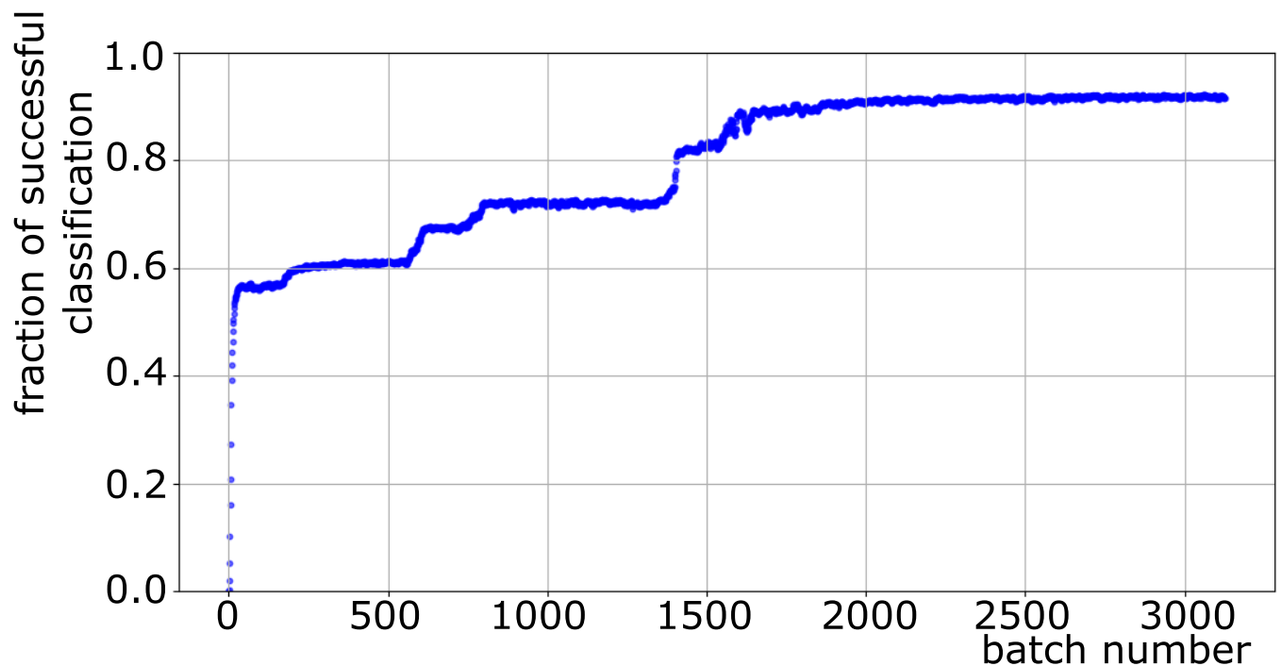
Validation done:

- 3 independent methods:
  - simplified sw
  - fw simulation
  - hw verification
- bit-compatible results

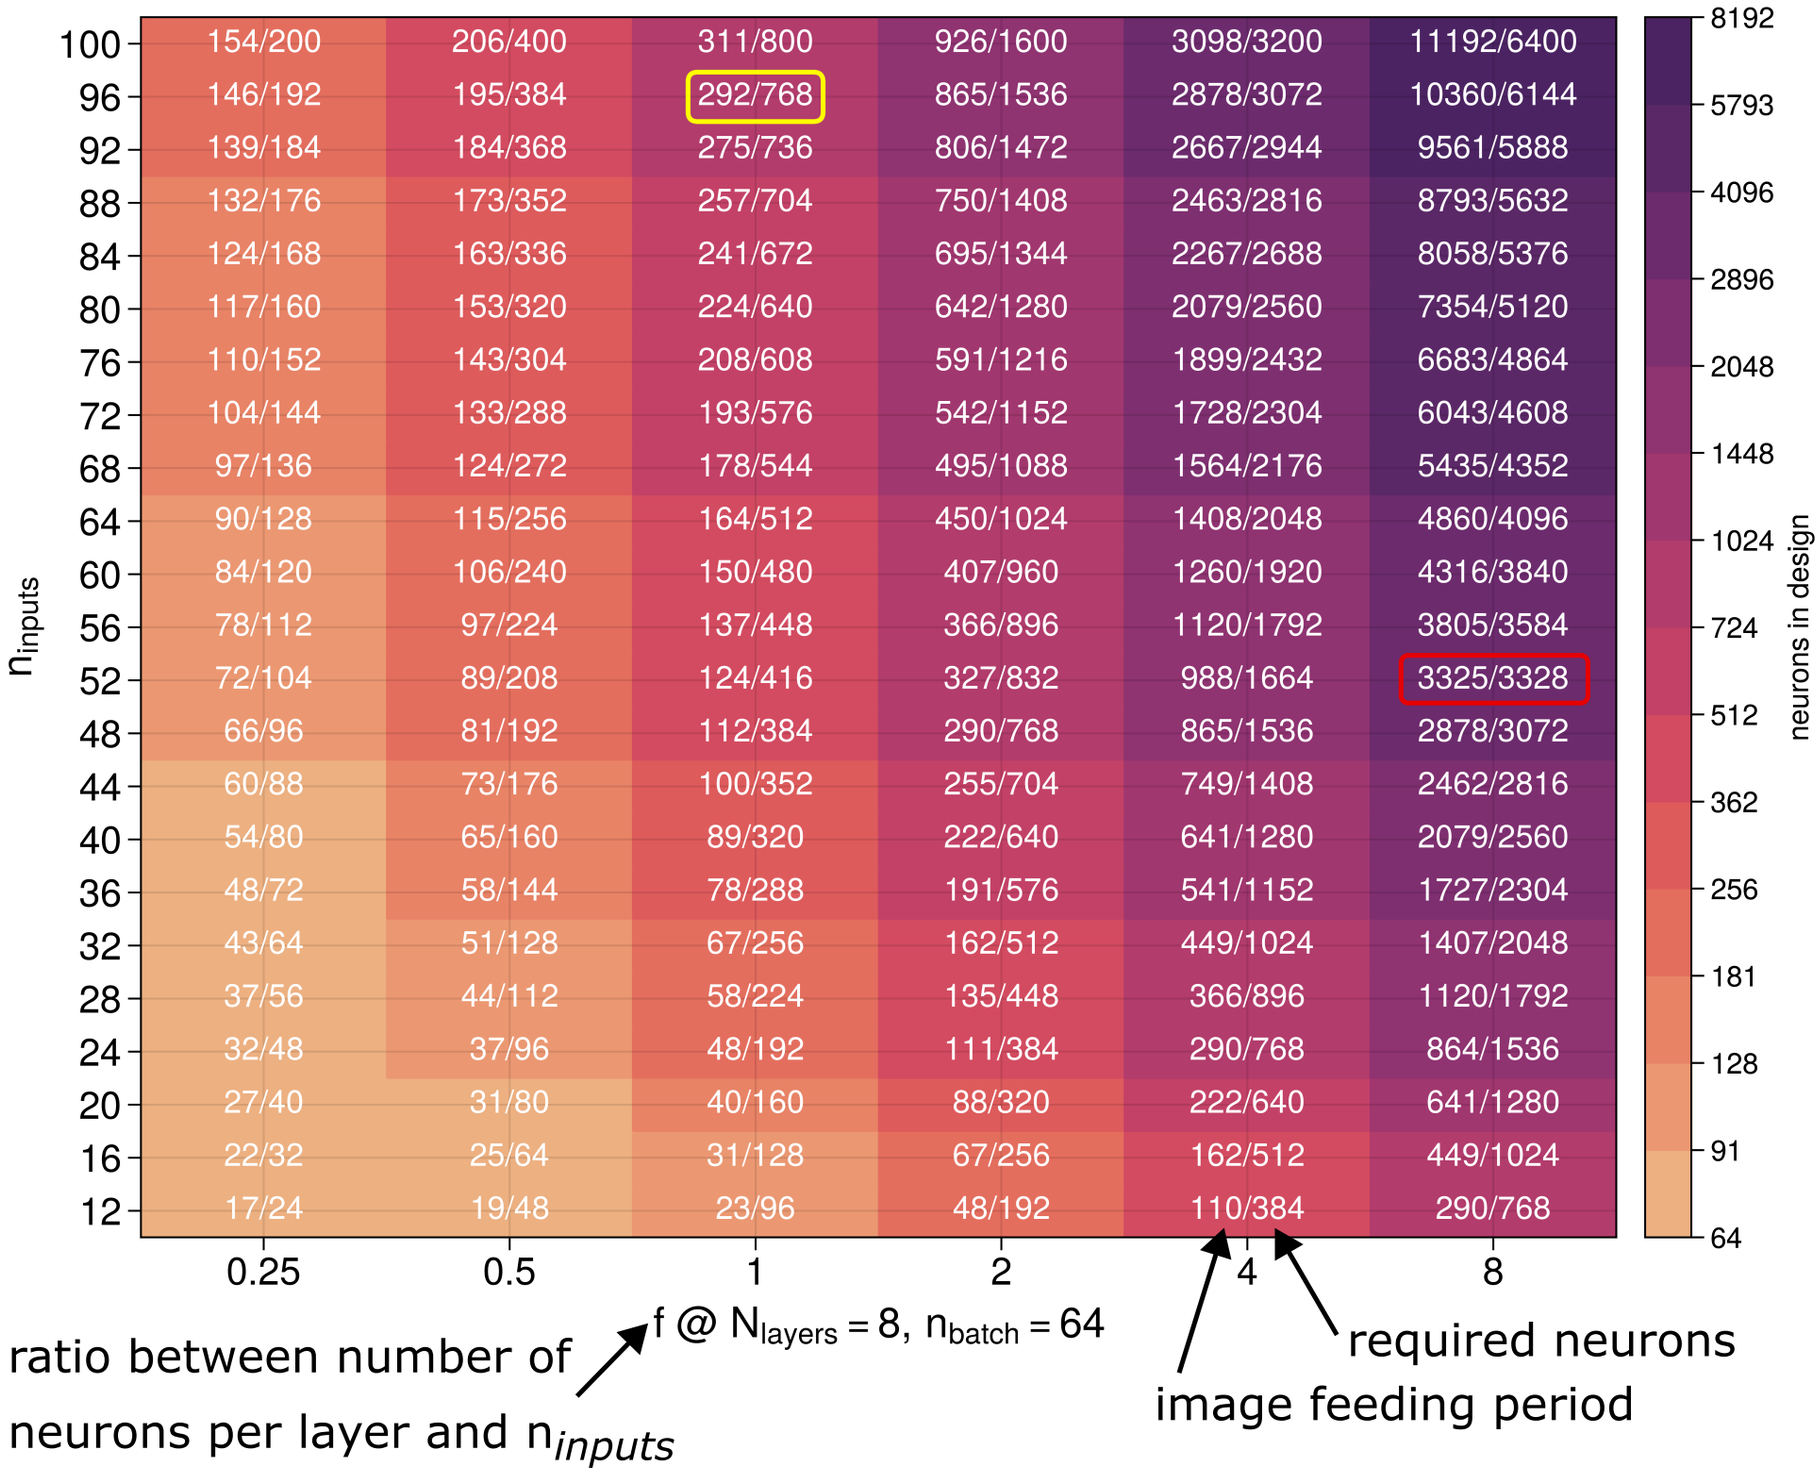
Demonstrated expected learning performance with no loss of precision.

Hardware implementation done:

- full design test with 151 neutrons (6 inputs & 7 outputs)
- versal VM1802 FPGA-fabric @ 160MHz
- uses 32-bit floating point (DSP blocks)
- resource usage per neutron: ~4DSCP58 + ~ 3 BRAM18
- could scale up to ~3000 neutrons in Versal family



Applicability landscape



Future application: detector-level anomaly detection

- select hardware-compatible unsupervised algorithms
- extend the training engine to anomaly detection
- integrate anomaly scoring within the detector readout
- validate robustness and real-time performance
- measure the effect on physics sensitivity

Thank you!

THANK YOU
FOR YOUR
ATTENTION