The new hardware Trigger Processor at NA62 experiment: status of the system and first results

Ottorino Frezza

(INFN Roma)

27th International Conference on Computing in High Energy & Nuclear Physics
CHEP 2024
Krakow – 21th October 2024

NA62 Experiment

- Fixed target experiment located in the North Area of CERN
- Measurement of very rare decay BR: $K^+ \rightarrow \pi^+ \nu\bar{\nu}$
- Search for new physics
- Nominal Beam Rate 750 MHz
  K^+ rate: 45 MHz
  K^+ decays in fiducial volume circa 5%

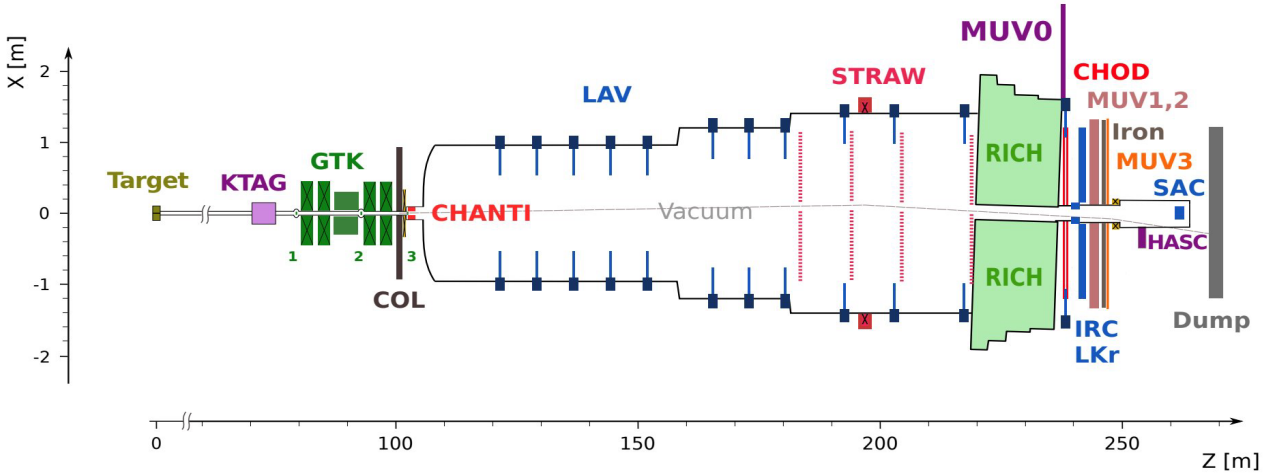

- Last result CERN press release :
https://home.cern/news/press-release/physics/na62-experiment-cern-observes-ultra-rare-particle-decay

NA62 TDAQ system

Multi-level trigger

- Level-0 (L0TP, Level 0 Trigger Processor) HW trigger on FPGA (input 10 MHz → output 1MHz)
- Level-1/2 software trigger running in DAQ farm (1MHz → 100kHz)

DAQ

- Data bursts are ~6s long
- Some detectors' primitives (generated from TEL62 readout boards) are sent to L0TP over 1GbE UDP channels
- L0TP generates triggers with max latency of 1ms
- 40 MHz synchronous operation

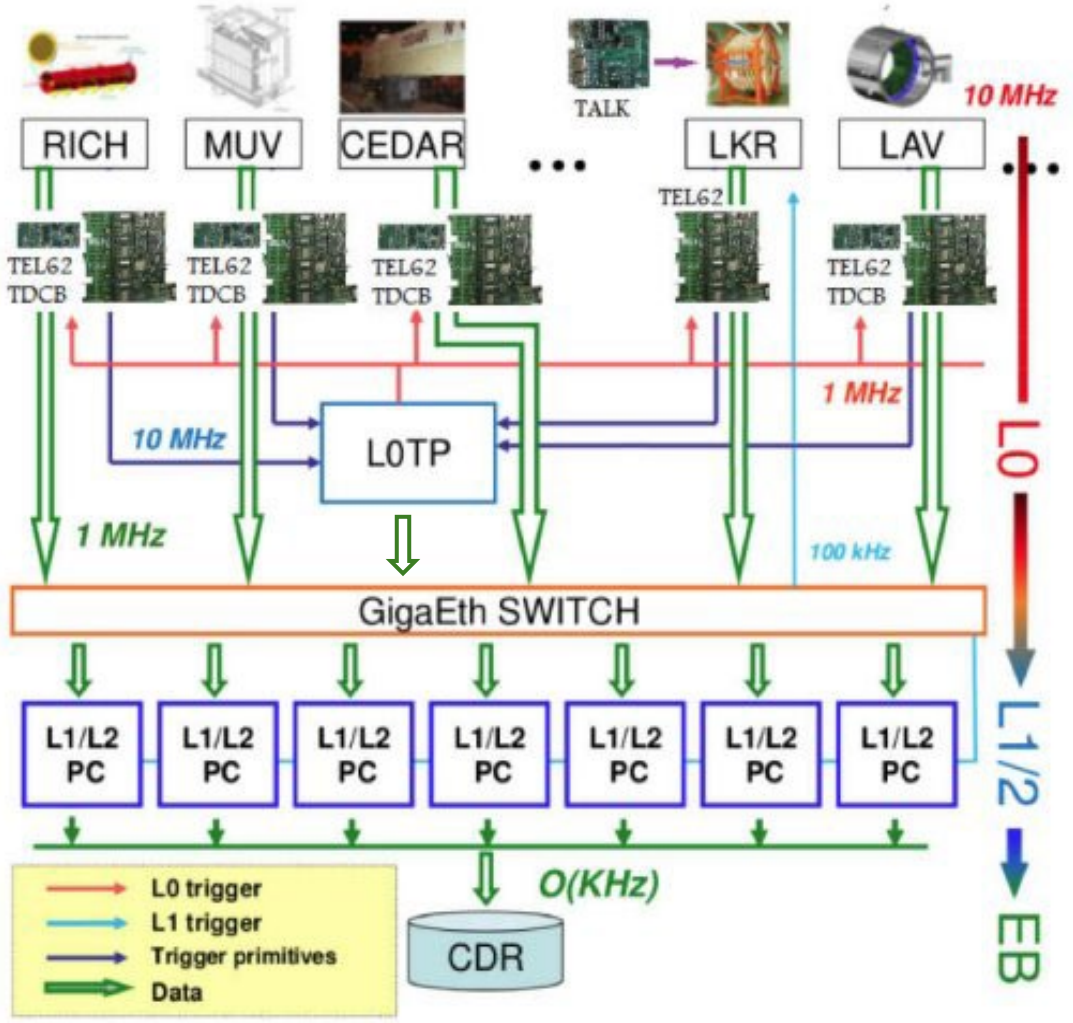

L0TP+ : Short history

- **2018**
  - *Trigger update proposal*
- **2019**
  - *HW procurement, testbed and simulation*
- **2022**
  - *Parasitic mode test*
  - *Beam test online in November*
- **2023**
  - *Dry test in January and in March*
- **Today Status:**
  - ***Integration completed***
  - ***L0TP+ Online***

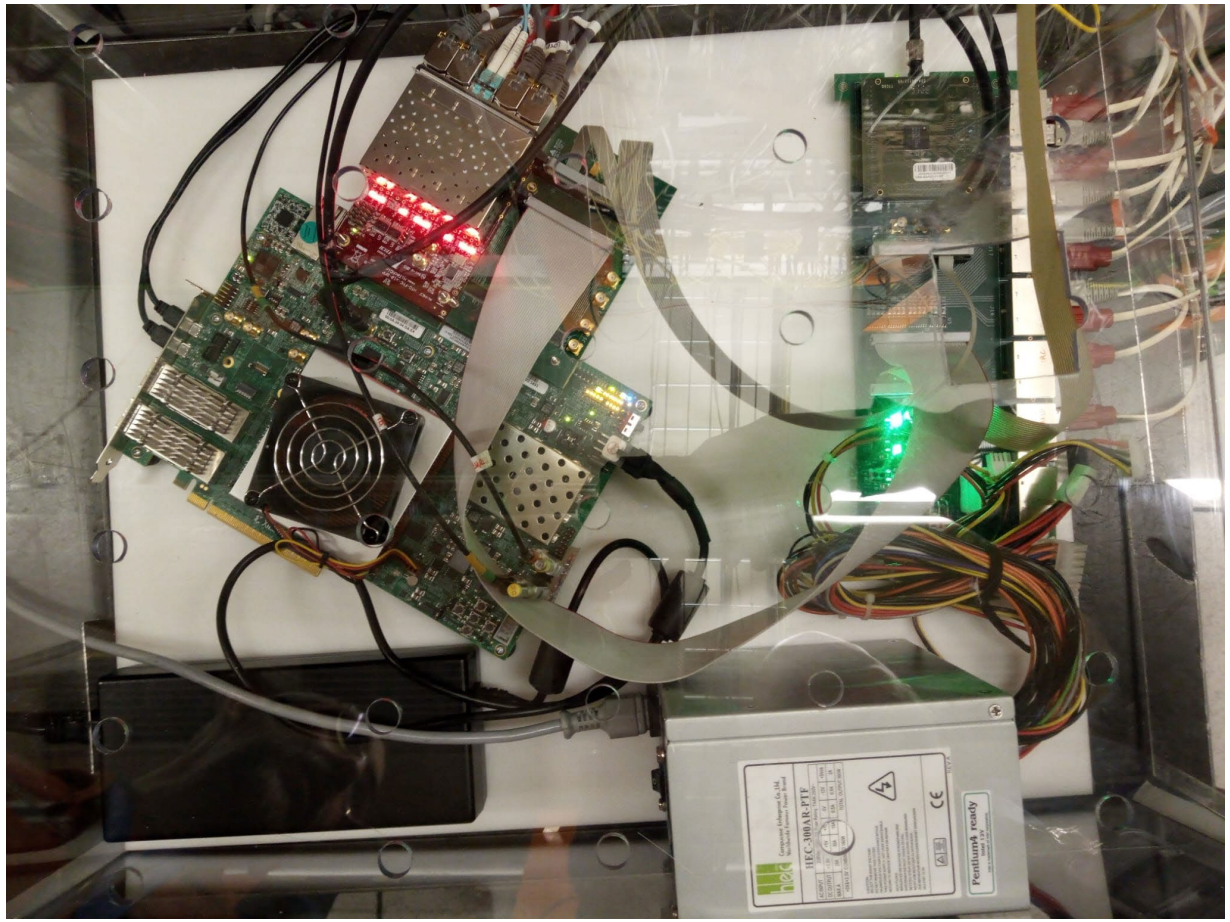

Hardware procurement

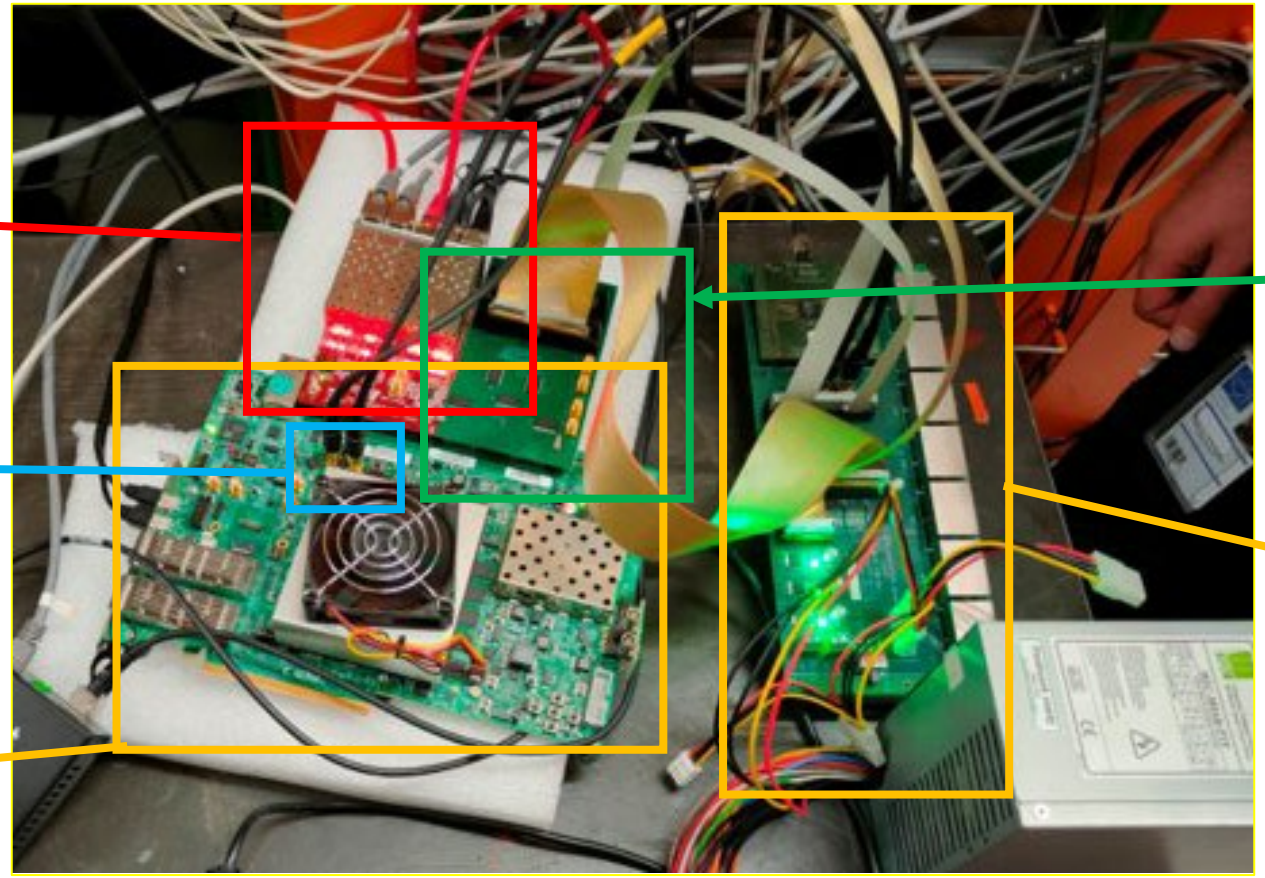

10-port SFP+(10x25Gb) FMC+ Module

System Clock @40MHz on SMA

VCU118 board equipped with Xilinx Virtex Ultrascale+

Custom Board to interface with TTC signals

Auxiliary board (to receive Clock, SOB, EOB, Trigger and CHOKE signals)

System Design

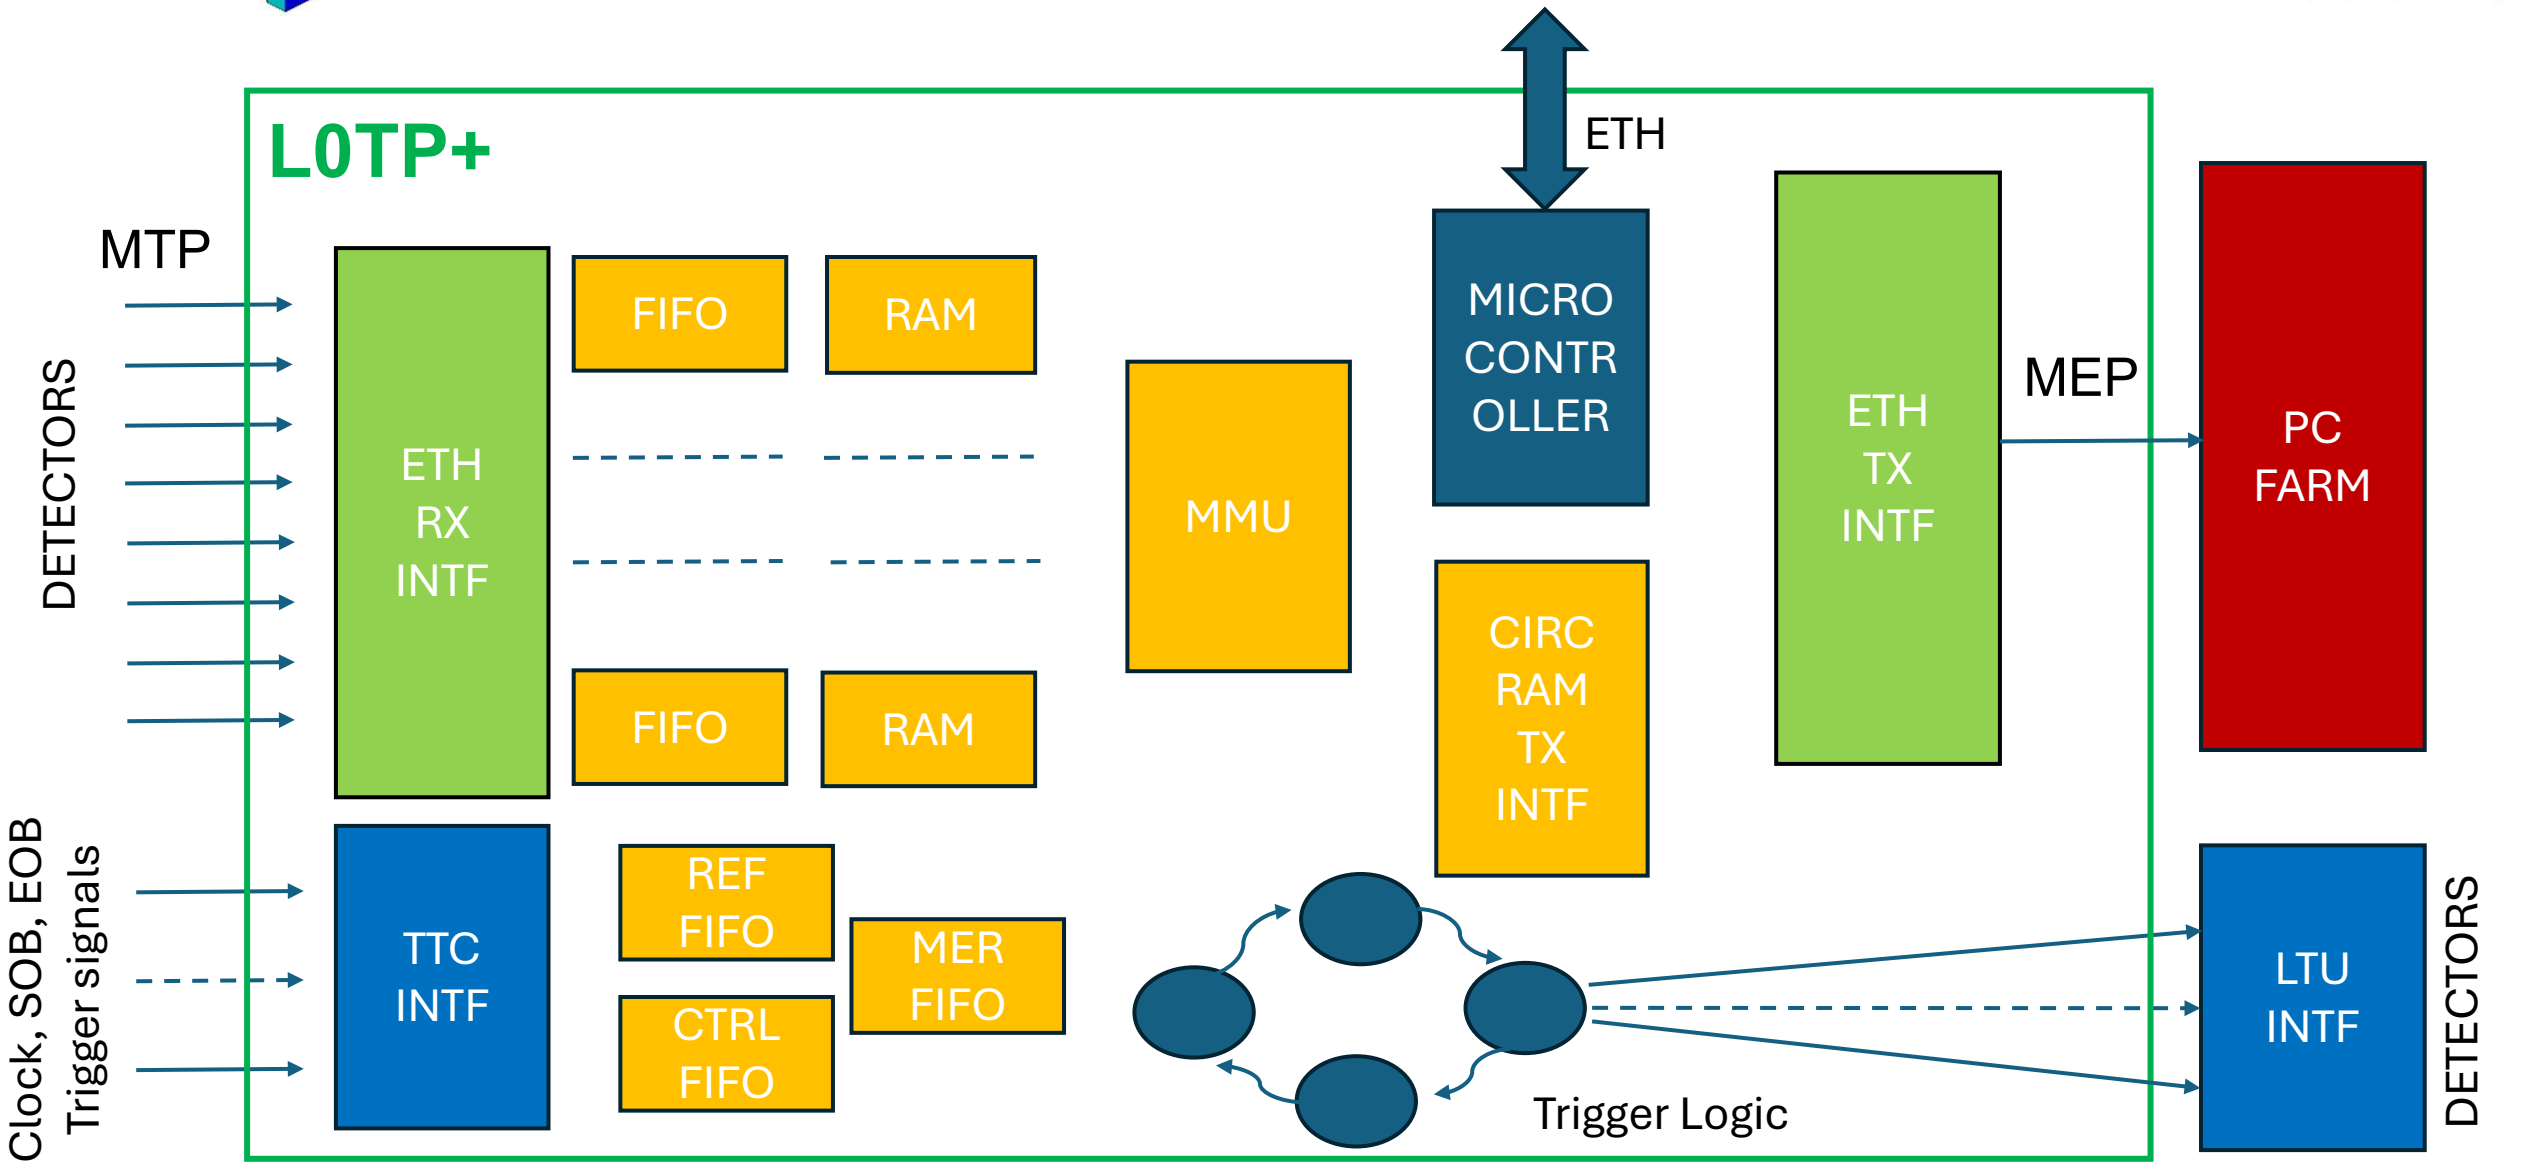

Validation and Test

- System validation performed in parasitic mode.
- Network TAPs (Test Access Points) to duplicate primitives and Triggers packets.
- Comparison of information about the generated trigger performed on dedicated server.

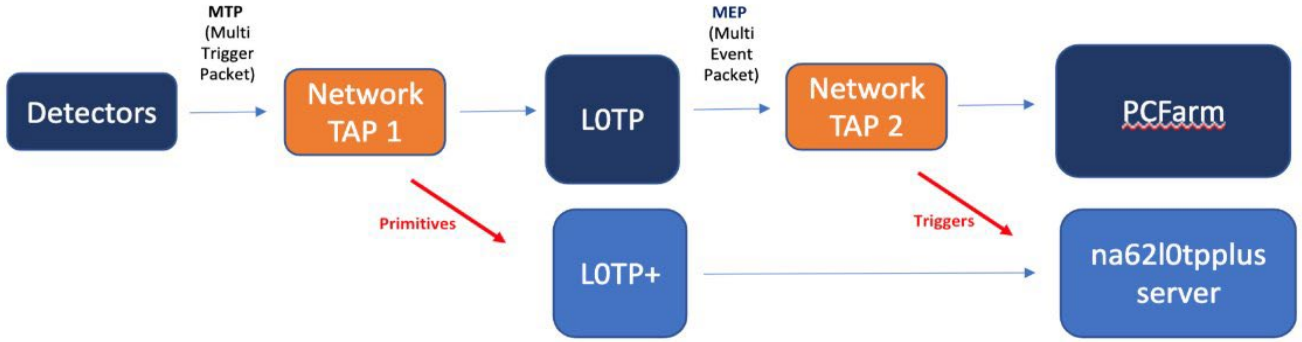

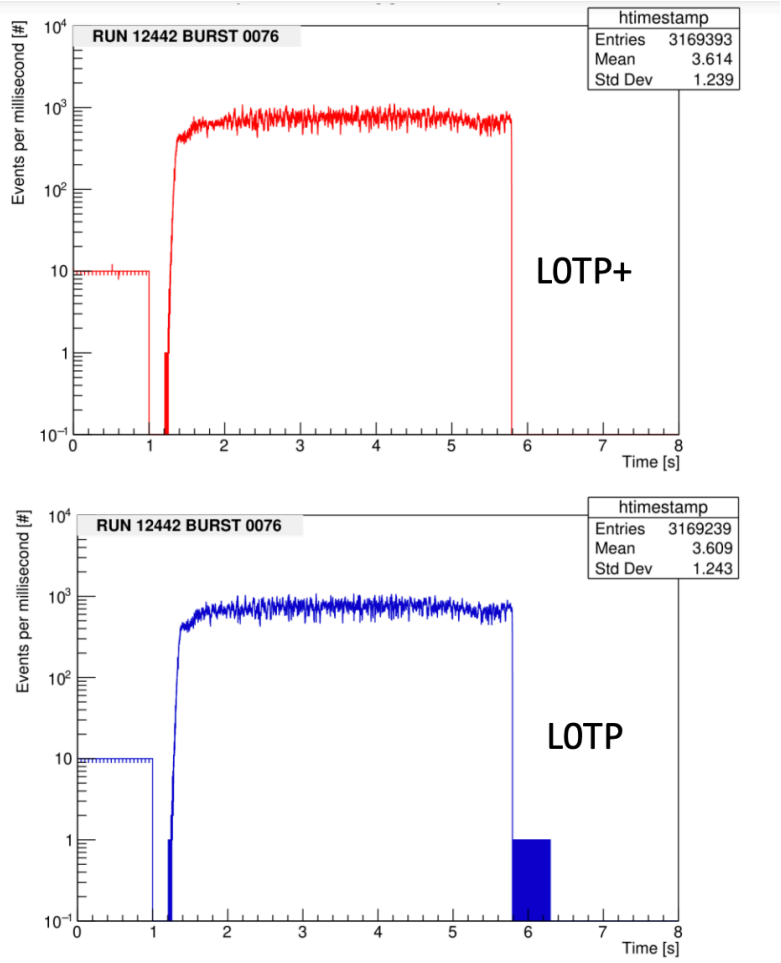

On the right, one Burst Trigger Timing reconstructed from UDP packets analysis received from na62l0tpplus Server. The profiles are the same except for calibration step at the end of the Burst.

Efficiency

Impact of L0TP+

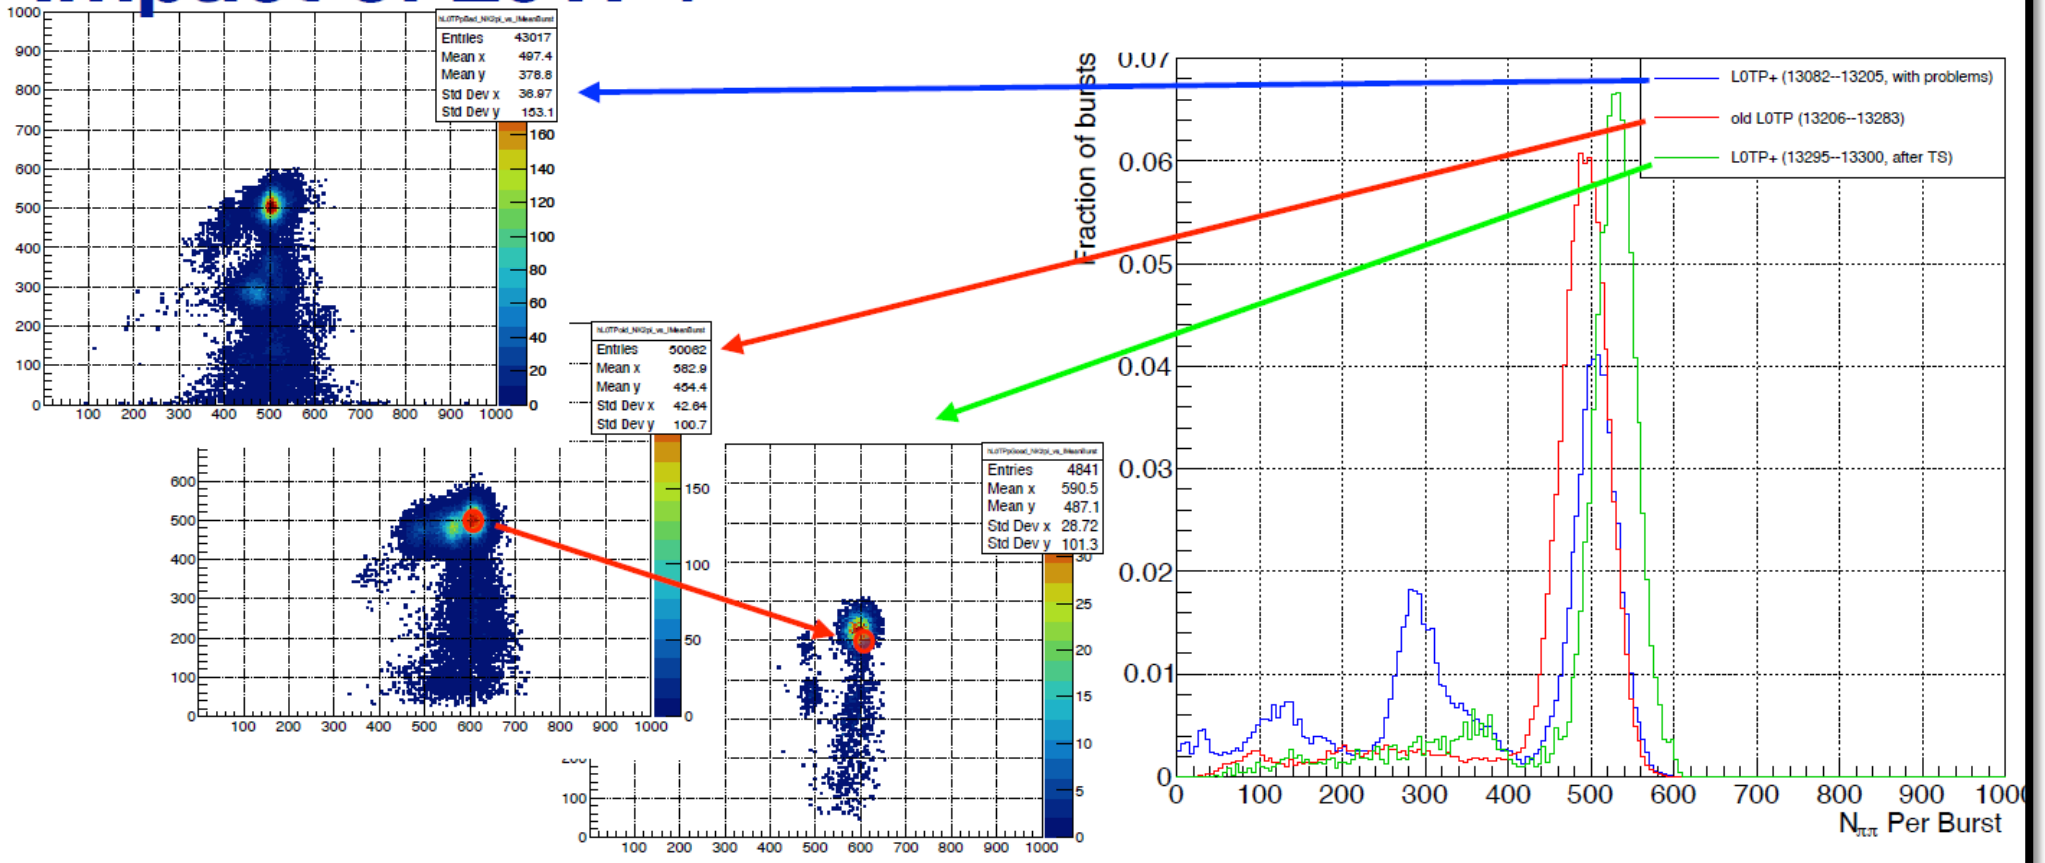

- Gain ~5% with L0TP+ . Slightly lower intensity but still higher Npipi with L0TP+.
- At start of run MANY bad events, solved with fw updates to L0TP+.

15

New functionalities

L0TP+ reproduces all L0TP functions but considering the huge amount of FGPA resources (only 30 % BRAM, 17 % LUT used in L0TP+) there is room to add several capabilities to the original design.

- **DATA LINKS**:
  the system is able to support ten 25GbE links through the FMC+ daughercard, and additional QSFP28, and FireFly ports can be used to connect additional data links from the detectors via 100 Gbps low latency links.
- **MICROCONTROLLER:**
  a 32-bit MicroBlaze Soft-Core Micro Controller was integrated for debug and configuration purposes. Applications can be deployed onto it either bare metal or by Xilinx Petalinux.
- **PCIe HOST INTERFACE**
- **STREAM PROCESSING MODULE:**
  with the outlook of processing primitive streams and thus improving the efficiency of the trigger (e.g. online PID in RICH via HLS4ML Neural Networks)

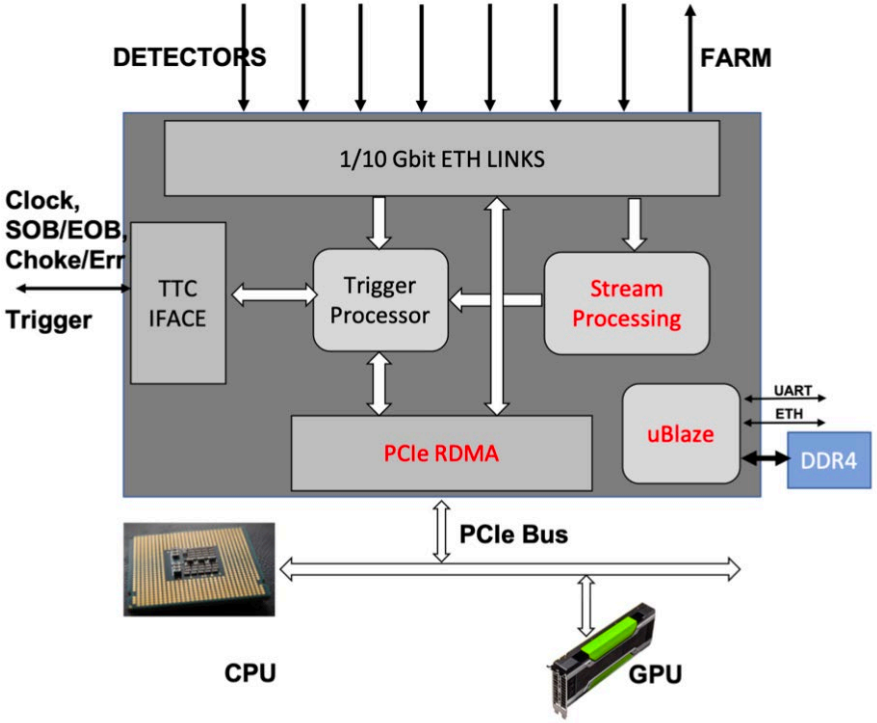

RICH PID using NN on FPGA at L0 Trigger

FPGA-RICH : Reconstruction of Cherenkov rings, produced by electrons in the RICH detector, using an AI algorithm on FPGA.

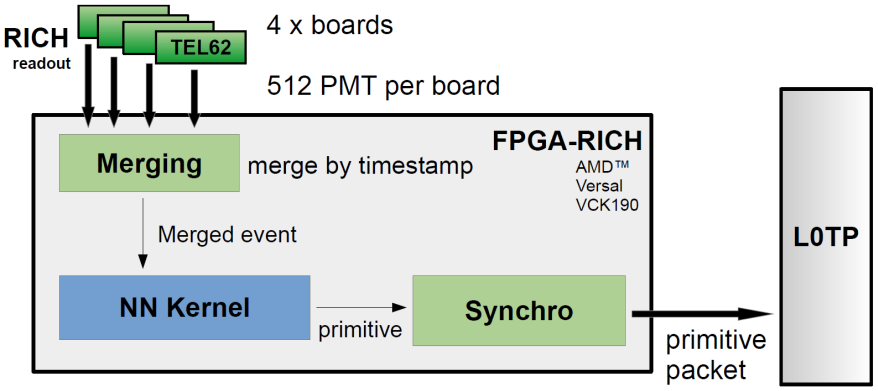

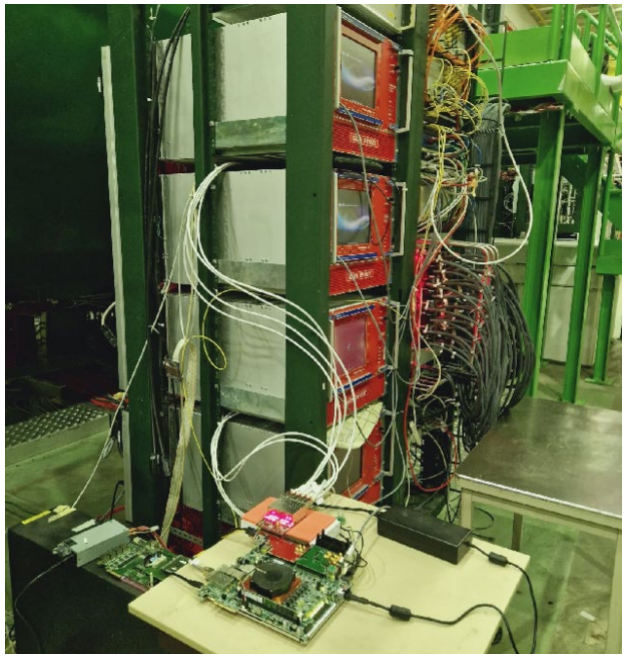

Prototype based on Alveo Board

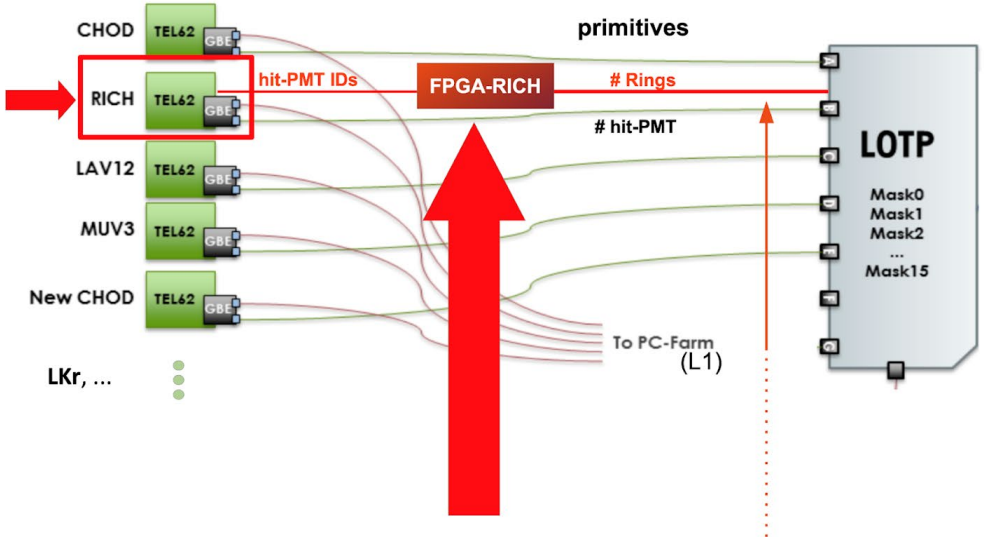

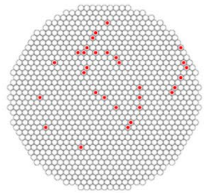

RICH primtives: # of hit-PMTs

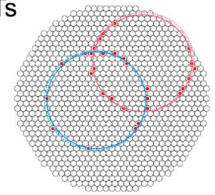

FPGA-RICH primtives:
# of rings (0, 1, 2, 3+)
(more in the future, e.g. # of e^-)

RICH PID using NN on FPGA at L0 Trigger

- NN based on dense layers that process an "unrolled" image obtained from mapping the RICH PMTs positions on the RICH disk plane (3 Dense layers: 64 (in) → 64, 64 → 16, 16 → 4 (out))

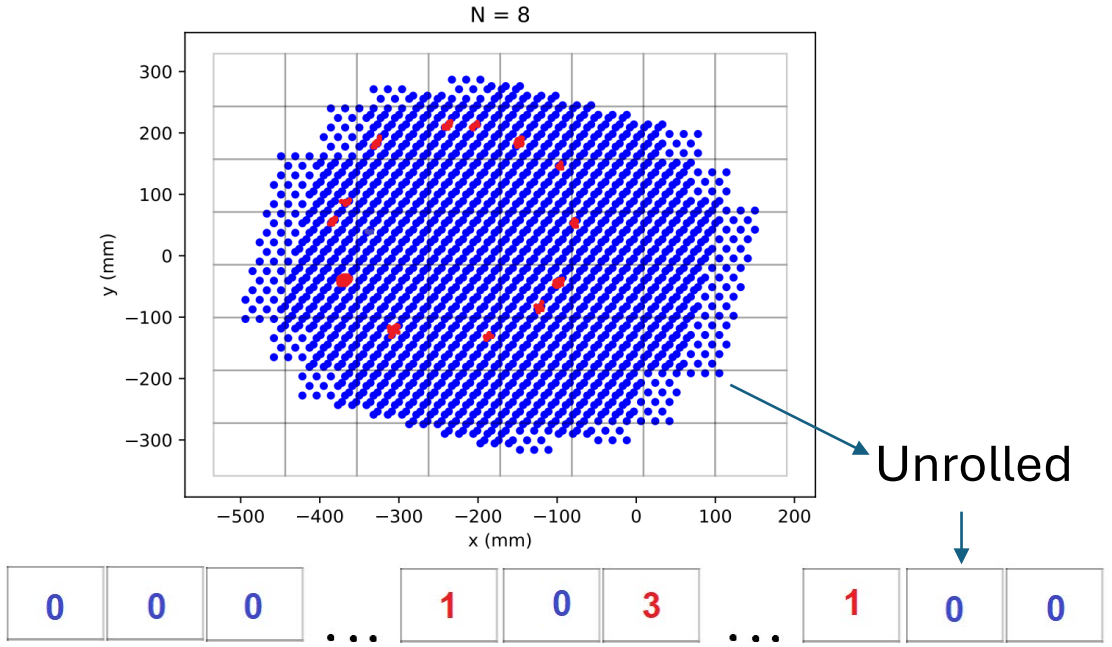

Classification performance

	Efficiency	Purity
0 Rings	88.88	95.04
1 Rings	88.90	86.48
2 Rings	76.34	72.22
>=3 Rings	77.12	84.62

See Perticaroli's talk on Thursday Track2
FPGA-RICH: A low-latency, high-throughput online particle identification system for the NA62 experiment

- **NN KERNEL *average throughput*** (depends on number of hits) **~20 MHz**, ***latency* = 29 clock cycles @**150 MHz (**193 ns**).
- **FPGA-RICH Utilization** LUT = 14%, BRAM = 3%, DSP = 7% FF=6% (VCK190).
- Full pipeline tested with artificial MGP streams in LAB.
- System deployed at the experiment, at the end of the RUN trying to find a workaround to TEL62 fw bugs.

New Trigger algorithm

The presence of halo muon causes random rejection of genuine events in some Trigger Masks

- mask0: RICH3*NewCHOD*!MUV3
- mask1: RICH3*UTMC*!QX*!MUV3*!LKr

The **!MUV3 on all MUV3 surface** can be reduced to quadrants to reduce the over veto probability

Introduce a **new logic signal at L0TP+** level:
Qi!Mi **= QI1*!MQI1+ QI2*!MQI2+ QI3*!MQI3+ QI4*!MQI4**

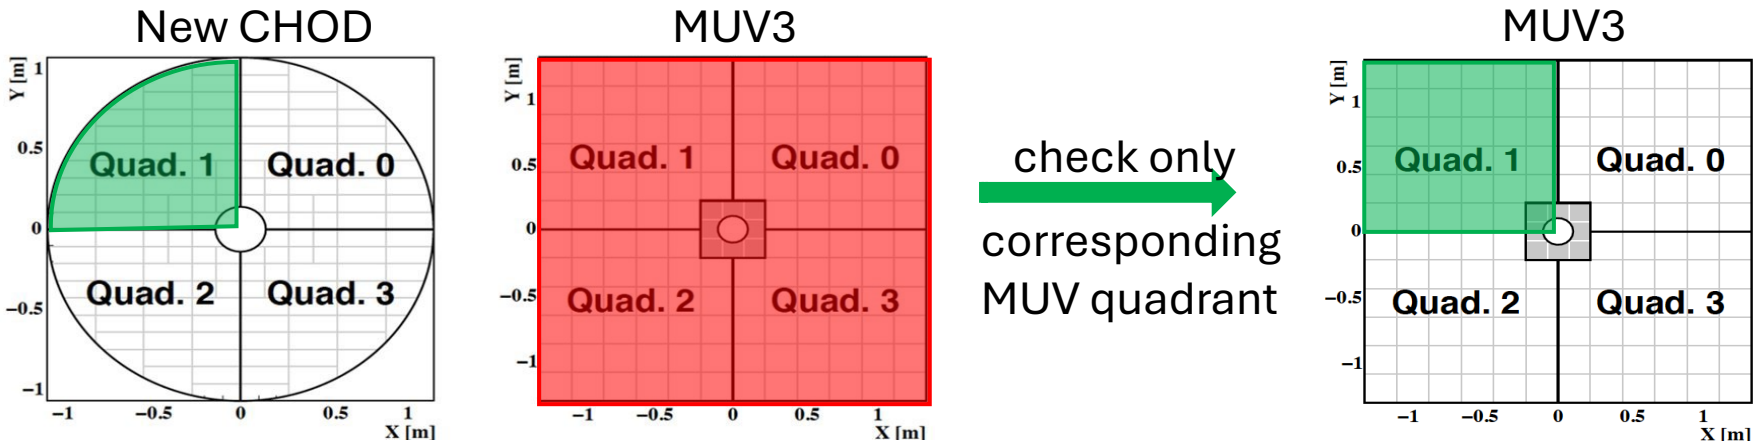

Try to implement separated **Qi veto conditions quadrants by quadrants** avoiding to veto on halo muons passing into other quadrants

Conclusion

- The system is on-line and the Trigger efficiency is better than previous system.
- The available resources on the new board allow implemtation of new features.
- New Trigger algorithm are under test with promising results
  - FPGA-RICH on L0TP+
  - New Trigger mask on MUV3 Veto detector
- Still two years of data taking left (2025, 2026)

Thank you!

BACKUP SLIDES

Conclusions

- New study of $K^+ \to \pi^+ \nu\bar{\nu}$ decay using NA62 2021–22 dataset:
  - Improved signal yield per SPS spill by 50%.
  - $N_{bg} = 11.0^{+2.1}_{-1.9}$, $N_{obs} = 31$
  - $\mathcal{B}_{21-22}(K^+ \to \pi^+ \nu\bar{\nu}) = \left(16.0^{+5.0}_{-4.5}\right) \times 10^{-11} = \left(16.0\ (^{+4.8}_{-4.2})_{stat}\ (^{+1.4}_{-1.3})_{syst}\right) \times 10^{-11}$

- Combining with 2016–18 data for full 2016–22 results:
  - $N_{bg} = 18^{+3}_{-2}$, $N_{obs} = 51$ (using 9+6 categories for BR extraction)
  - $\mathcal{B}_{16-22}(K^+ \to \pi^+ \nu\bar{\nu}) = (13.0^{+3.3}_{-2.9}) \times 10^{-11} = (13.0\ (^{+3.0}_{-2.7})_{stat}\ (^{+1.3}_{-1.2})_{syst}) \times 10^{-11}$
  - Background-only hypothesis rejected with significance Z>5.
- **First observation of $K^+ \to \pi^+ \nu\bar{\nu}$ decay: BR consistent with SM prediction within 1.7σ**
  - **Need full NA62 data-set to clarify SM agreement or tension.**

2023–LS3 data-set collection & analysis in progress...

Abstract

- The NA62 experiment is designed to study kaon's rare decays using a decay-in-flight technique. Its Trigger and Data Acquisition (TDAQ) system is multi-level, making it critically dependent on the performance of the inter-level network. To manage the enormous amount of data produced by the detectors, three levels of triggers are used. The first level L0TP, implemented using an FPGA device, has been in operation since the start of data taking in 2016. To increase the efficiency of the system and implement additional algorithms, an upgraded system (L0TP+) was developed starting in 2018. This upgrade utilizes a high-end FPGA available on the market, offering more computing power, larger local memory, and higher transmission bandwidth. We have planned tests for a new trigger algorithm that implements quadrant-based logic for the veto systems. This new approach is expected to improve the main trigger efficiency by several percent. Extensive tests were conducted using a parasitic setup that included a set of Network TAPs and a commodity server, allowing for proficient comparison of trigger decisions on an event-by-event basis. The experience gained from this parasitic mode operation can be leveraged for the next datataking period as a development setup to implement additional features, thereby accelerating the TDAQ upgrade. After the testing period, the new system has been adopted as the online processor since 2023. Preliminary results on the efficiency of the new system will be reported. Integration with the new AI-based FPGA-RICH system, which performs online partial particle identification, will also be discussed

L0TP+: selected device

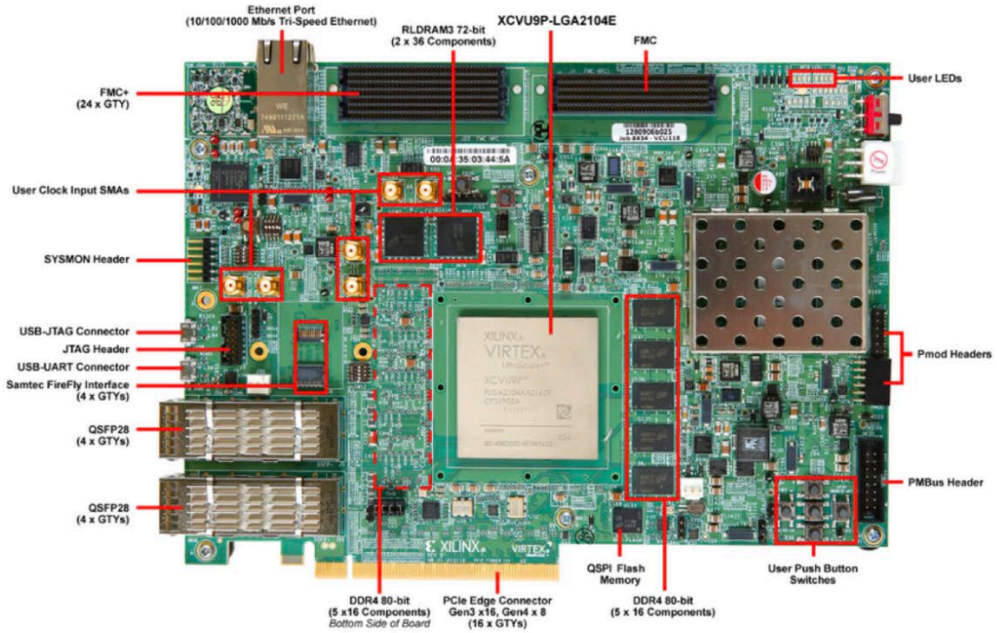

Xilinx Virtex UltraScale+ FPGA VCU118 Evaluation Kit

Featured Xilinx Devices

Featuring the **VCU118 XCVU9P-L2FLGA2104E FPGA**

System Logic Cells (K)	2,586
DSP Slices	6,840
Memory (Mb)	345.9
GTY 32.75 Gb/s Transceivers	120
I/O	832

Communication & Networking

- 10/100/1000 Mbps Ethernet (SGMII)
- Dual 4x28Gbps QSFP28 cages
- Samtec FireFly 4x28Gbps Interface
- Dual USB-to-UART Bridge with mico-B USB connector
- RJ45 Ethernet connector
- PCI Express endpoint Gen3 x 16

Clocking

- SI5335A Quad Clock Generator
- Si570 IIC Programmable LVDS Clock Generator
- SI5328C Clock Multiplier and Jitter Attenuator
- 2x SMA MGT Reference Clock inputs
- 1 SMA User Clock input

Memory

- Two 4 GB DDR4 component memory interfaces (five [256 Mb x 16] devices each)
- 4 MB RLD3 component memory interfaces (five [256 Mb x 16] devices each)IIC EEPROM: 8Kb
- Micro Secure Digital (SD) connector 1Gb Quad SPI Flash

Expansion Connectors

- FMC+ HSPC connector (24 – 28Gbps GTY Transceivers, 80 differential user defined pairs)
- FMC HPC1 connector (58 differential user defined pairs)
- PMOD header
- IIC

L0TP+ new features: ML on FPGA for trigger computing

TensorFlow → Keras → QKeras → hls4ml → Vivado™ HLS

- **Exploration of Vivado HLS-based method to derive ML FPGA embedded hardware for trigger computing on data stream**
- **Goal**: for any event detected by the RICH provide an estimate for the **number charged particles** and the **number of electrons ➔ (e.g.) Dense NN model**

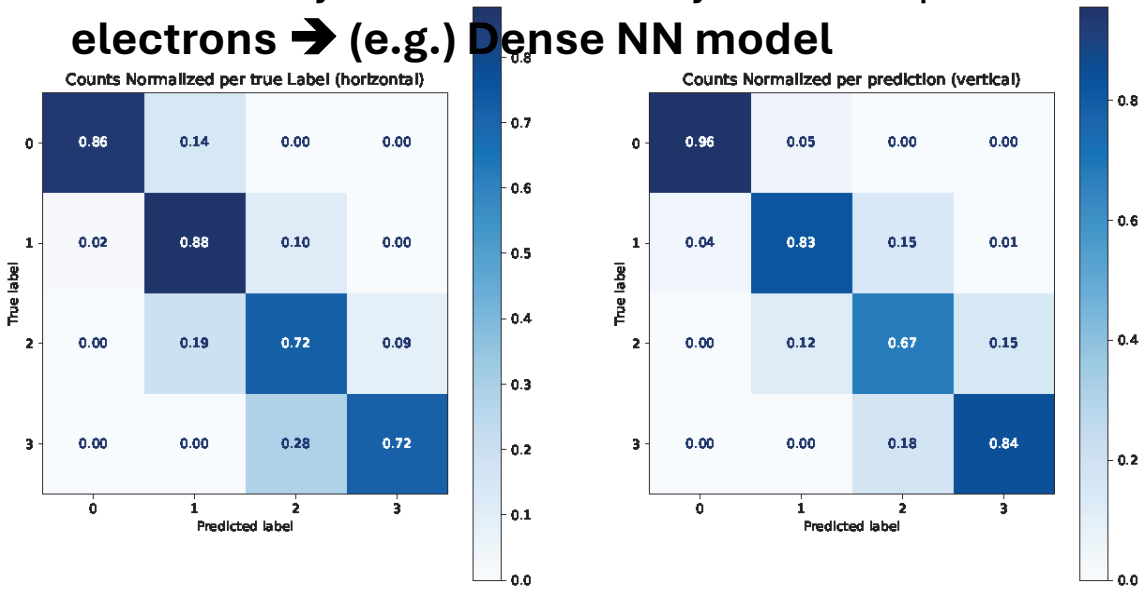

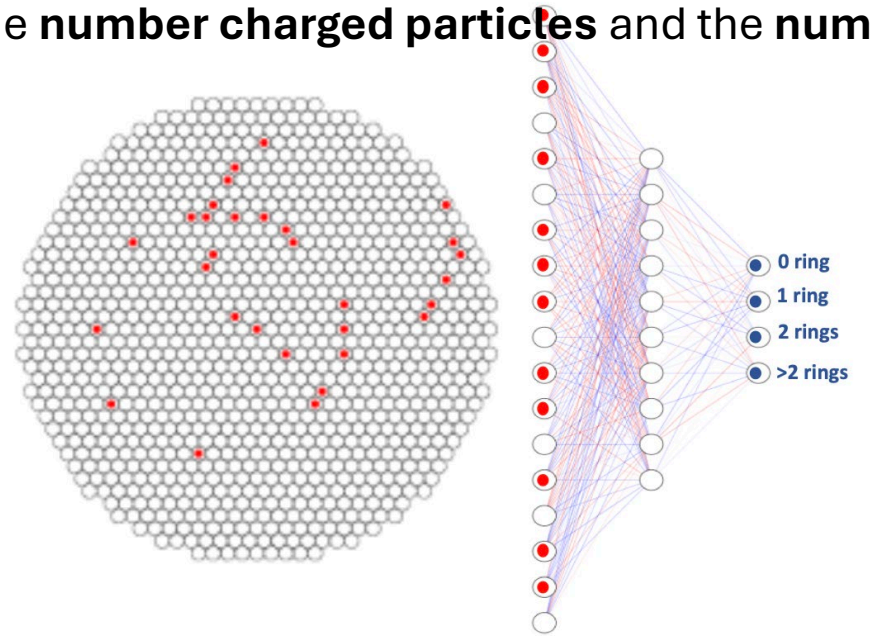

L0TP+ new features: ML on FPGA for trigger computing

- Preliminary results for online classification of the **number of "electrons"** show that even the very simple NN architectures that we tested is able, below 35 GeV/c momentum, to achieve interesting performance (see terminal picture below).
- It can be improved for the online unfiltered event stream using a **dedicated NN receiving in input data from other detectors** (e.g. L0CALO).

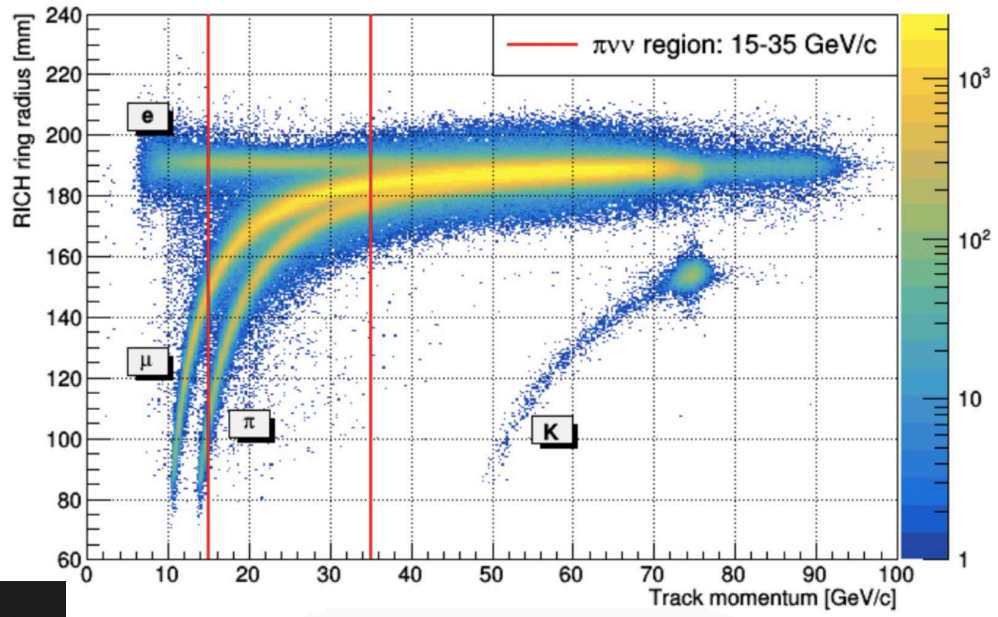


```
Total Events   163905
Total events of class 0 is    84628   (51.63 %)
Total events of class 1 is    76822   (46.87 %)
Total events of class 2 is     2432   (1.48 %)
Total events of class 3 is       23   (0.01 %)
Total events classified as 0 is    75533   (46.08 %)
Total events classified as 1 is    75209   (45.89 %)
Total events classified as 2 is    11920   (7.27 %)
Total events classified as 3 is     1243   (0.76 %)
Class  0  Efficiency 82.6  Purity 92.5  OverContamination  7.5  UnderContamination   0.0
Class  1  Efficiency 80.6  Purity 82.3  OverContamination  0.2  UnderContamination  17.5
Class  2  Efficiency 74.6  Purity 15.2  OverContamination  0.0  UnderContamination  84.8
Class  3  Efficiency 91.3  Purity  1.7  OverContamination  0.0  UnderContamination  98.3
```

L0TP+ validation steps

2022

- Parasitic mode test
- Beam test online in Novembe

2023

- Dry test in January/February
  (Integration completed)

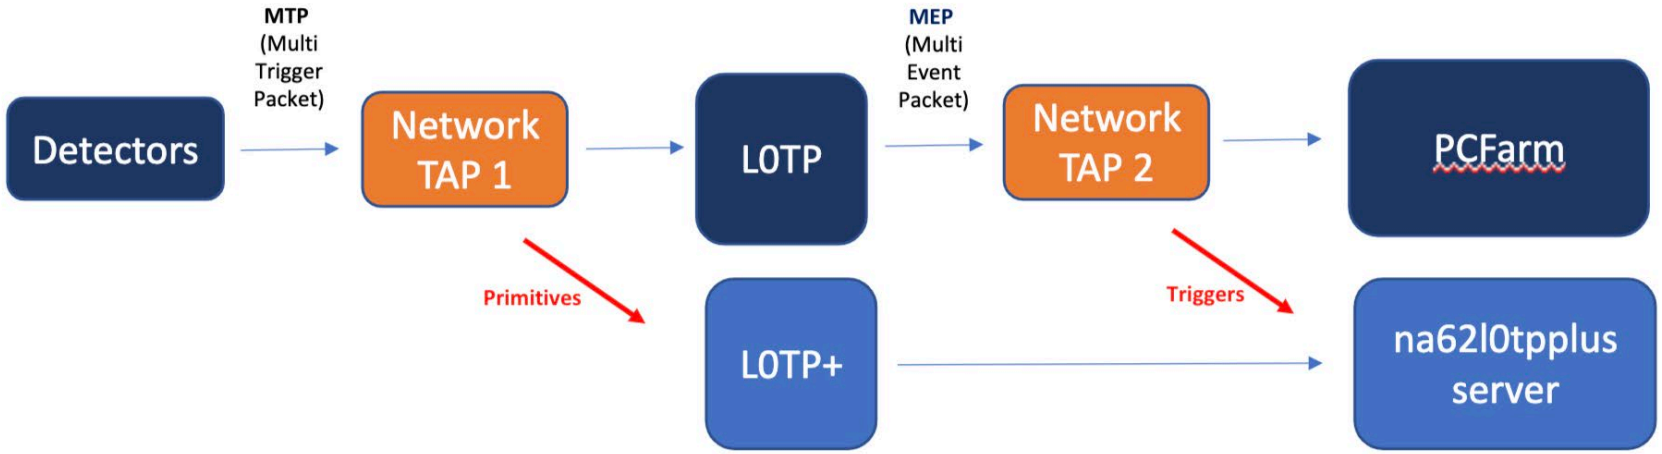

L0TP+ validation steps

2022

- Parasitic mode test
- Beam test online in November

2023

- Dry test in January/February
  (Integration completed)

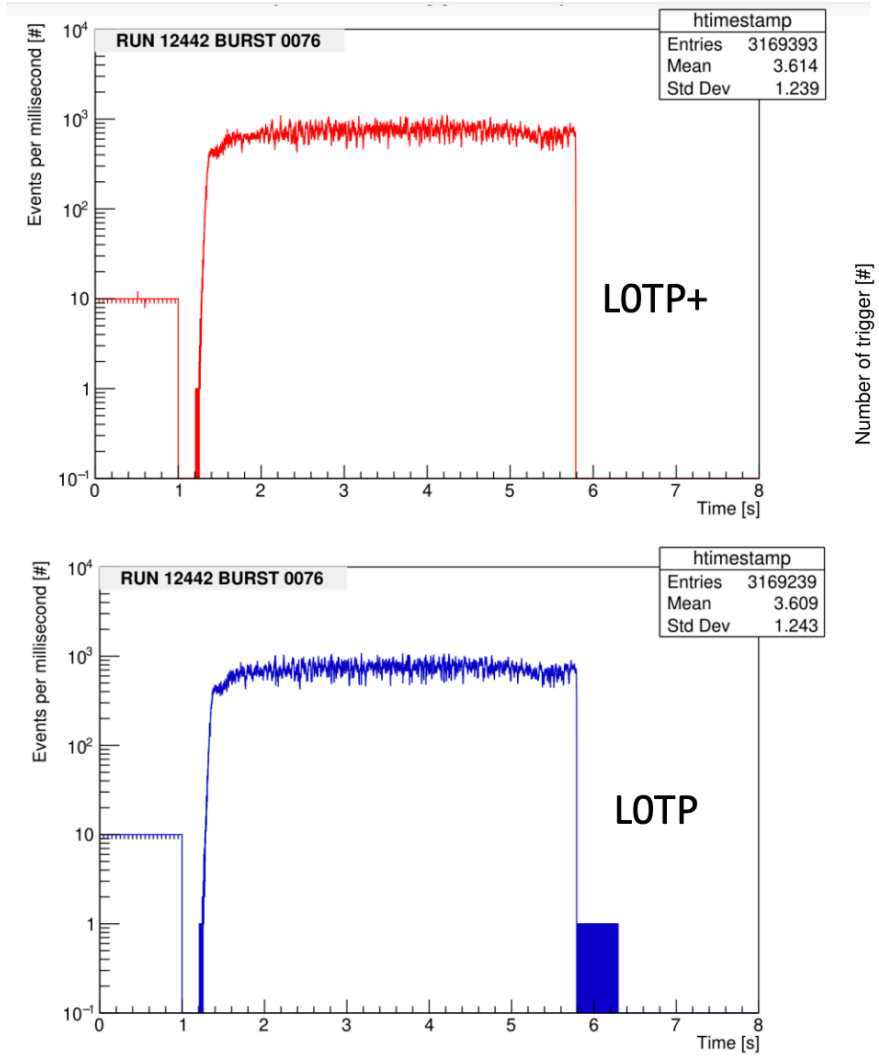

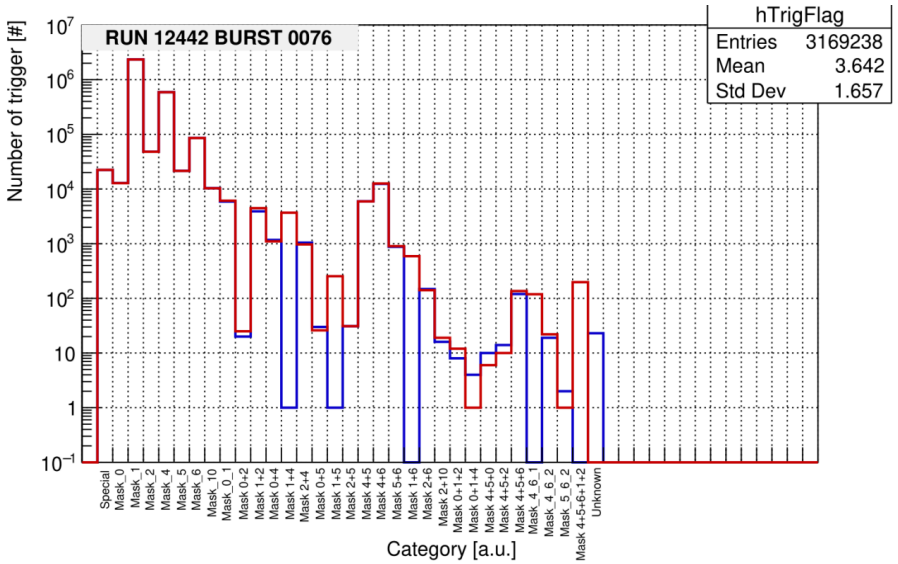

L0TP+ validation steps

2022
- Parasitic mode test
- Beam test online in November

2023
- Dry test in January/February
  (Integration completed)

- Last tests were performed in no beam (**dry run**) condition, swapping L0TP+ online and disconnecting L0TP

- System output was aligned with the standard NA62 data flow (from L0TP+ to PCFarms, fixed old issue on IP and MAC addresses) ➔ L0TP+ supports load balancing to PC Farm directly in hardware while L0TP uses a software program on server)

L0TP+ validation steps: Hardware

2022

- Parasitic mode test
- Beam test online in November

2023

- Dry test in January/February
  (Integration completed)

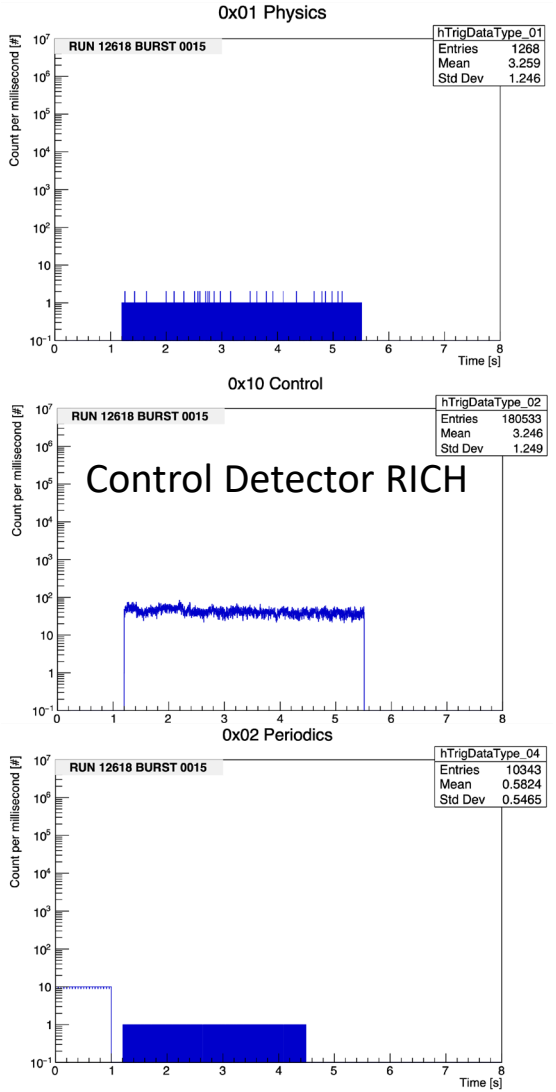

Few TDAQ boards on (RICH,MUV3,NCHOD) were used to send primitives to L0TP+.
The HV was switched off, so data acquired came from **electronic noise** due to low thresholds on TEL62

Conclusions

- Validation moved forward ➔ Fake EOB issue solved, CHOKE ON/OFF validated, Lkr Calib validated
- Fully integrated in NA62 data flow ➔ PC Farm Load balancing, RC command validated
  ➔ **L0TP+ IS NOW ONLINE** (since February last tests)

- **Next Step 1**: L0TP equivalent system at day zero (beam engineering run) ➔ in process during these days
- Continue testing until the start of Run ➔ increase confidence and test stability of L0TP+, debugging some new issues that may occurs before

- **Next Step 2**: keep parasitic setup as development platform during 2023
  ➔ ML PID on FPGA, takin

New Trigger algorithm

- My proposal is not to touch at all the pnn standard trigger masks
- mask0: RICH3*NewCHOD*!MUV3
- mask1: RICH3*UTMC*!QX*!MUV3*!LKr

- Introduce a **new logic signal at L0TP+** level:
  Qi!Mi **= QI1*!MQI1+ QI2*!MQI2+ QI3*!MQI3+ QI4*!MQI4**

- Add 2 new trigger **masks 8-9** build as follow (we have up to 16 slots):
- **Mask 8 = RICH3*!Qi!Mi**
- **Mask 9 = RICH3*UTMC*!QX*!(Qi!Mi)*!LKr**

- **Mask 8:** even with very high downscale factor ~400 same as mask0 will allow to **extract immediately the trigger efficiency increase** by comparing number of **K2π in mask0 and mask8**

- **Mask 9:** initial DS depending on the rate increase observed. After validation and optimization assign proper bandwidth is significant efficiency increase is observed

L0TP+ validation steps

2022

- Parasitic mode test
- Beam test online in November

2023

- Dry test in January/February
  (Integration completed)

Number of trigger per millisecond, superimposed L0TP (blue) and L0TP+ (red)

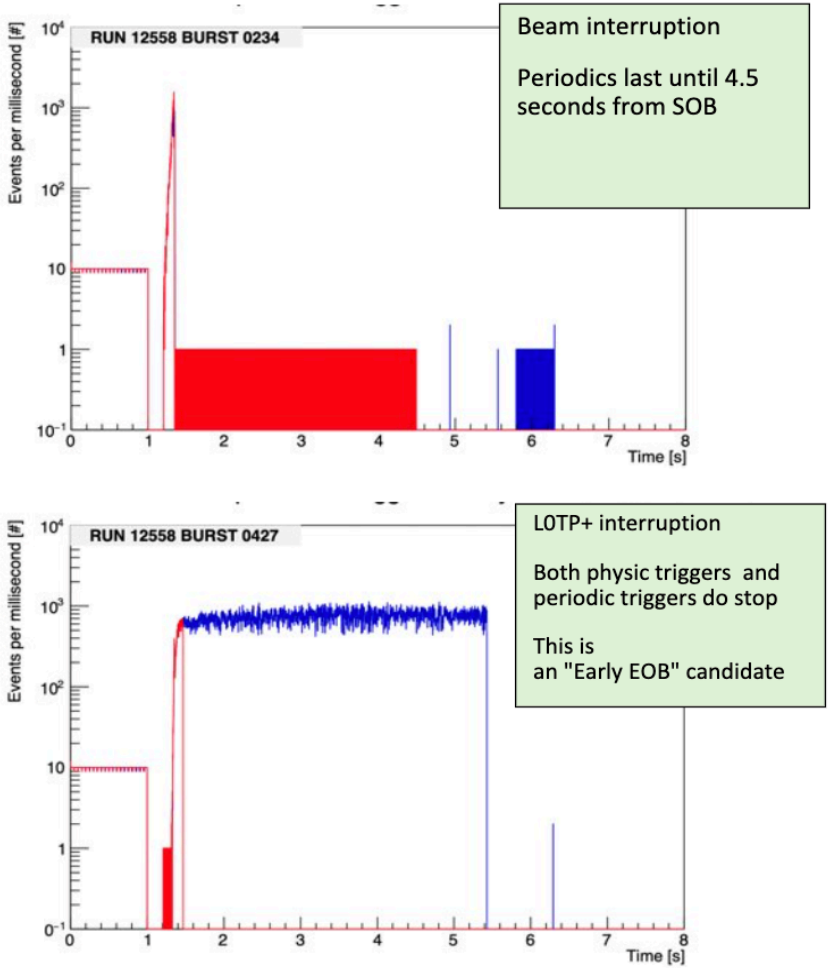

What was observed was compatible with **fake SOB** (11 induced) and with **fake EOB** (01 induced)
This also match with the presence of a modified ribbon cable to reduce **crosstalk** in L0TP old setup, differently from the flat ribbon cable of the first L0TP+ parasitc setup

L0TP+ validation steps: Hardware

2022

- Parasitic mode test
- Beam test online in Novembe

2023

- Dry test in January/February
  (Integration completed)

"SOB-EOB" issues: **Crosstalk Hypothesis on flat ribbon cables**

Confirmation strategy

Fake EOB counter logic on L0TP+ to quantify their relative frequency and to reproduce the Crosstalk issue

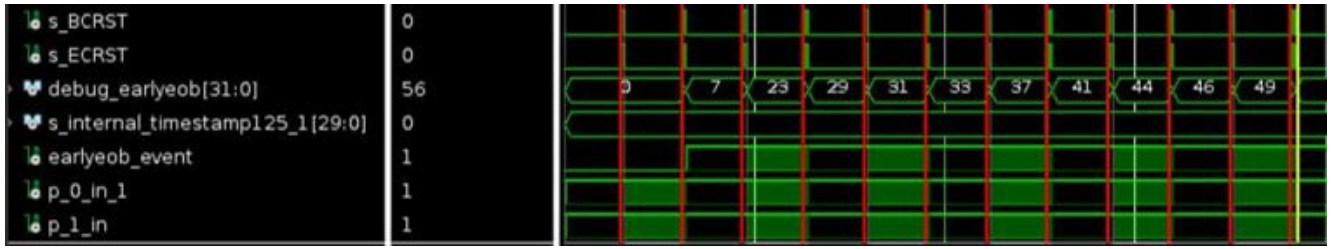

Solution

Switching to the custom L0TP Cable, Crosstalk issue seemed to be solved from the counter logic side.

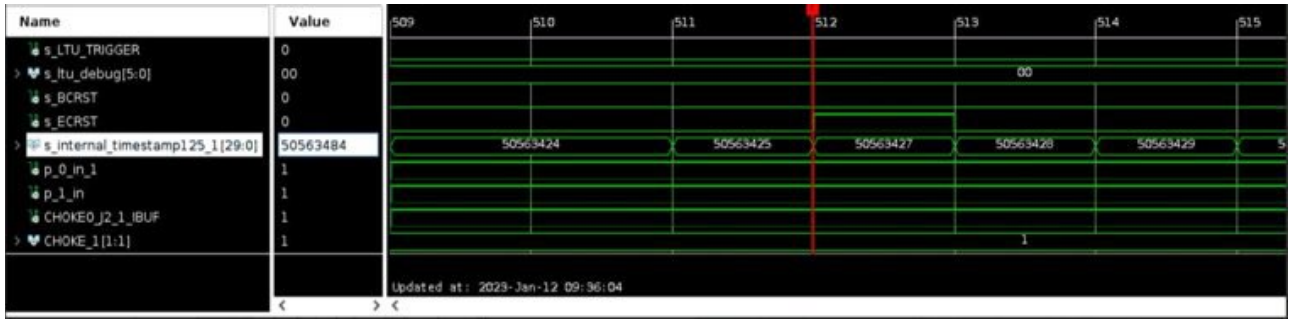

L0TP+ validation steps: Software

2022
- Parasitic mode test
- Beam test online in November

2023
- Dry test in January/February
  (Integration completed)

Big Screen display:
- Correct propagation of RC GUI changes through L0TP+ DIM Server (and last, its registers).
- Big Screen correctly reports all meaningful values for L0TP+ (as for L0TP)

PC Farm load balancing (DIM):
Since L0TP+ supports load balancing (in a round-robin fashion) to PC Farm directly in hardware, once validated the correct utilization of the total number of PC Farms by L0TP+, we verified [illegible] PC Farms and then reincluding them [illegible]

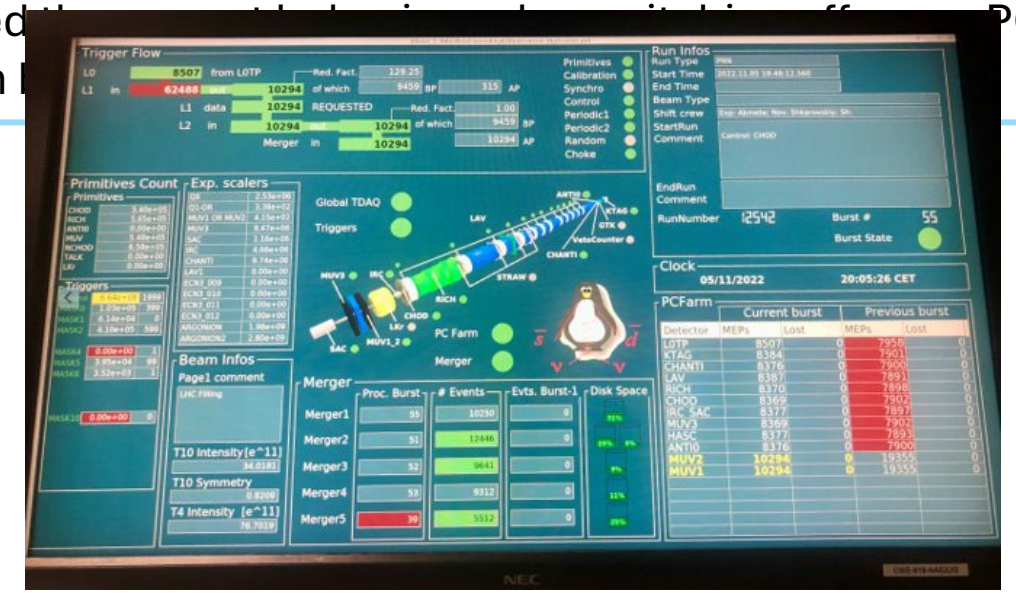

L0TP+ new features: GPU for trigger computing

Heterogenous Computing Node for HEP low level trigger

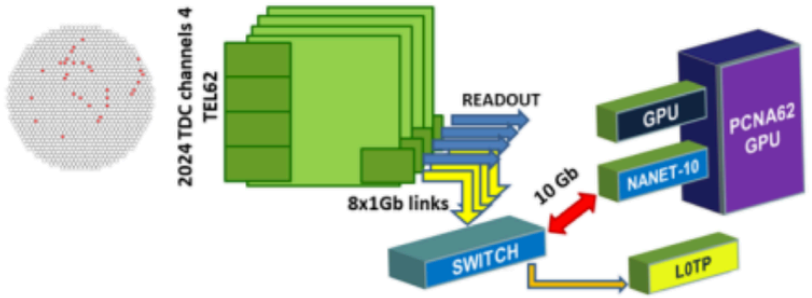

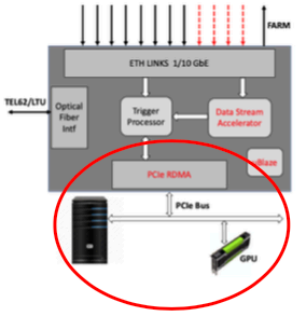

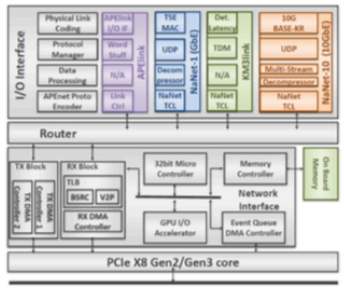

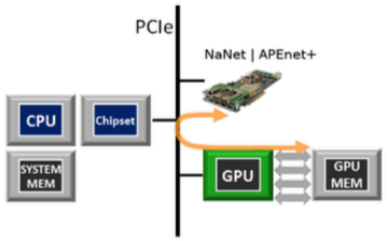

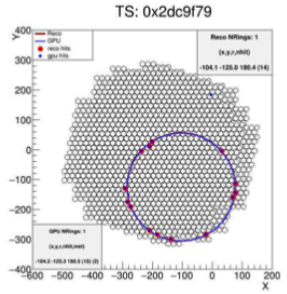

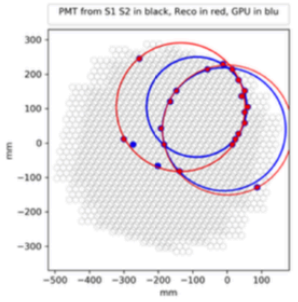

Leverage on NaNet project design

- Hardwired GPU Direct engine in FPGA implementing a low latency, low jitter data transfer between FPGA-based readout channels (Tel62) based on commodity protocol (UDP) and GPU
- GPU optimized application for identification of RICH Cherenkov ring
- Demonstrated in parasitic mode on NA62 experiment

A single device for L0TP GPU-accelerated to lower system complexity, to reduce trigger latencies and to enhance trigger performance. By product:

- L0TP+ complex diagnostic for hardware verification
- Backup primitives on trigger host server to get large statistics of trigger behavior and performance

L0TP+: porting of L0TP design

- The original design of L0TP processor has been ported on L0TP+ testbed
- TP coded in behavioral VHDL --> porting was "painless"
  - Synthesis ok, simulation in progress
- Integration with new ETH-UDP block (inherited from NaNet project) is in progress
- Achievements:
  - Clock speed > 130MHz (vs 125MHz old L0TP clock)
  - Very low resource occupancy --> room for hardware implementing additional features

Name	CLB LUTs (1182240)	Block RAM Tile (2160)	DSPs (6840)	Bonded IOB (832)
top_L0tp	2.15%	8.13%	0.15%	4.57%
CDC_inst2 (NIMInterface)	<0.01%	0.00%	0.00%	0.00%
clocks_resets_inst (axi_eth_1G_clocks_resets)	<0.01%	0.00%	0.00%	0.00%
CTSTMP (altcountertimestamp)	0.00%	0.02%	0.00%	0.00%
dbg_hub (dbg_hub)	0.05%	0.00%	0.00%	0.00%
ethlink_inst (ethlink)	1.99%	7.38%	0.15%	0.00%
jtag_inst (jtag_axi_master)	0.04%	0.12%	0.00%	0.00%
pll_40Mhz_inst (pll_40Mhz_diff_in)	0.00%	0.00%	0.00%	0.00%
Reg_inst (Register_intf_v1_0_S00_AXI)	0.01%	0.00%	0.00%	0.00%
syn_top.virt_IO (vio_0)	0.04%	0.00%	0.00%	0.00%
trigger_inst (Trigger)	0.03%	0.60%	0.00%	0.00%