Istituto Nazionale di Fisica Nucleare

Ministero degli Affari Esteri
e della Cooperazione Internazionale

Jefferson Lab

Streaming Readout for next generation e- scattering experiments

Supported by Italian Ministry of Foreign Affairs (MAECI) as Projects of great Relevance within Italy/US Scientific and Technological Cooperation under grant n. MAE0065689 - PGR00799

Mariangela Bondì for JLAB-INFN SRO team
INFN - Sezione di Catania

SRO vs Traditional DAQ

TRADITIONAL DAQ

- Data path different from trigger path
- Trigger decision based on a limited information
- Trigger logic takes time to decide and if trigger decision is satisfied
  - trigger signal back to the FEE
  - a new event is defined
  - data read from memory and stored on tape

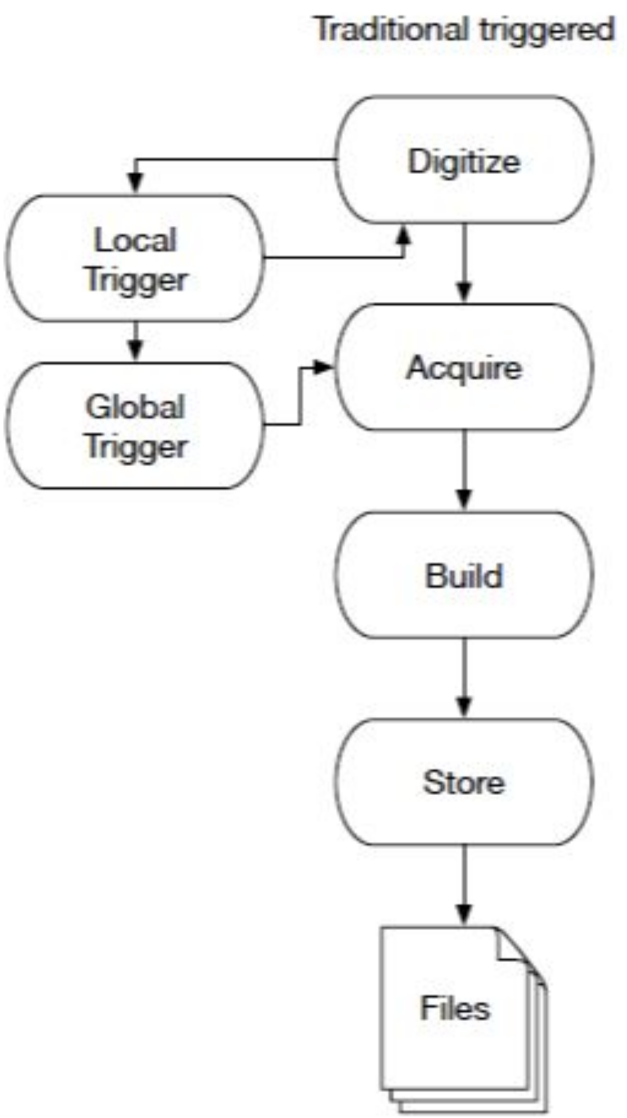

Pro

- we know it works reliably !

Drawbacks

- only few information form the trigger
- Trigger logic not easy to change and adapt to different conditions

SRO vs Traditional DAQ

SRO DAQ

- Data path equal to trigger path
- Trigger decision based on complete detector information
- “Event” defined at software level

Pro

- All channels can be part of the trigger
- Sophisticated tagging/filtering algorithms
  - Use of high level programming languages
  - Use of AL/MI tools
- Scalability

Drawbacks

- we don’t have the same experience as for triggered DAQ

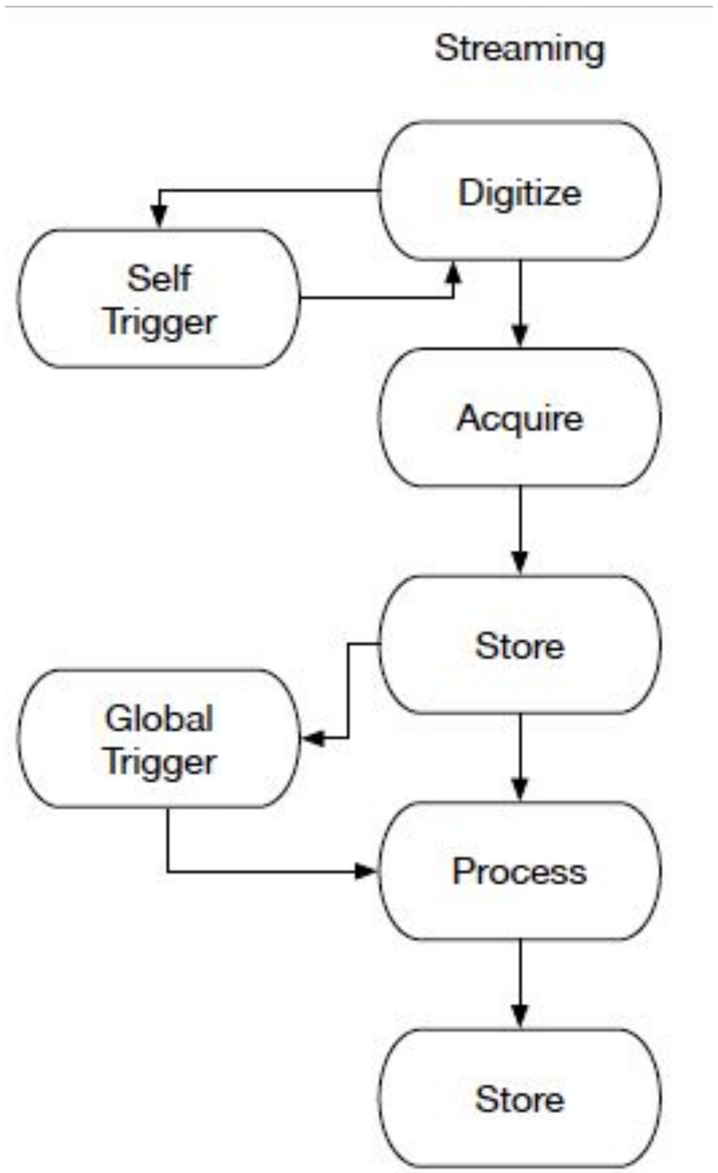

SRO for EIC

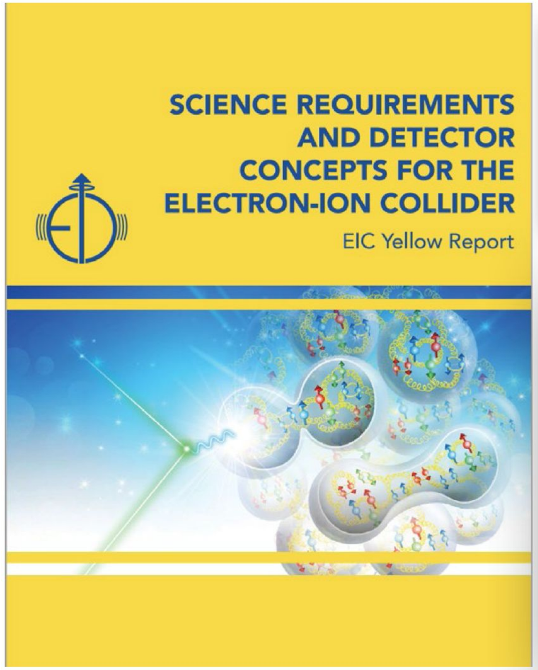

Streaming Readout for EIC Detectors

Proposal submitted 25 May, 2018

STREAMING READOUT CONSORTIUM

S. Ali, V. Berdnikov, T. Horn, I. Pegg, R. Trotta
Catholic University of America, Washington DC, USA
M. Battaglieri (Co-PI)[1], A. Celentano
INFN, Genova, Italy
J.C. Bernauer* (Co-PI)[2], D.K. Hasell, R. Milner
Massachusetts Institute of Technology, Cambridge, MA
C. Cuevas, M. Diefenthaler, R. Ent, G. Heyes, B. Raydo, R. Yoshida
Thomas Jefferson National Accelerator Facility, Newport News, VA

* Also Stony Brook University, Stony Brook, NY

ABSTRACT

Micro-electronics and computing technologies have made order-of-magnitude advances in the last decades. Many existing NP and HEP experiments are taking advantage of these developments by upgrading their existing triggered data acquisitions to a streaming readout model. A detector for the future Electron-Ion Collider will be one of the few major collider detectors to be built from scratch in the 21st century. A truly modern EIC detector, designed from ground-up for streaming readout, promises to further improve the efficiency and speed of the scientific work-flow and enable measurements not possible with traditional schemes. Streaming readout, however, can impose limitations on the characteristics of the sensors and sub-detectors. Therefore, it is necessary to understand these implications before a serious design effort for EIC detectors can be made. We propose to begin to evaluate and quantify the parameters for a variety of streaming-readout implementations and their implications for sub-detectors by using on-going work on streaming-readout, as well as by constructing a few targeted prototypes particularly suited for the EIC environment.

14.6 Data Acquisition

14.6.1 Streaming-Capable Front-End Electronics, Data Aggregation, and Timing Distribution

A streaming readout is the likely readout paradigm for the EIC, as it allows easy scaling to the requirements of EIC, enables recording more physics more efficiently, and allows better online monitoring capabilities. The EIC detectors will likely be highly segmented,

EIC R&D

Streaming Readout Consortium

eRD23

ECCE

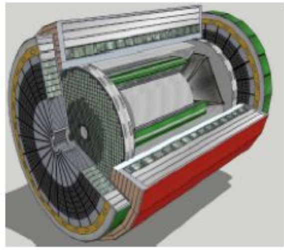

ATHENA

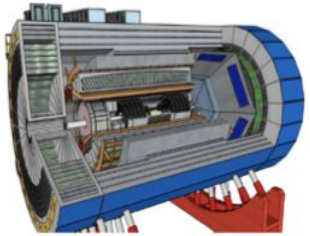

CORE

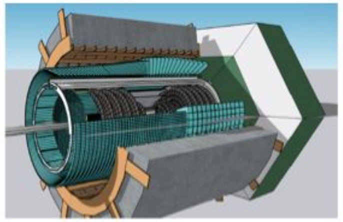

The three projects shared the same SRO concept → DETECTOR I

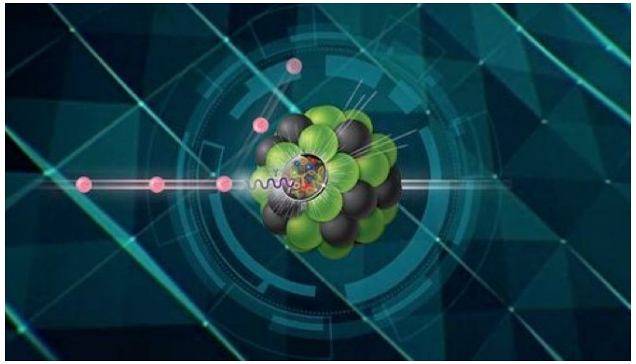

Streaming Readout DAQ @ JLAB

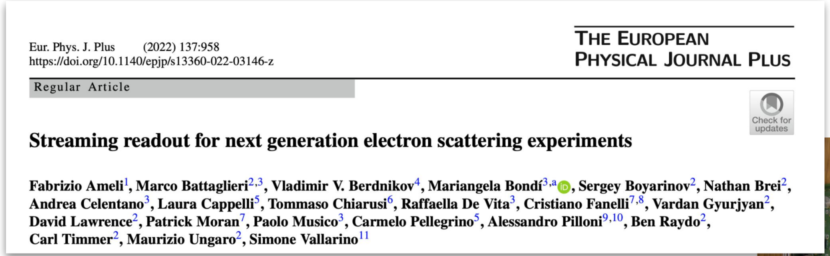
Eur. Phys. J. Plus (2022) 137:958
https://doi.org/10.1140/epjp/s13360-022-03146-z

THE EUROPEAN PHYSICAL JOURNAL PLUS

Regular Article

Check for updates

Streaming readout for next generation electron scattering experiments

Fabrizio Ameli[1], Marco Battaglieri[2,3], Vladimir V. Berdnikov[4], Mariangela Bondì[3,a], Sergey Boyarinov[2], Nathan Brei[2], Andrea Celentano[3], Laura Cappelli[5], Tommaso Chiarusi[6], Raffaella De Vita[3], Cristiano Fanelli[7,8], Vardan Gyurjyan[2], David Lawrence[2], Patrick Moran[7], Paolo Musico[3], Carmelo Pellegrino[5], Alessandro Pilloni[9,10], Ben Raydo[2], Carl Timmer[2], Maurizio Ungaro[2], Simone Vallarino[11]

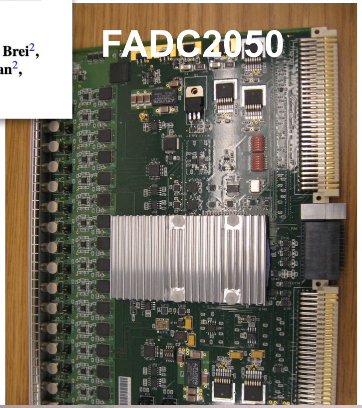

Jefferson Lab tests a next-generation data acquisition scheme

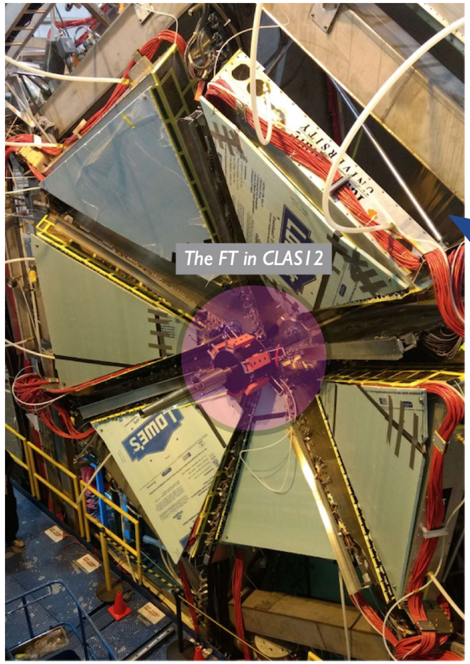

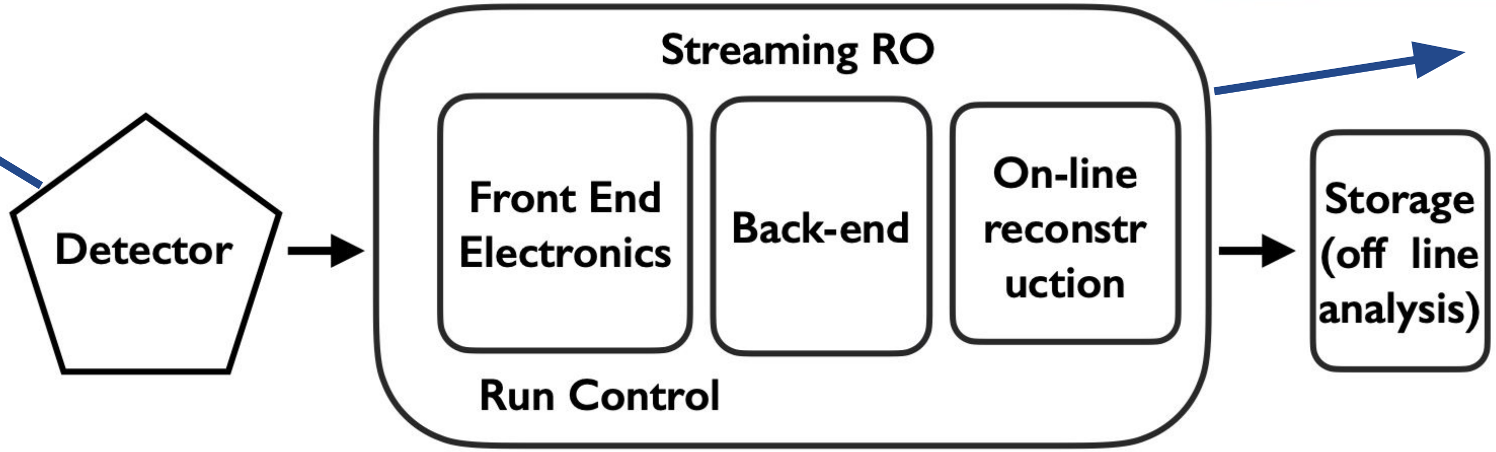

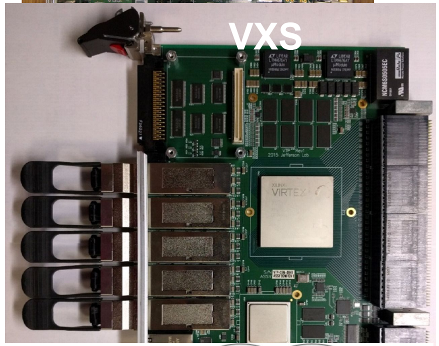

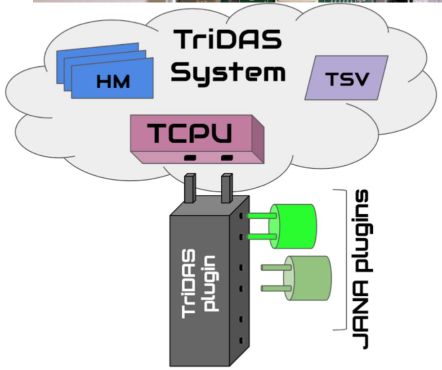

- SRO advantages are evident buy it needs to be demonstrated by the use in real experimental conditions
- To validate SRO concept:
  - Assemble SRO components
  - Test SRO DAQ
- **JLAB-INFN efforts to develop a prototype SRO DAQ**

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

SRO test @ JLAB: Experimental setup in HALL B

- **On-beam tests:**
    - 10.4 GeV electron beam on thin Pb/Al target
- **Hall-B CLAS12 Forward Tagger: Calorimeter + Hodoscope**
    - FT-CAL: 332 PbWO4 crystals (APD)
        - 10 +12 FADC250 boards + 2VTPs (in 2 crates/ROCs)
    - FT-HODO: 232 scintillator tiles (SiPM)
        - 15 FADC250 boards
    - FT-Tracker: MicroMegas
- **SRO DAQ full chain: JLAB-FADC250, TRIDAS, JANA2**

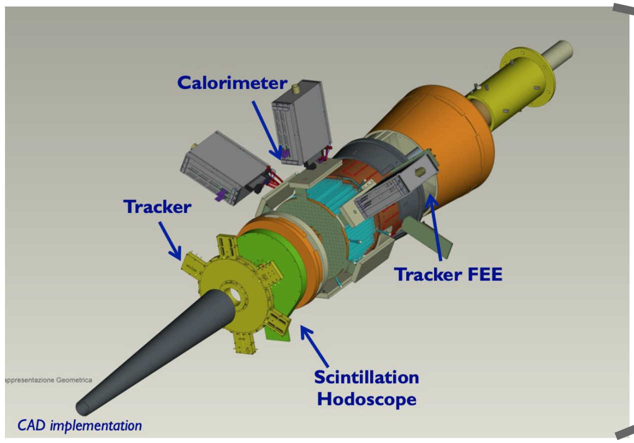

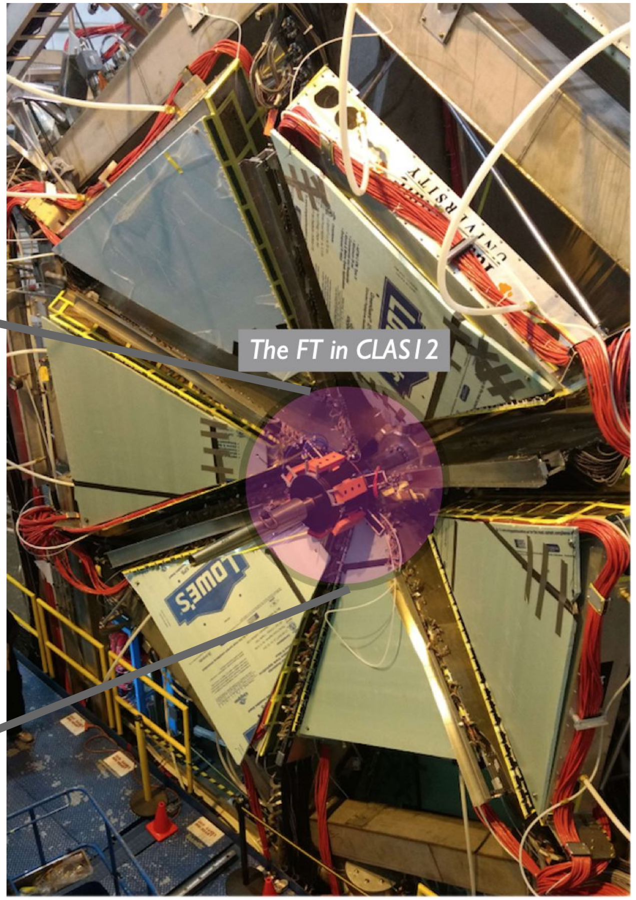

GOAL:

- **DAQ performance**
- **First AI application in streaming readout on real data tested online**
- **Physics channel identification: pi0 production**

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

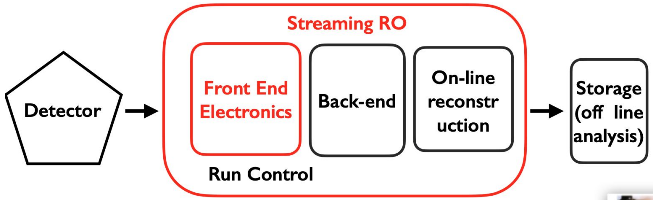

SRO test @ JLAB: Front-end electronics

D.Abbott, C.Cuevas, B.Raydo

JLAB FADC – Streaming mode

A 250 MHz FADC generates a 12 bit sample every 4ns. That's 3 Gb/s for one channel. 16 channels is 48 Gb/s. Currently, we identify a threshold crossing (hit) and integrate charge over a ROI and send only a sum and timestamp for each hit.

Available bandwidth will allow for 1 hit every 32ns from all channels.
A data frame (Time Slice) for all available hits is generated in the VTP every 65µs

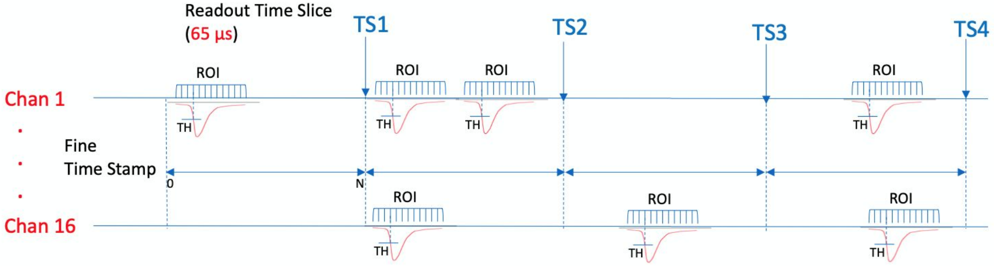

The next revision to the firmware will have an option for full ROI wave forms to be streamed, but this will allow possible dropped hits due to bandwidth limitations

The FADC can still simultaneously operate in triggered mode with an 8µs pipeline and 2µs readout window.

Jeffers

JLAB – VTP Board

Linux OS on the Zync-7030 SoC
(2-core ARM 7L , 1GB DDR3)
10/40Gbps Ethernet option
(runs the CODA ROC)

Xilinx Virtex 7 FPGA
Serial Lanes from both the VXS backplane and the Front panel
4GB DDR3 RAM

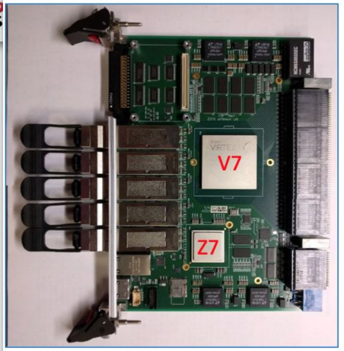

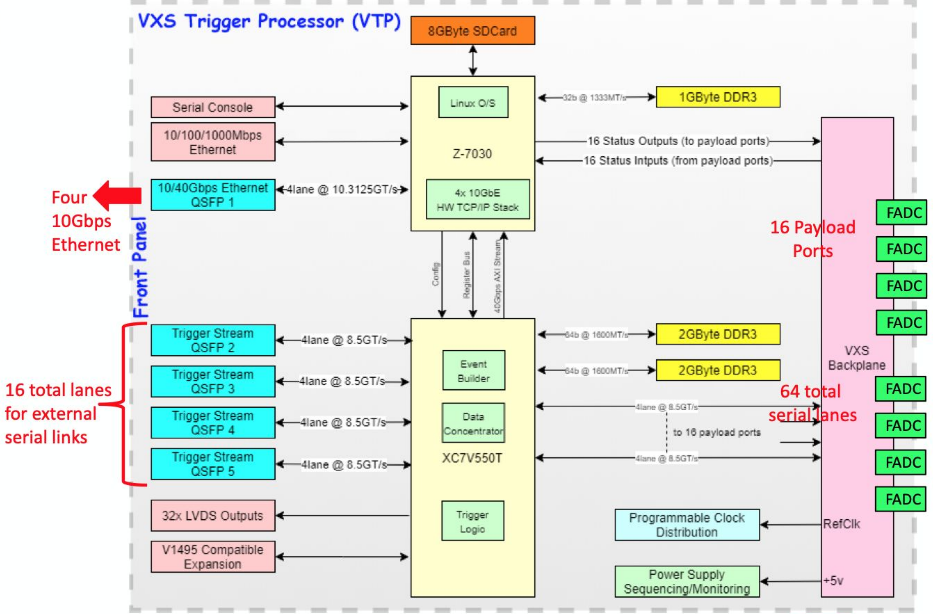

Jefferson Lab

SRO test @ JLAB: Back-end software

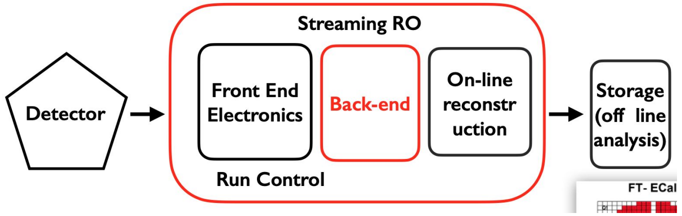

T. Chiarusi, C. Pellegrino, L. Cappelli

The TriDAS framework

Farnesina
Ministero degli Affari Esteri e della Cooperazione Internazionale
INFN CNAF

- TriDAS characteristics:
  - C++17 multithreaded software framework
  - Dependencies: CMake, ZeroMQ, Boost
  - State machine driven process
  - Flexible design:
    - Configurable via datacard (e.g. detector geometry)
    - L2 trigger algorithms in standalone plugins
    - Data format
- Composed by 5 modules:
  - HM (*Hit Manager*)
  - TCPU (*Trigger CPU*)
  - TSV (*TriDAS SuperVisor*)
  - EM (*Event Manager*)
  - TSC (*TriDAS System Controller*)
- The TriDAS code is available here

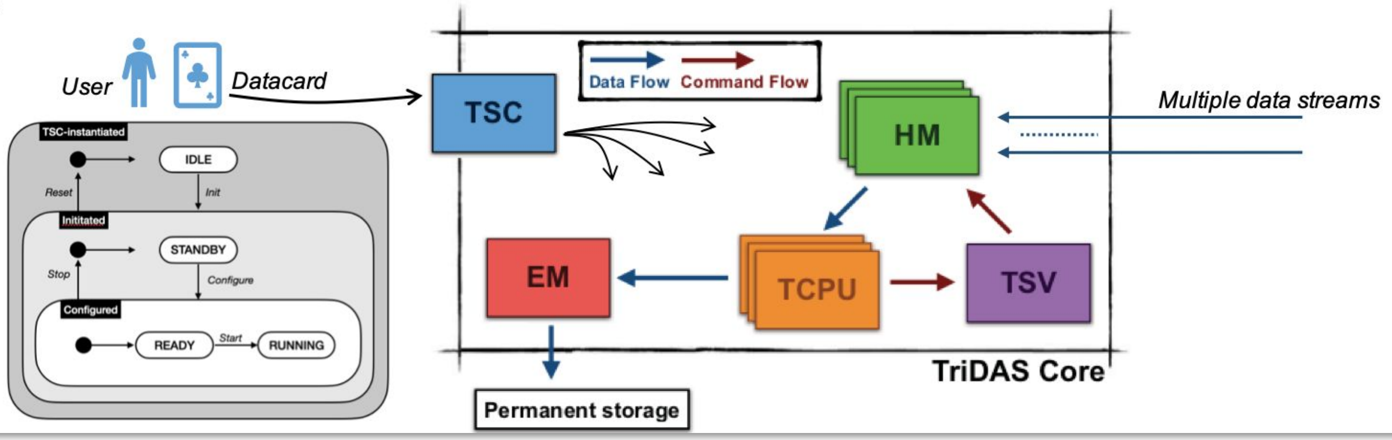

SRO test @ JLAB: Back-end software

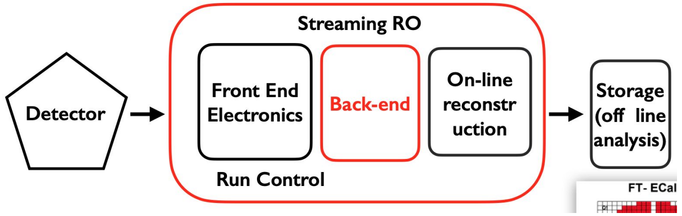

T. Chiarusi, C. Pellegrino, L. Cappelli

Data Flow: the HM aggregation

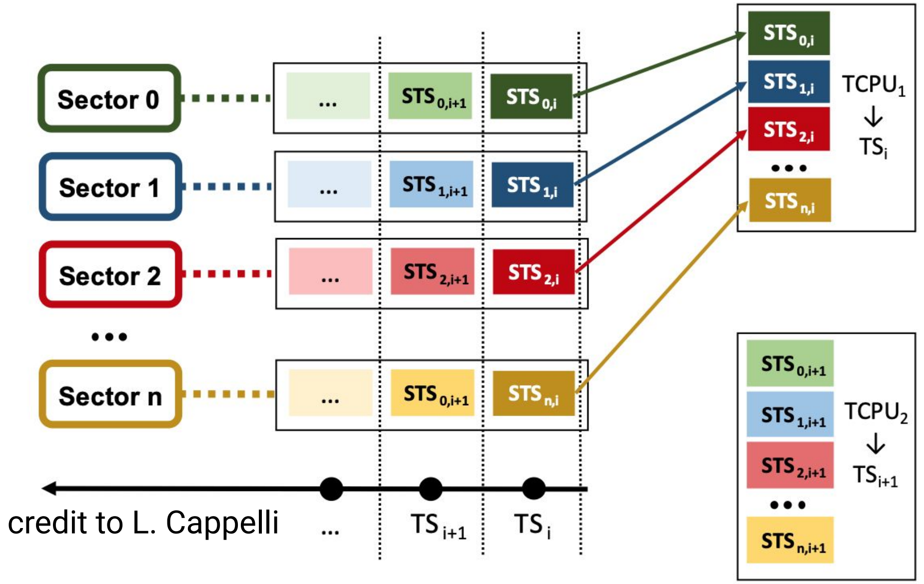

- Each HM:
  - Collects data from a specific sector of the detector
  - Subdivides the data into a sequence of time-ordered bunches called **Sector Time Slices (STSs)**
    - Fixed time duration called **Time Slice (TS)** chosen at run start time via datacard parameter (50ms in CLAS12)
  - Sends the STSs to a TCPU according to the token received from the TSV
- A TCPU receives all the STSs of a TS

SRO test @ JLAB: Back-end software

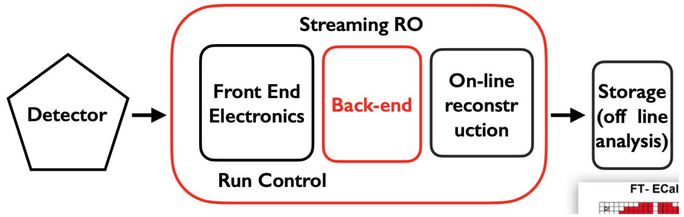

T. Chiarusi, C. Pellegrino, L. Cappelli

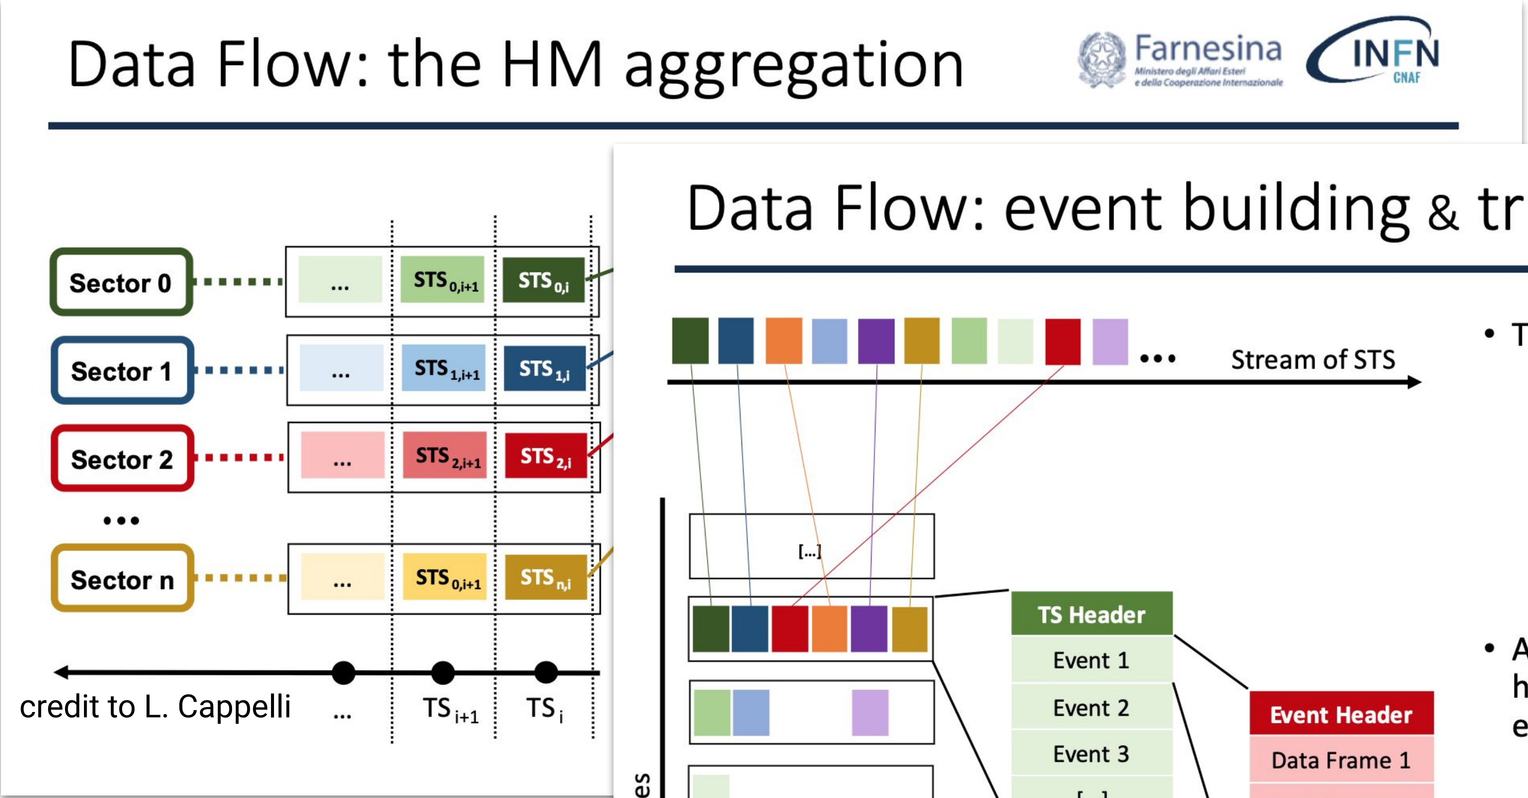

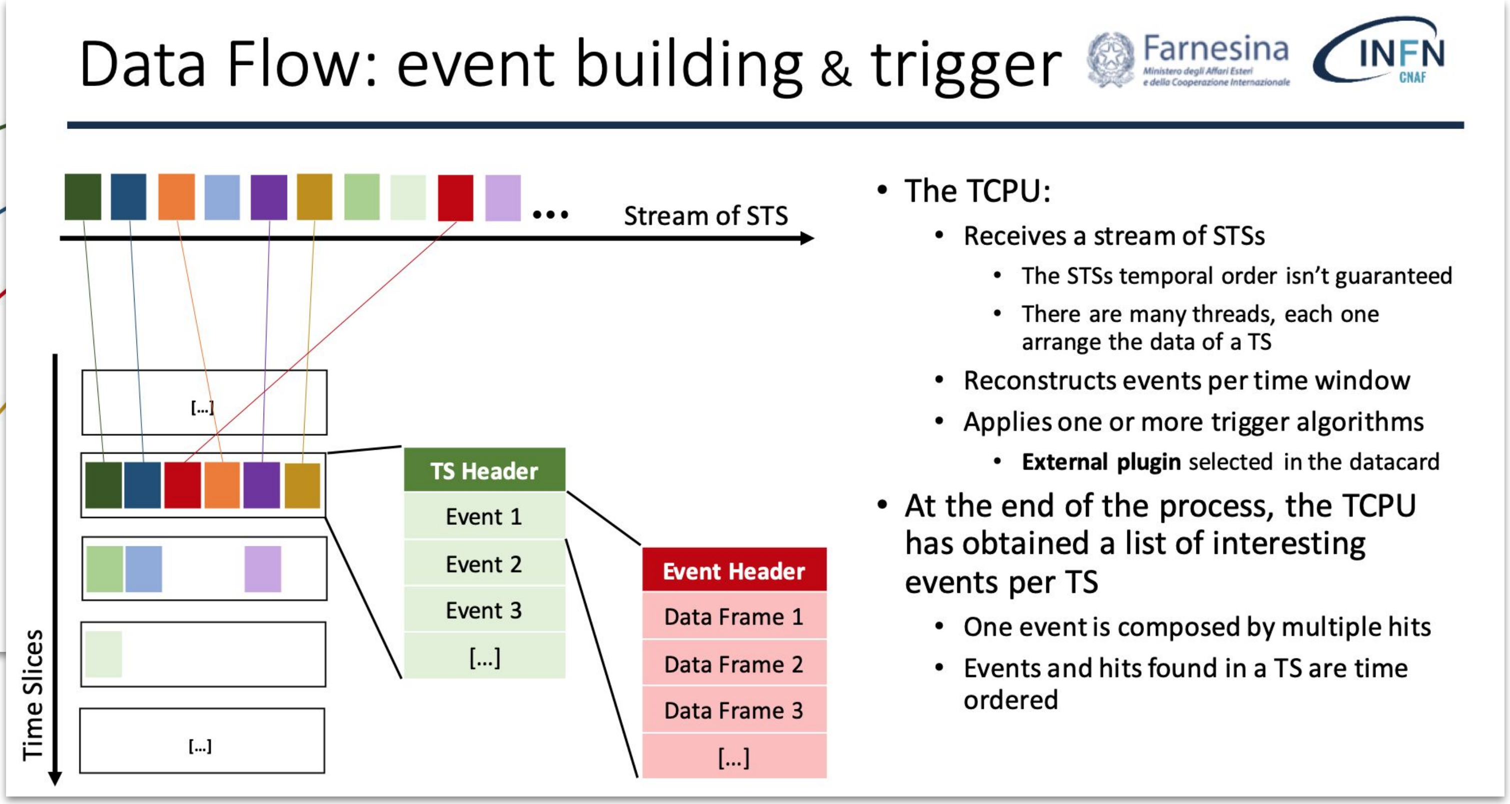

SRO test @ JLAB: Back-end software

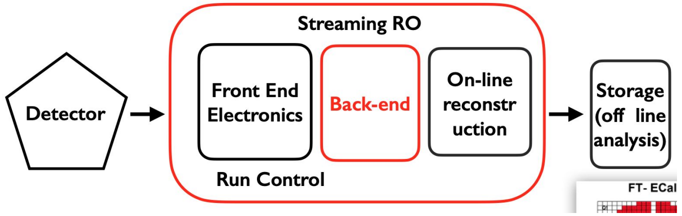

T. Chiarusi, C. Pellegrino, L. Cappelli

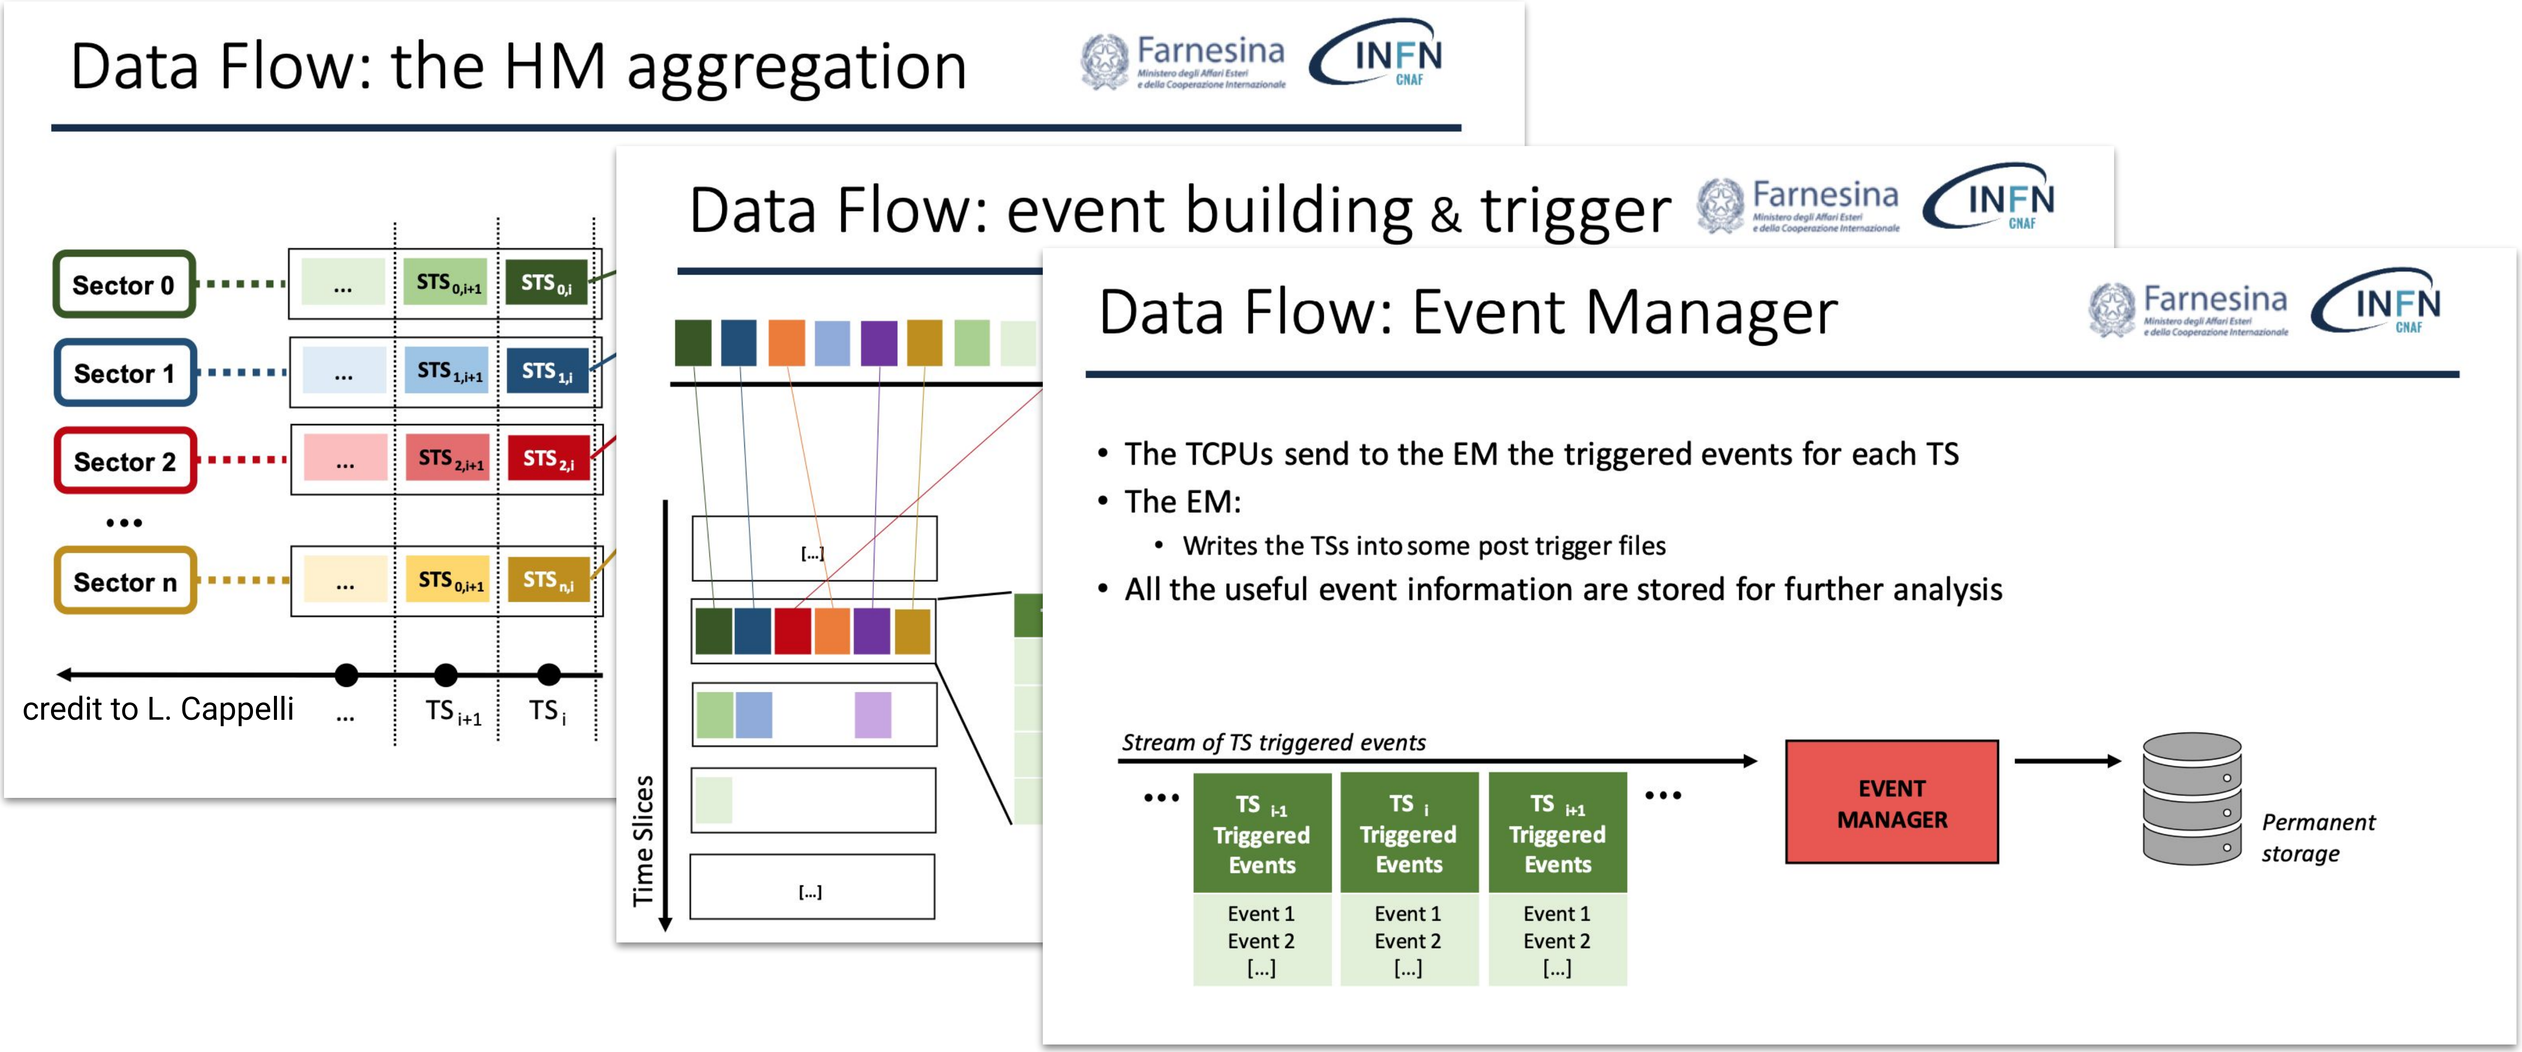

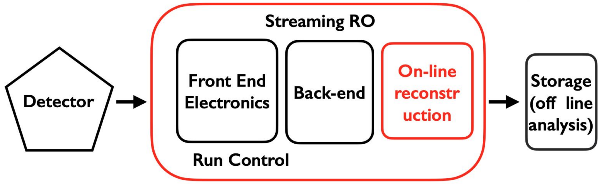

SRO test @ JLAB: Reconstruction framework

N.Brei, D.Lawrence

- Same reconstruction algorithms (software trigger) for both online and offline analysis
  - Cluster reconstruction with ML
  - Standard clustering algorithm
- Real-time tagging/filtering data
- Offline algorithm development immediately available for use in Software Trigger

Software triggers:

- TRIDAS L1 “minimum bias”: at least 1 crs with energy > 2 GeV
- JANA L2 triggers:
  - cosmic tracking
  - Number of cluster (>1, 2 ,3)
    - AI clustering reco algorithm
    - Standard clustering algorithm

TriDAS + JANA2

- JANA2: C++ framework
  - Full event reconstruction
    - Calibrations
    - Translation table
    - Multi-threading
  - Software trigger
    - Summed energy threshold
    - Single/Double cluster
    - Coincidence FT + FH
    - Prescale
  - Trigger decisions recorded in output stream

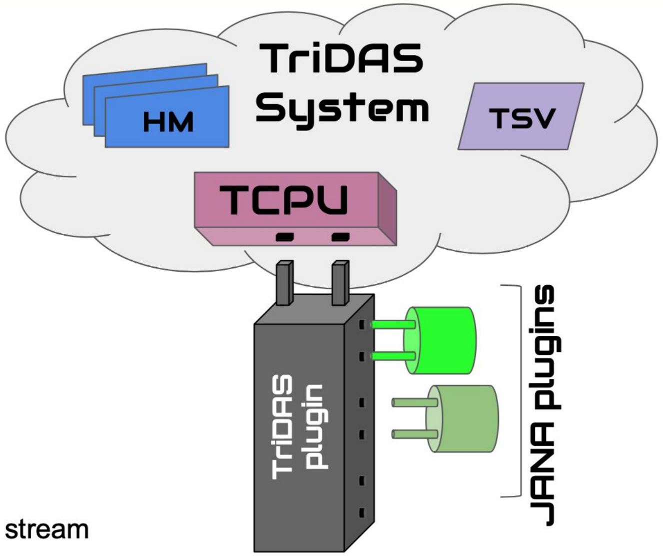

EPSCI https://jeffersonlab.github.io/JANA2/ Jefferson Lab

SRO test @ JLAB results: TRIDAS + JANA2 perfomance

T. Chiarusi, C. Pellegrino

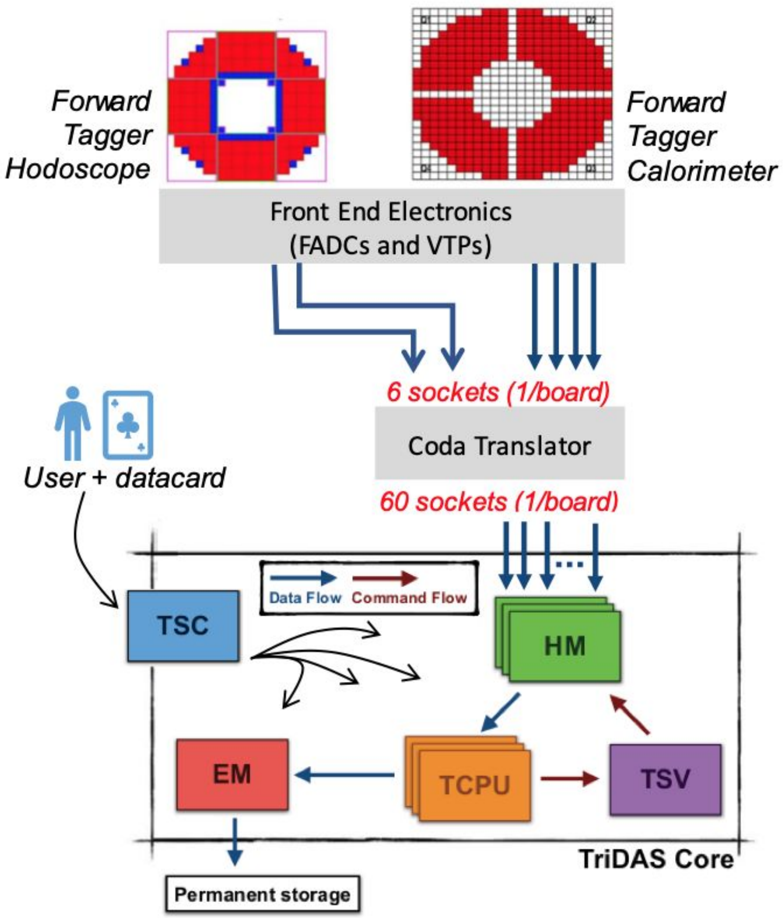

- Linux servers used:
  - 48 cores, 1GHz each, 64 GB RAM
  - 3 servers used for all modules
- HM instances: from 5 to 20
  - CPU consumption linear with the number of instances (500% – 1600%)
  - Memory occupancy constant (12-13 GB per run)
- TCPUs instances: 10 instances on 2 servers = 20 instances
  - 5 Time Slices at the same time on each instance
  - Trigger: **Jana2 plugin** (rudimental reconstructions and clustering)
  - CPU consumption: depending on the trigger algorithms (400% – 1600%)
  - Memory occupancy: 20-24 GB

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

SRO test @ JLAB results: off-line data analysis

- **On-beam tests:**
  - 10.4 GeV e- beam on thin Pb/Al target
  - Inclusive pi0 production
    - e + Pb/Al -> Xeπ^0 -> (X)e$\gamma\gamma$
  - Two gammas detected in FT-CAL

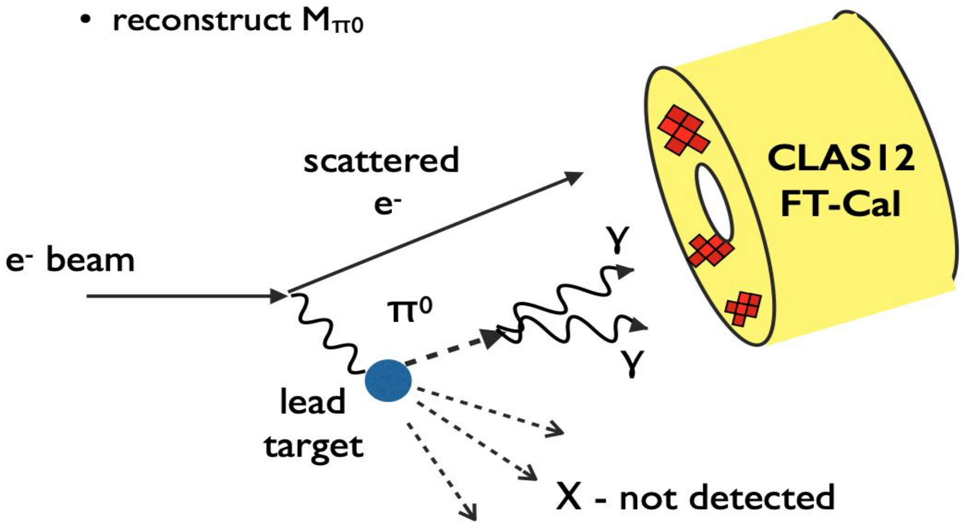

SRO test @ JLAB results: AI vs standard clustering

C. Fanelli

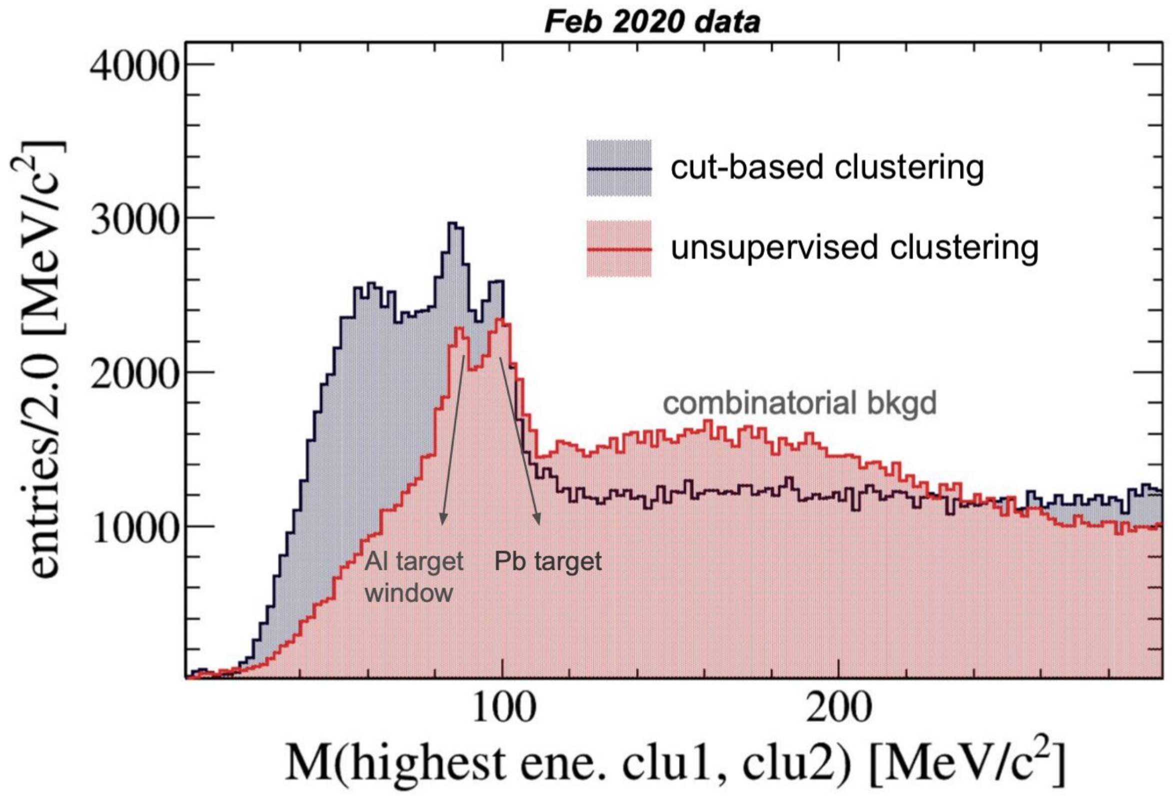

- AI clustering inspired by *Hierarchical Density-Based Spatial Clustering of Applications with Noise* (HDBSCAN)
  - It is not cut-based
  - it is able to cope with a large number of hits
- Compared γγ-invariant mass spectrum obtained utilizing both the standard and the HDBSCAN clustering algorithm
  - AI significantly improves signal-to-background ratio in the π0 region
  - A longer runtime of ~30% relative to the standard clustering algorithm
- AI clustering approach promising alternative to traditional cut-based approaches

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

SRO test @ JLAB results: physics channel

M. Battaglieri, M.Bondì, A. Celentano, R. De Vita, A. Pilloni, S. Vallarino

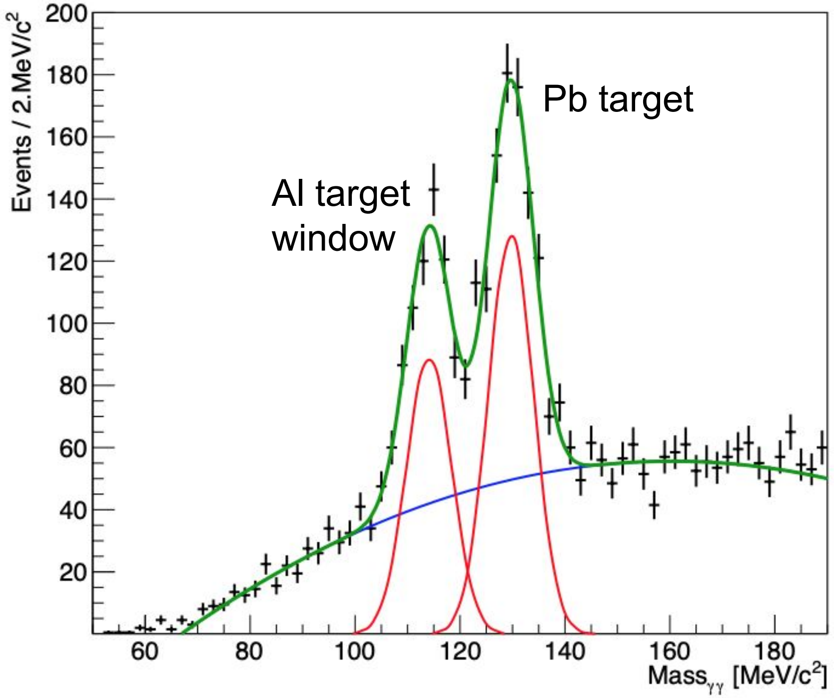

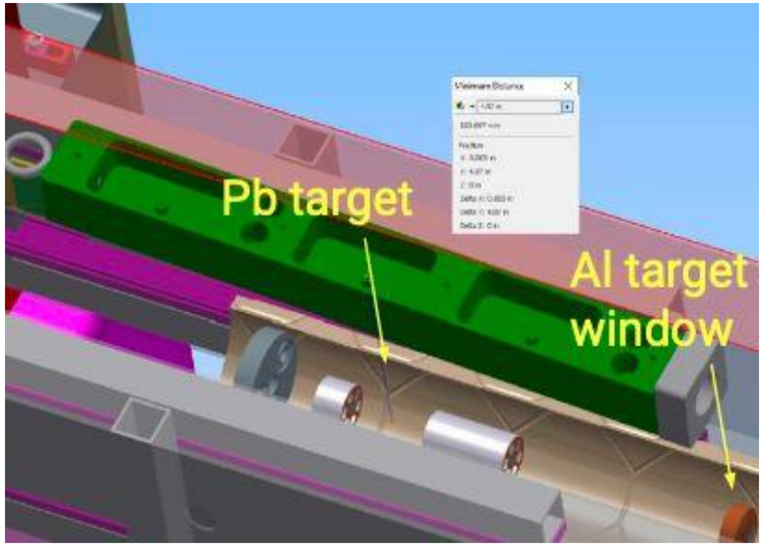

- Measured (expected) pi0 yield
  - Peak 1 : 1365 +/- 140 (~1800)
  - Peak 2: 930 +/- 100 (~420)
- Good agreement provides a significant validation of the SRO DAQ performance

Two pi0 peaks corresponding to two vertices (and a wrong assumption on the vertex position)

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

SRO @ JLAB: Hall D test

M. Battaglieri, V. Berdnikov, T. Horn

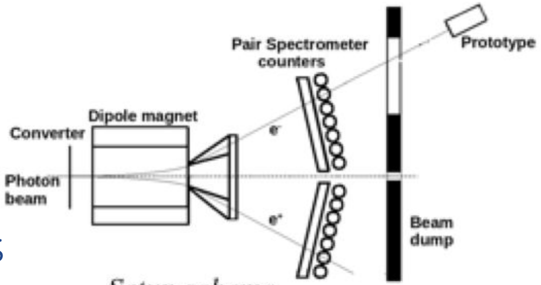

Setup scheme

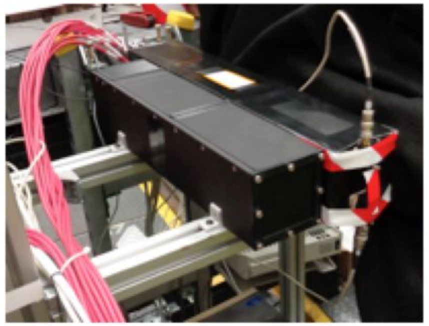

SiPM(left) & PMT(right) cal. prot.

Waveboard

- **EIC ECAL PbWO prototype**
  - Use the Hall D Pair Spectrometer setup
  - EIC-ECAL prototype irradiated with a 4.7 GeV e- beam
  - Simple setup to compare TRIGGERED to TRIGGER-LESS
  - 3x3 PbWO crystals, PMT and SiPM readout

- **SRO - DAQ setup:**
  - FEE: Waveboard
  - Back-end software: TRIDAS
    - L1 event: threshold equivalent to 2GeV
  - Reconstruction: JANA2

GOAL: compare TRIGGERED to TRIGGER-LESS

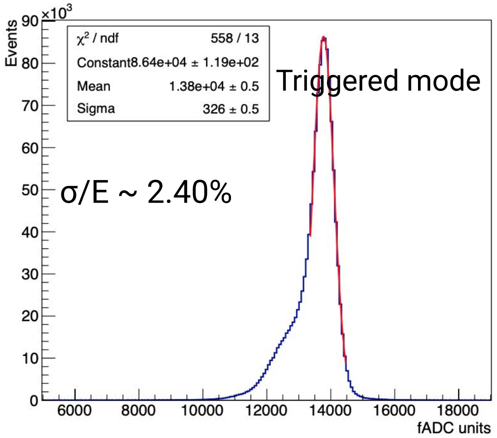

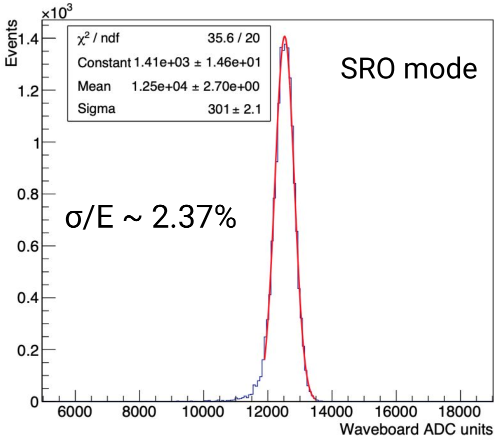

F. Ameli et al., Eur. Phys. J. Plus (2022) 137: 958
https://doi.org/10.1140/epjp/s13360-022-03146-z

SRO @ JLAB news: TRIDAS + ERSAP

ERSAP

- Reactive, event-driven data-stream processing framework that implements micro-services architecture
- Provides basic stream handling services (stream aggregators, stream splitters, etc.)
- Adopts design choices and lessons learned from TRIDAS, JANA, CODA and CLARA

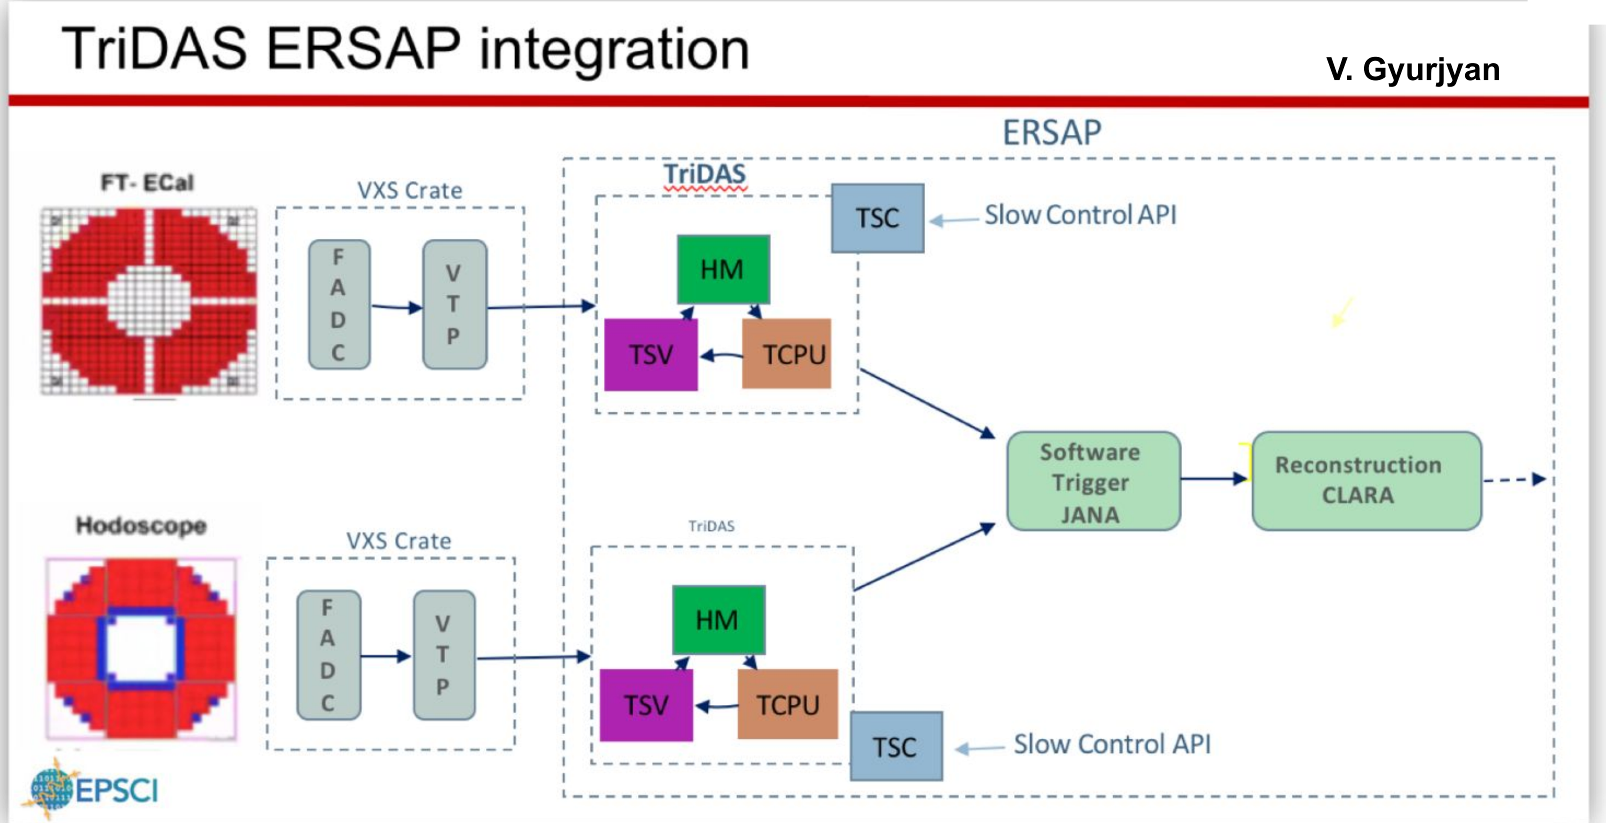

- **Next SRO - JLAB tests:**
  - cosmic tests with EIC-Cal prototype using ERSAP (November 2022)
  - on-beam tests with EIC-Cal prototype using ERSAP (Spring 2023)

Summary

- SRO is the option for future EIC
- Take advantage of the full detector's information for an optimal (smart) tagging/filtering
- So many advantages: performance, flexibility, scaling, upgrading …

  … but, has to demonstrate to be as effective (or more!) than triggered systems
- Streaming Readout on-beam tests performed in Hall-D and Hall-B at JLab
- First SRO chain (FE + SRO sw + ON-LINE REC) tested with existing hardware
- Deployment of JLab SRO framework based on micro-services architecture (ERSAP)
- Cosmic and on-beam (JLab) tests of EIC Cal prototype with AI-supported tagging/filtering (implemented in Jana2)

Many thanks to the whole JLAB SRO team: F.Ameli (INFN), M. Battaglieri (INFN), V.Berdnikov (CUA), S.Boyarinov (JLab) M.B. (INFN), N.Brei (JLab), A.Celentano (INFN), T.Chiarusi (INFN), C.Cuevas(JLab), R. De Vita (INFN), C.Fanelli (MIT), G.Heyes (JLab), T.Horn (CUA), V.Gyurjyan(JLab), D.Lawrence (JLab), L.Marsicano (INFN), P.Musico (INFN), C.Pellegrino (INFN), B.Raydo (JLab), M.Ungaro (JLab), S.Vallarino (INFN)

Many thanks to CLAS12 collaboration as well as JLAB technical staff for their accommodation and support of this effort