Streaming data acquisition system for future electron 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
Università di Roma Tor Vergata
INFN - Sezione di Roma2

Farnesina
Ministero degli Affari Esteri e della Cooperazione Internazionale

INFN
Istituto Nazionale di Fisica Nucleare

Jefferson Lab

Few comments

- **For many years the design of DAQ system has assumed that:**
  - The data rate from a large detector is impossible to acquire without a trigger to reduce event rates
  - Even if the untriggered data rate could be acquired it would be impossible to store.
  - Even if it could be stored the full dataset would represent a data volume that would require impractically large computing resources to process.

- **One of the biggest changes in the last 25 years is the ubiquitous presence of computing in 21st century life**.
  - Cost has gone down
  - Processor performance and bandwidth has consistently improved
  - Storage is cheaper and faster every year

- **It is time to design a new DAQ framework for the future.**

Traditional (Triggered) DAQ

- Data path is different from trigger path
    - All channels continuously measured and hits storec in short term memory by the FEE
    - Channels participating to the trigger send (partial) information to the trigger logic
- Trigger decision based on a limited information
- Trigger logic takes time to decide and if the trigger conditi is satisfied:
    - trigger signal back to the FEE
    - a new event is defined
    - data read from memory and stored on tape

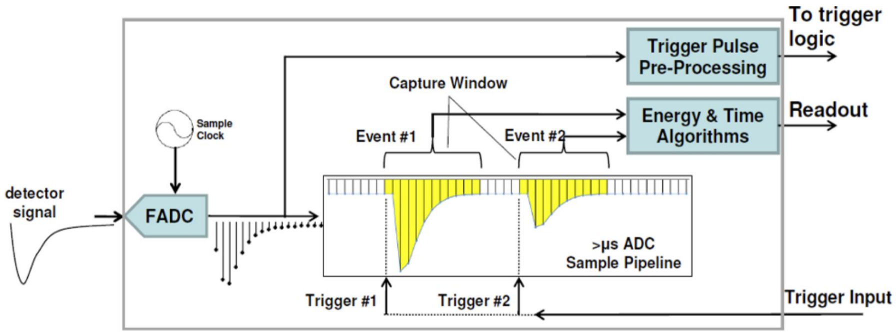

- **Pros**
    - we know it works reliably!
- **Drawbacks**:
    - only few information form the trigger
    - Trigger logic (FPGA) difficult to implement and debug
    - not easy to change and adapt to different conditions

Streaming readout: starting with an example

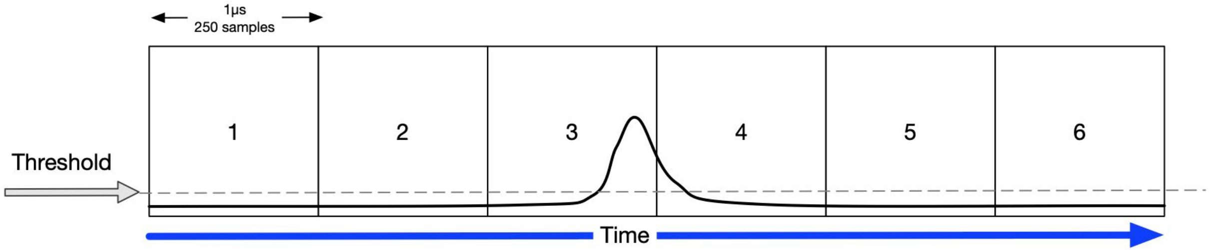

- In the example an ADC clocks a 32-bit sample every 4 nS into a memory.
- 250 samples represents 1µs of time - call this a “frame”
- The memory can be read as a sequence of six data frames each starting with a length and offset.
  - Frames 1, 2, 5, and 6 contain no data above threshold. So, length 0, offset 0.
  - Frame 3 has length 100, offset 150.
  - Frame 4 has length 50, offset 0.
  - **If length and offset are encoded into a single 32-bit word this simple format encodes the useful data from 6 µS in only 156 words.**
- Encoding data as a continuous sequence of data blocks, each representing a period of time is **streaming readout.**
  - Fixed rate of variable length blocks
  - Thresholds and zero suppression are applied locally.

Length = 0	Offset = 0	Count = 1
Length = 0	Offset = 0	Count = 2
Length = 100	Offset = 150	Count = 3
100 Data words		
Length = 50	Offset = 0	Count = 4
50 Data words		
Length = 0	Offset = 0	Count = 5
Length = 0	Offset = 0	Count = 6

Streaming DAQ

- Data path equal to trigger path
    - each channel that exceeds a threshold (as implemented on the front-end board) is labeled with a time-stamp and then transferred to an on-line CPU farm.
- Trigger decision based on complete detector information, possibly with the same reconstruction software used offline
- “Event” defined at software level

- **Pros**
    - All channels can be part of the trigger
    - Sophisticated tagging/filtering algorithms
    - scalability
- **Drawbacks**:
    - we don’t have the same experience as for Triggered DAQ

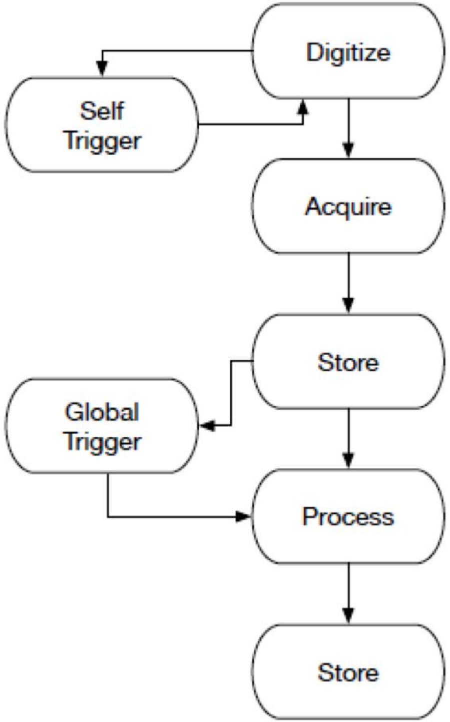

Why SRO is so important?

- High luminosity experiments
  - Current experiments are limited in DAQ bandwidth
  - Reduce stored data size in a smart way (reducing time for off-line processing)

- Shifting data tagging/filtering from the FEE (hw) to the back-end (sw)
  - Optimize real-time rare/exclusive channels selection
  - Use of high level programming languages
  - Use of existing/ad-hoc CPU/GPU farms
  - Use of AL/MI tools

- Scaling
  - Easier to add new detectors in DAQ pipeline
  - Easier to scale
  - Easier to upgrade

- Many NP and HEP experiments adopt the SRO scheme (with different solutions)
  - CERN: LHCb, ALICE, AMBER
  - FAIR: CBM
  - DESY: TPEX
  - BNL: sPHENIX, STAR,EIC
  - JLAB: SOLID, BDX, CLAS12,...
  - …
  -

SRO for CLAS12

- ➢ Installed at Jefferson Lab's experimental HALL-B
- ➢ Expansive program of physics topics:
  - investigation of the structure of the proton and neutron both in their ground state, as well as their many excited states
  - searches for exotic meson and baryon configuration
  - …
  - ….

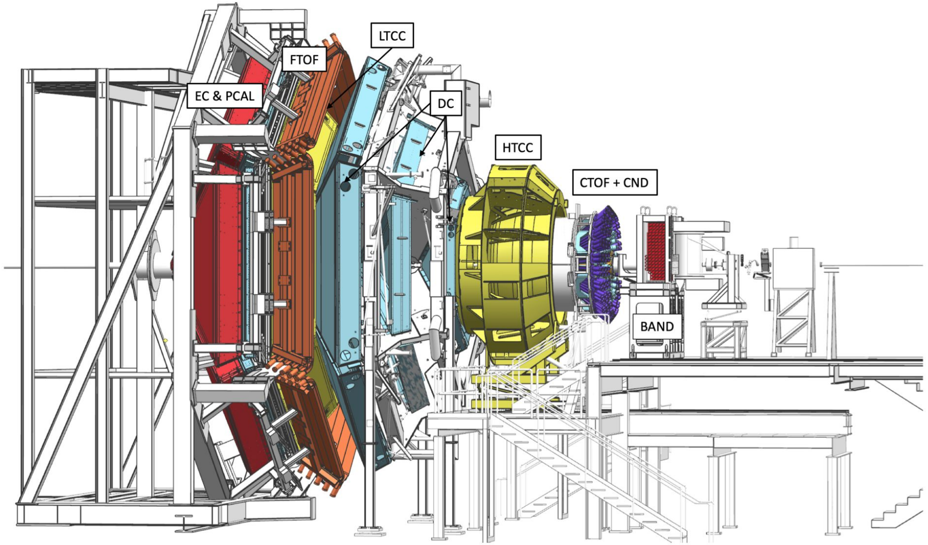

- Luminosity : $10^{35}cm^{-2}s^{-1}$
- Standard DAQ system
  - Event rate: 15-30kHz
  - Data rate: 500-1000 MByte/s
- Upgrade in luminosity: x2 in 2-3 years (Phase 1) and x100 in 5-7 years (Phase2)
- With the current triggered technology the maximum possible event acquisition rate for CLAS12 is ~100 kHz
- Streaming RO is necessary for a long-term HI-LUMI upgrade of CLAS12

SRO for EIC

- ➢ Future Electron ion collider @ BNL
- ➢ Purpose: to explore the quark and gluon structure of protons and nuclei
- ➢ Machine parameters
  - Luminosity up to 10^{34} $cm^{-2}s^{-1}$
  - Highly polarized (~70%) electron and nucleon beams
  - Center of mass energy: from 40GeV to 140GeV
- ➢ Detector requirements:
  - Large acceptance
  - Precise vertexing
  - Frwrd/Bckw angles
  - HRes Tracking
  - Excellent PID
- ➢ Three detector project shared the same SRO concept
  - Number of channels O(2M)

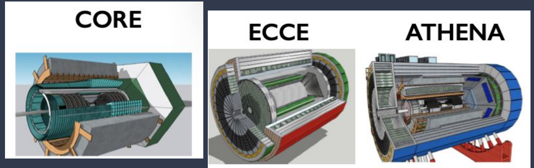

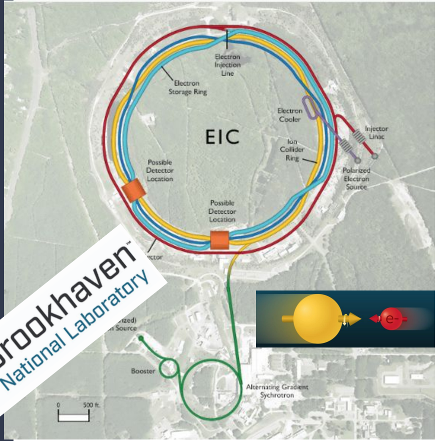

Yellow report streaming DAQ design

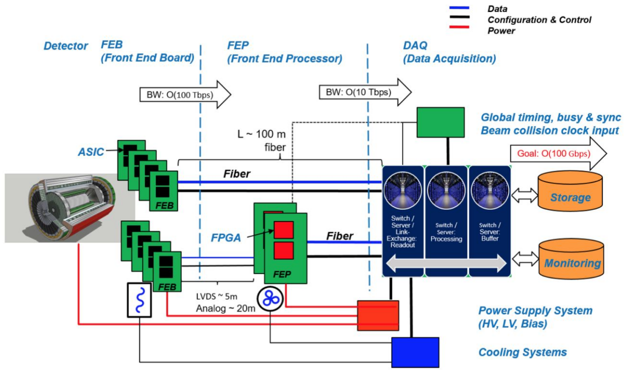

DAQ system requirements

- Collider
  - EIC will deliver beam in up to 1160 bunches (~100MHz)
  - Expected physics rate ~500KHz
- Detector
  - Streaming Design
    - Physics requires min-bias data
    - High collision rates
    - low detector occupancy
  - O(100Gbps) output rates

Streaming RO- Components

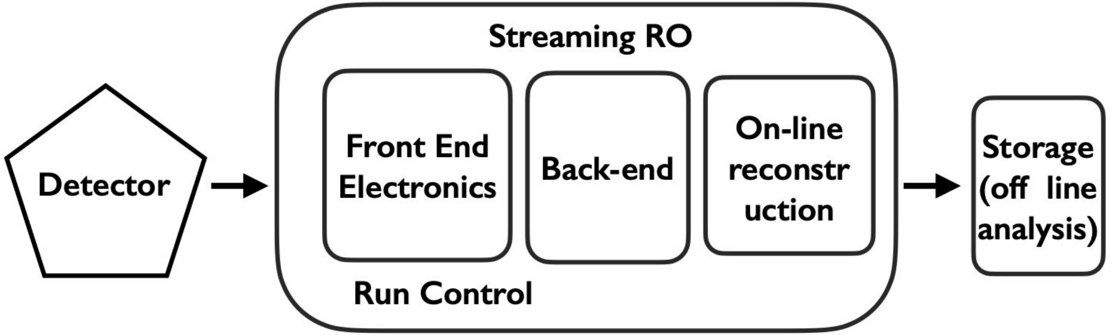

- SRO advantages are evident buy it needs to be demonstrated by the use in real experimental conditions
- To validate SRO concept:
  - Assemble SRO components
  - Test SRO DAQ
- **JLAB-INFN efforts to develop a prototype SRO DAQ**

Streaming RO prototype– FEE
JLAB FADC250 digitizer and VTP board

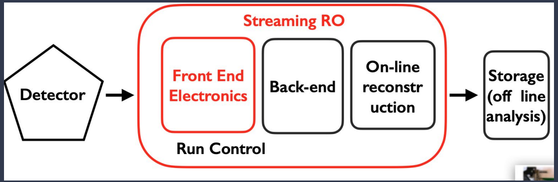

D.Abbott, C.Cuevas, B.Raydo

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

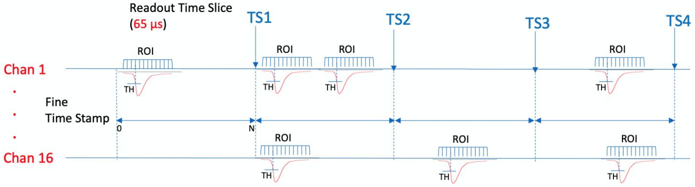

The next revision to the firmware will have an option for full ROI wave forms to be streamed, but this will allow possible dropped hits due to bandwidth limitations

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

Jefferson Lab

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

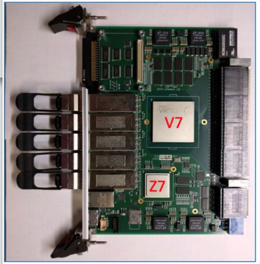

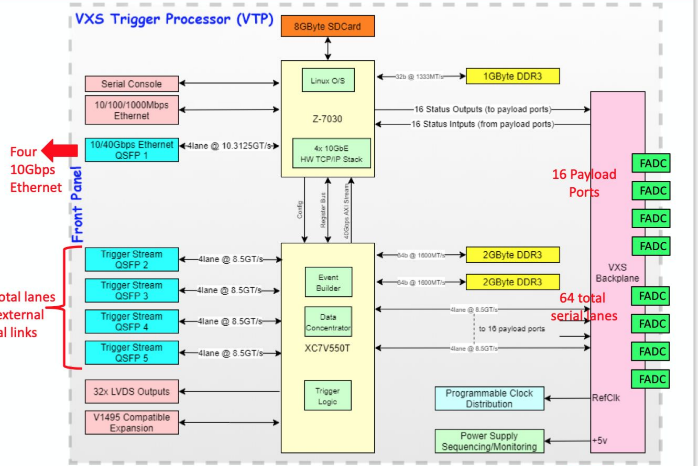

Jefferson Lab

Streaming RO prototype– Backend TRIDAS software

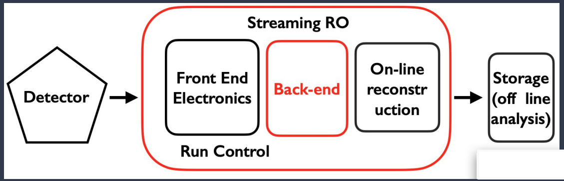

T. Chiarusi, C. Pellegrino, L. Cappelli

- Designed for the astrophysical neutrino detector prototype NEMO

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]

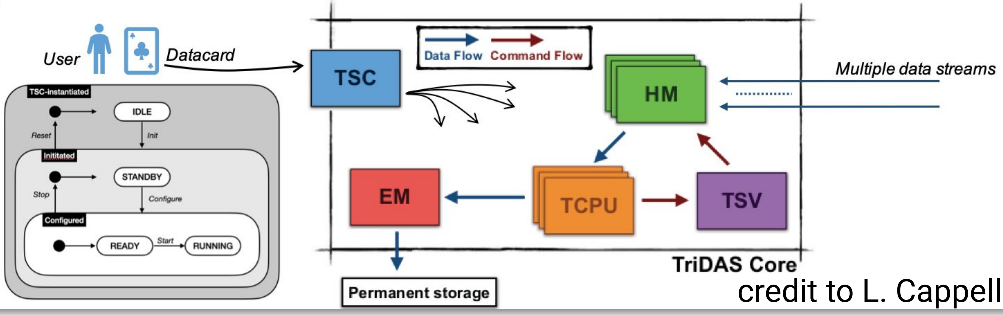

credit to L. Cappelli

Streaming RO prototype– Backend TRIDAS software

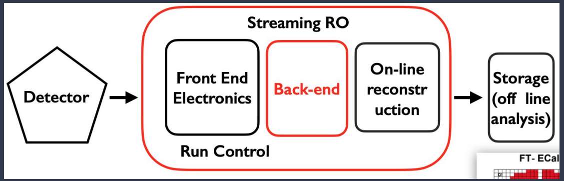

T. Chiarusi, C. Pellegrino, L. Cappelli

Data Flow: the HM aggregation

Farnesina Ministero degli Affari Esteri e della Cooperazione Internazionale INFN CNAF

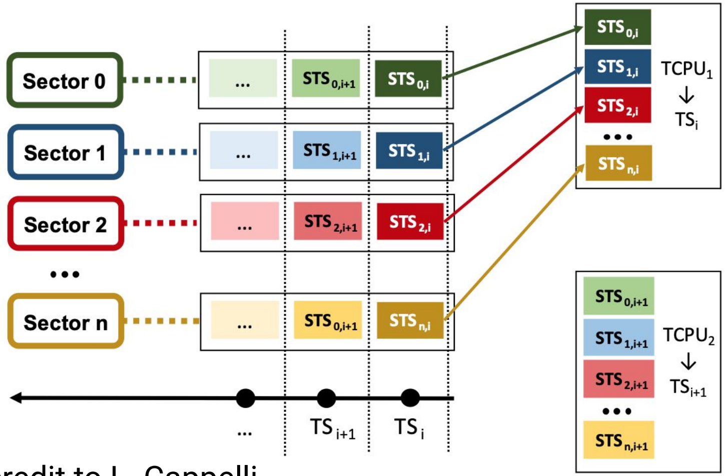

- Each HM:
  - Collects data from a specific sector of the detector
  - Subdivides the data into a sequence of time-ordered bunches called **Sector Time Slices (STSs)**
    - Fixed time duration called **Time Slice (TS)** chosen at run start time via datacard parameter (50ms in CLAS12)
  - Sends the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

credit to L. Cappelli

Streaming RO prototype– Backend TRIDAS software

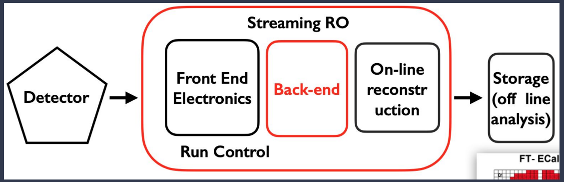

T. Chiarusi, C. Pellegrino, L. Cappelli

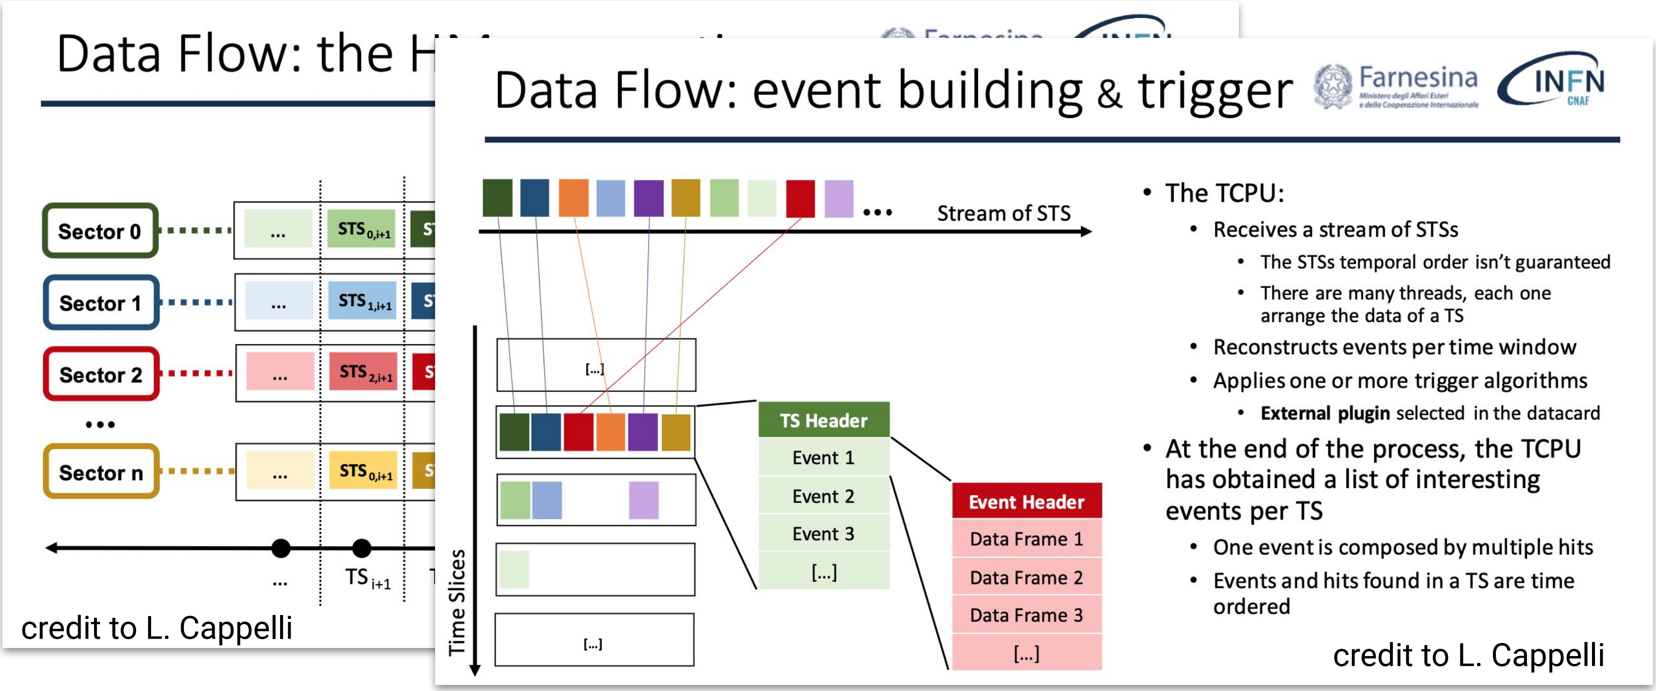

Streaming RO prototype– Backend TRIDAS software

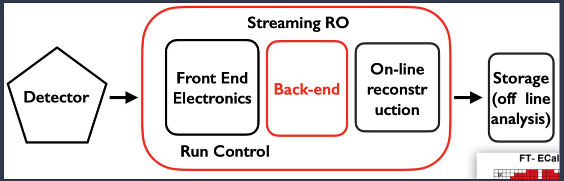

T. Chiarusi, C. Pellegrino, L. Cappelli

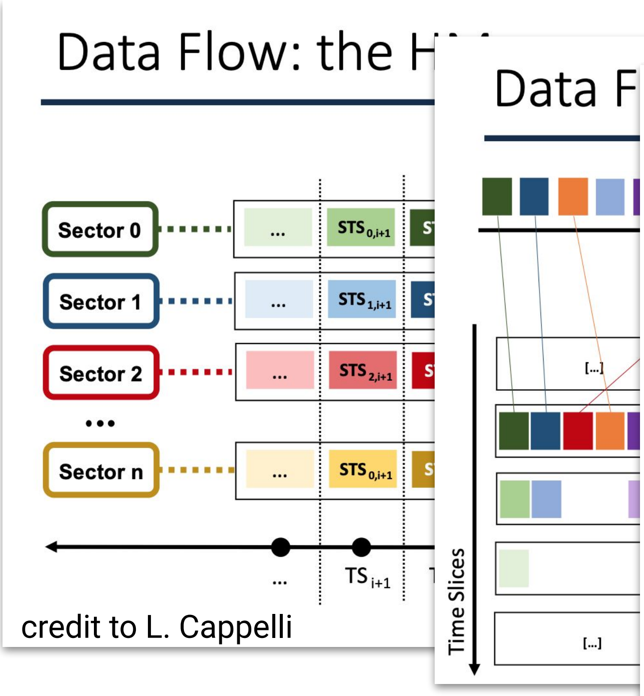

Data Flow: Event Manager

- The TCPUs send to the EM the triggered events for each TS
- The EM:
  - Writes the TSs into some post trigger files
- All the useful event information are stored for further analysis

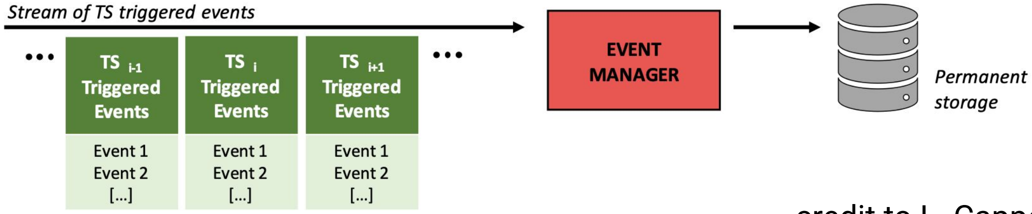

credit to L. Cappelli

Streaming RO prototype- Online-reco JANA2 software

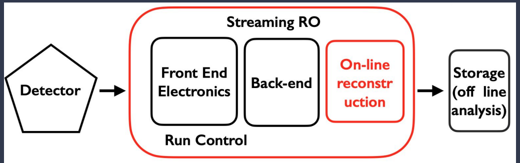

N.Brei, D.Lawrence

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

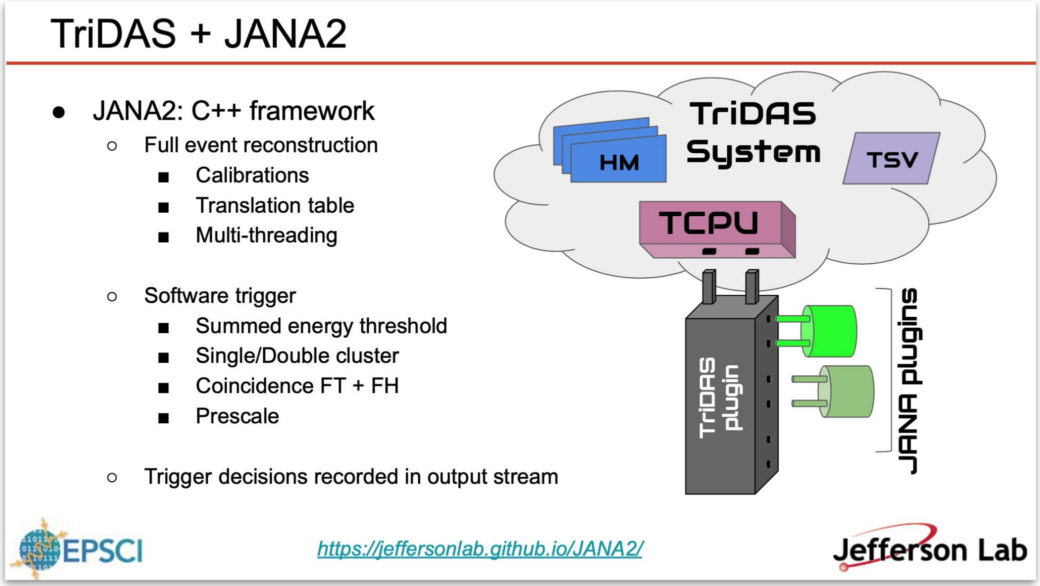

Streaming RO - CLAS12 FT test:

M. Battaglieri, M.Bondi, R. De Vita, S.Vallarino, A.Celentano, A.Pilloni, P.Moran

- **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 system performance
- Physics channel identification: pi0 production

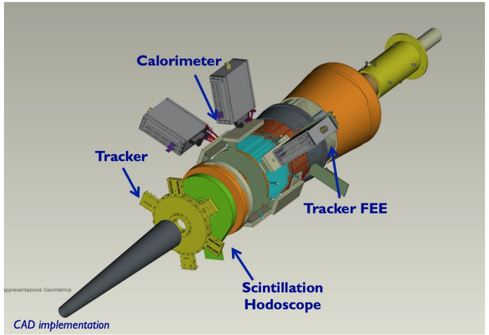

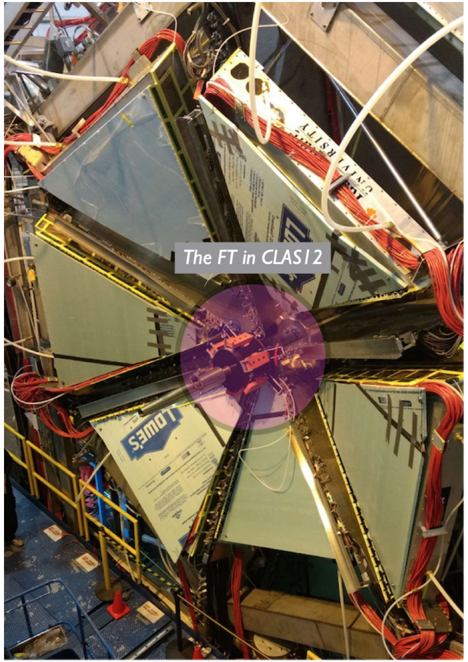

Streaming RO - CLAS12 FT test:

T. Chiarusi, C. Pellegrino, L. Cappelli, D.Abbott, C.Cuevas, B.Raydo, S. Boyarinov

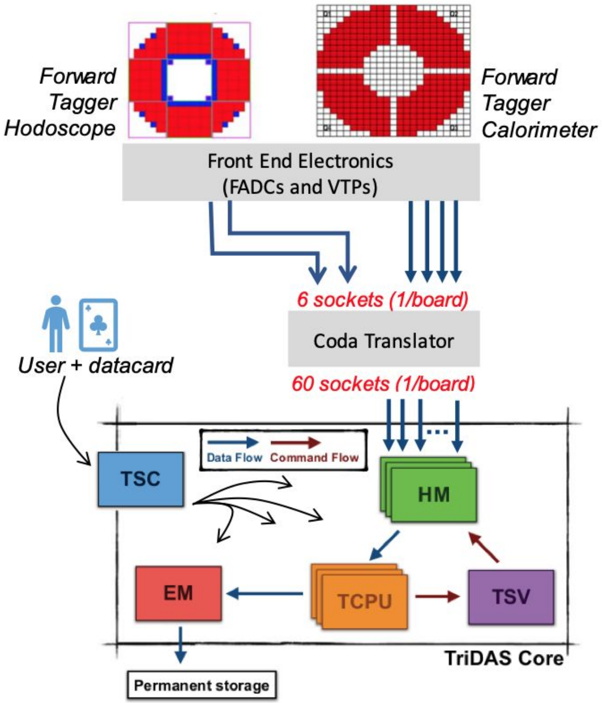

- System tested on the Forward Tagger sub-detectors
  - 6 streams: 2 streams from 3 VXS Trigger Processors (**VTP**)
    - Throughput: max 4 GBps per stream, set from few tens of MBps to 100 MBps
  - 16 Flash ADCs 250 (**FADC**) per VTP, 8 per stream
  - 16 channels per FADC
  - Total channels: 768 (16 ch * 8 FADC * 2 streams * 3 VTP)

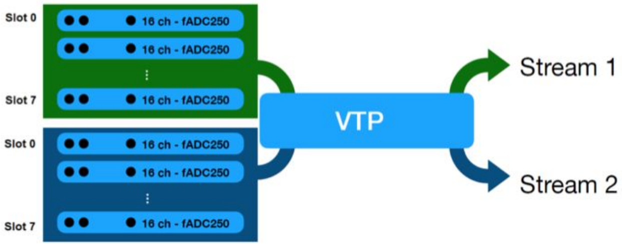

- Between the front-end electronics and TriDAS there was the **CODA translator** layer

Streaming RO - CLAS12 FT test:

T. Chiarusi, C. Pellegrino, L. Cappelli, D.Abbott, C.Cuevas, B.Raydo, S. Boyarinov

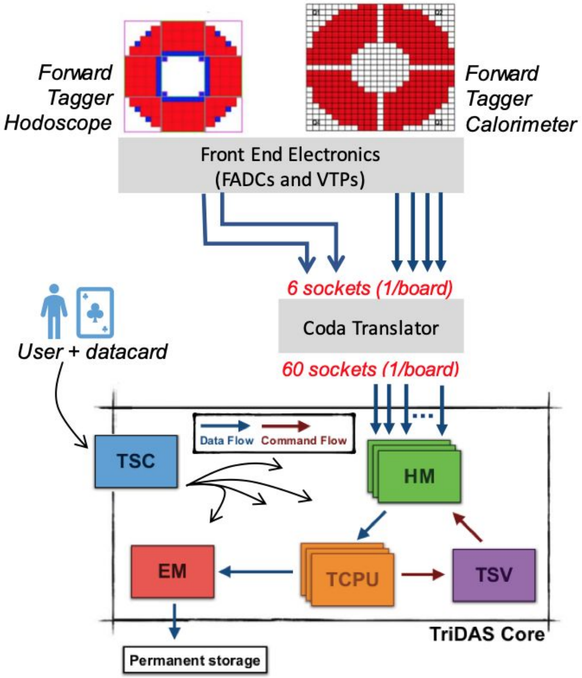

- Linux servers used:
  - 48 cores, 1GHz each, 64 GB RAM
  - 3 servers used for all modules
- HM instances: from 5 to 20
  - CPU consumption linear with the number of instances (500% – 1600%)
  - Memory occupancy constant (12-13 GB per run)
- TCPUs instances: 10 instances on 2 servers = 20 instances
  - 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

From:
F. Ameli et al. "Streaming readout for next generation electron scattering experiments". arXiv: 2202.03085

Streaming RO - CLAS12 FT test:

M. Battaglieri, M.Bondi, R. De Vita, S.Vallarino, A.Celentano, A.Pilloni, P.Moran

- **On-beam tests:**
  - 10.4 GeV electron beam on thin Pb/Al target
  - Inclusive pi0 production
    - e + Pb/Al ->(X) eπ^0 -> (X)e$\gamma\gamma$
  - Two gammas detected in FT-CAL
  - EM clusters identification, anti-coincidente with FT-Hodo

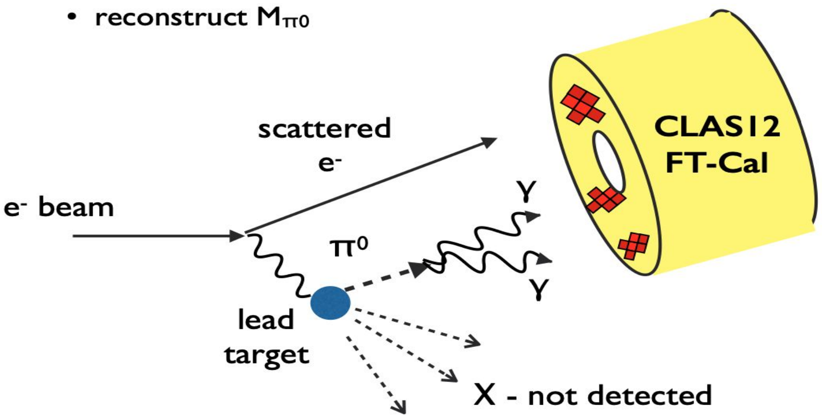

Streaming RO - CLAS12 FT test:

M. Battaglieri, M.Bondi, R. De Vita, S.Vallarino, A.Celentano, A.Pilloni, P.Moran

- **Realistic simulations:**
  - Realistic exclusive pi0 electro-production provided by JPAC
  - Realistic GEANT4 model of the FT detector
  - Contributions considered: electro-photoproduction by Pb

- Physics model of π^0 real photoproduction from JPAC (arXiv:1505.02321)
- Electroproduction simulated as quasi-real ph.prod. as in Tsai
- $2 < k_\gamma < 10$ GeV
- Acceptance $2^\circ < \theta_{\pi^0} < 6^\circ$, quite larger than the real one;
- Real acceptance (different for each target) from GEANT
- Other cuts from GEANT

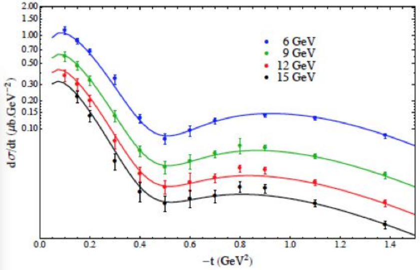

Expected yield (20mn run L=1e35 cm-2 s-1)

- From Lead ~1800
- From 160μm Al+glue ~420

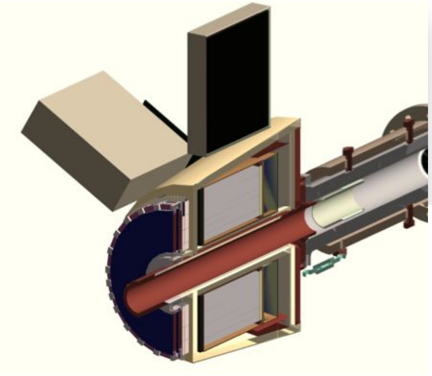

Streaming RO - CLAS12 FT test:

M. Battaglieri, M.Bondi, R. De Vita, S.Vallarino, A.Celentano, A.Pilloni, P.Moran

- off-line data analysis:
  - same full suite of reconstruction algorithm used in the on-line analysis
  - energy-calibration and time-walk correction
  - two targets: Pb(primary) + Al scattering chamber window
    - lower invariant mass due to the assumption that the vertex is located at the ... on when calculating the invariant mass.

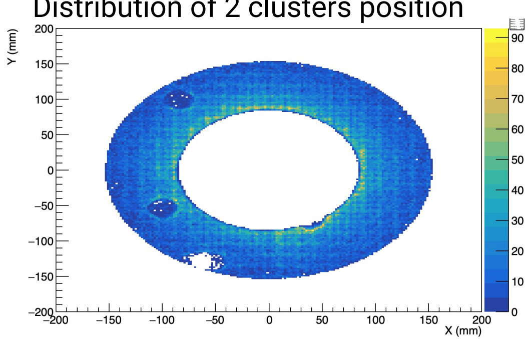

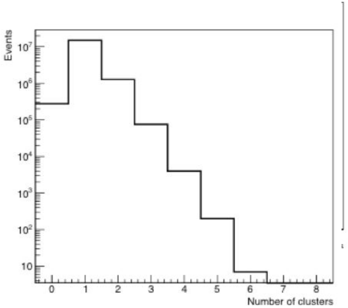

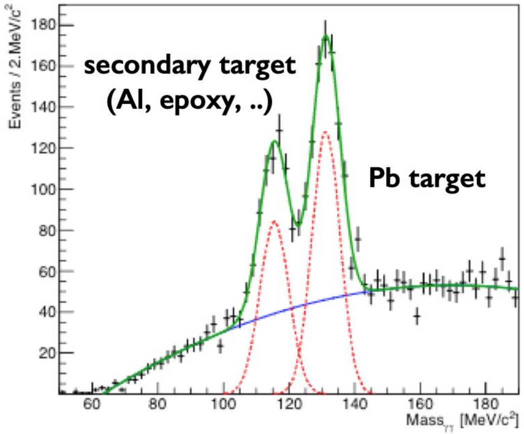

- Measured (expected) pi0 yield
  - Peak1 = 1365+-140 (~1800)
  - Peak 2 = 930+-100 (~420)

Summary

- SRO is the option for future electron scattering experiments
- 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 at JLab
- First SRO chain (FE + SRO sw + ON-LINE REC) tested with existing hardware

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