TDAQ group report

A. Gioiosa (Pisa), G. Pezzullo (Yale)

for the Mu2eII tdaq subgroup

Trigger/DAQ Group

Mailing list: mu2eii-tdaq@listserv.fnal.gov

- Antonio Gioiosa, Convenor, INFN Pisa
- Gianantonio Pezzullo, Convenor, Yale
- Richard Bonventre, LBNL
- Rebecca Chislett, UCL, Tracker
- Raffaella Donghia, LNF
- Bertrand Echenard, Caltech
- Ryan Rivera, FNAL
- Roberto Soleti, LBNL
- Franco Spinella, INFN Pisa – Calorimeter
- Craig Dukes, UVa – CRV
- Jinyuan Wu, FNAL

Architectures under study

Two TDAQ architectures proposed so far:

1. 2-level Trigger (L1 Trigger + HLT)
2. Software Trigger using GPUs

contributed paper for Snowmass22

IX. Trigger and data acquisition

A. Requirements
B. Architectures
    1. 2-level trigger, L1(FPGA)+HLT
    2. Software trigger with GPUs
C. Radiation
D. R&D

Requirements

IX. TRIGGER AND DATA ACQUISITION

A. Requirements

In order to set the requirements for the TDAQ system, we assume that Mu2e-II will adopt a similar experimental setup as Mu2e, but will improve granularity of detector elements up to a factor of 2. The direct consequences of these assumptions are:

- an increase in the event data size of a factor of $\sim \times 6$; $\times 3$ due to the instantaneous rate and $\times 2$ due to the number of channels, reaching a level of 1 MB/event;
- a reduced period when no beam is delivered to the apparatus, which in Mu2e is 1 s out of 1.4 s;
- a factor of $\sim \times 10$ larger dose on the electronics;

Assuming that the Mu2e-II storage capacity on tape will be twice that of Mu2e, reaching $\sim$ 14 PB/year (equivalent to a few kHz), the required trigger rejection needs to be a factor of ~ 5 better than in Mu2e, which is at the level of a few hundreds.

Generic Data Readout Topology

- Data Concentrator Layer
  - Aggregate small front-end fragments into larger chunks for efficient event building
- Event Builder Layer
  - Data is switched from Concentrator Layer to Event Builder Layer such that full events arrive at Event Builder Layer and are buffered
    - Preprocessing or filtering could occur
- Storage Decision Layer
  - Available decision nodes make high level storage decision on full events retrieved from Event Builder Layer buffer

Architectures

B. Architectures

Mu2e uses *artdaq*[423] and *art*[424] software as event filtering and processing frameworks respectively. The detector Read Out Controllers (ROC) from the tracker and calorimeter, stream out continuously the data, zero-suppressed, to the Data Transfer Controller units (DTC). The data of a given event is then grouped in a single server using a 10 GBytes switch. Then, the online reconstruction of the events starts and makes a trigger decision. If an event gets triggered, we pull also the data from the CRV and we aggregate them in a single data stream. Figure 24 shows a scheme of the Mu2e data readout topology described above.

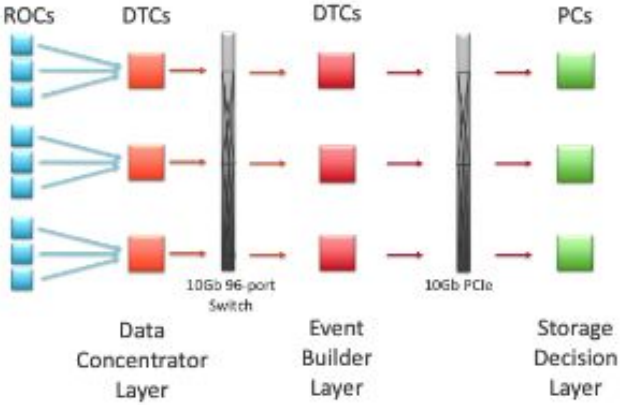

FIG. 24. Mu2e data readout topology.

The Mu2e main physics triggers use the info of the reconstructed tracks to make the final decision. The Mu2e Online track reconstruction is factorized into three main steps [425]: (i) hits preparation, where the digitized signals from the subdetectors are converted into reconstructed hits, (ii) pattern-recognition to identify the group of hits that form helicoidal trajectories, and finally (iii) track fit through the hit wires, which performs a more accurate reconstruction of the track.

L1 + High Level Trigger

- Aggregate the data into a board equipped with multiple FPGAs
- Run the early stage of the track reconstruction + full calorimeter reconstruction
- Apply a L1 decision and move data to the HLT farm which runs full track reco

1. 2-level trigger, L1(FPGA)+HLT

For Mu2e-II one of the ideas is to implement a L1 hardware trigger that exploits the first two stages of the Online track reconstruction on a dedicated FPGA based board and then exploit the rest of the reconstruction on the commercial servers. The major challenges are represented by: (a) the amount of data that needs to be concentrated on a single board and (b) the migration of a non negligible part of the Online reconstruction onto an FPGA. For the first one, the system will need to use more performant rad-hard optical transceivers (an R&D is already ongoing at CERN), which are needed to stream the data from the ROCs to the data-concentrator layer, and a more powerful switch (100 Gb switch are already available). For the second one, it's important to realize that FPGA development can take place now – hardware is not needed! Starting now would help the understanding of required resources and in consideration of topology trade offs. For example, what

L1 + High Level Trigger

- Aggregate the data into a board equipped with multiple FPGAs
- Run the early stage of the track reconstruction + full calorimeter reconstruction
- Apply a L1 decision and move data to the HLT farm which runs full track reco

eration of topology trade offs. For example, what size FPGA is best suited, or what are the advantages and disadvantages of commercially available hardware versus established custom boards in the community versus creating a new custom board. In the last decade, a new tool named High Level Synthesis (HLS) [126] has been developed to rival manual VHDL or Verilog algorithm development. The major HLS features are: (i) it allows non specialists to easily understand and develop low and fixed latency FPGA algorithms, (ii) it simplifies offline emulation, (iii) it facilitates debug and verify in a software environment (often 10x faster iterations than firmware simulation tools). We also note that other HEP collaborations, like the CMS experiment at CERN, have been heavily investing in HLS approach to FPGA algorithm development. Interestingly, the Mu2e run plan offers the possibility to test (parasitically) a prototype of a L1 trigger board in the second phase (after the LBNF shutdown). Leveraging Mu2e as a live data source would give valuable feedback for advancing Mu2e II's R&D phase.

Software Trigger on GPUs

Parallelize the reconstruction Algorithm:

- *Referring to Amdahl's Law, we could consider the KinKal parallelizable if the time execution on GPUs is at least one order of magnitude lower than the execution with CPUs only*

- *Algorithm depending*

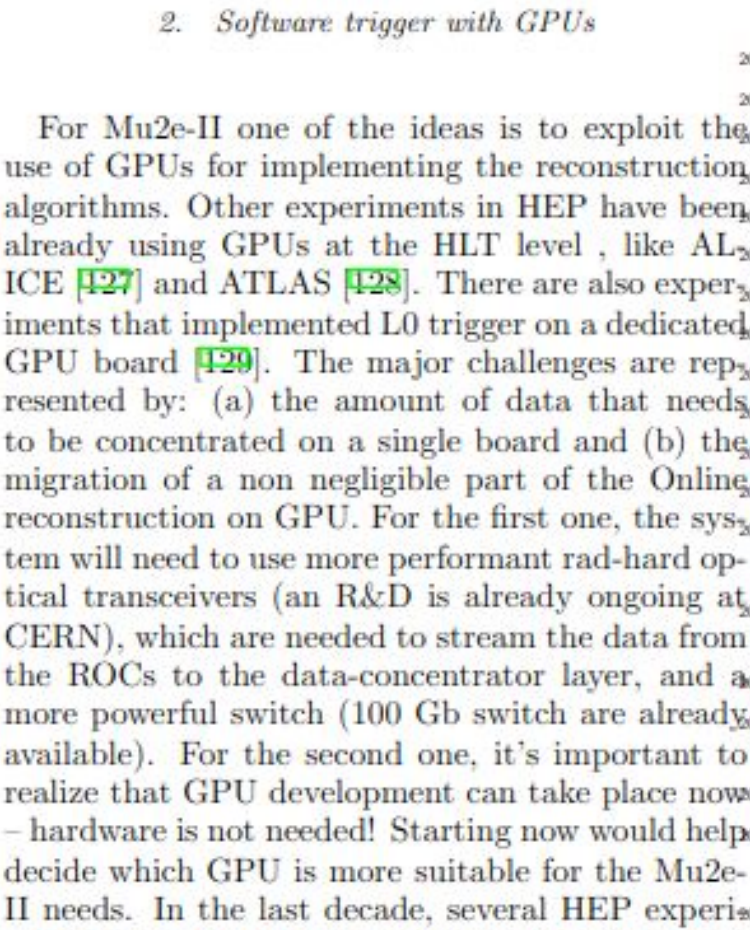

2. *Software trigger with GPUs*

For Mu2e-II one of the ideas is to exploit the use of GPUs for implementing the reconstruction algorithms. Other experiments in HEP have been already using GPUs at the HLT level , like ALICE [127] and ATLAS [128]. There are also experiments that implemented L0 trigger on a dedicated GPU board [129]. The major challenges are represented by: (a) the amount of data that needs to be concentrated on a single board and (b) the migration of a non negligible part of the Online reconstruction on GPU. For the first one, the system will need to use more performant rad-hard optical transceivers (an R&D is already ongoing at CERN), which are needed to stream the data from the ROCs to the data-concentrator layer, and a more powerful switch (100 Gb switch are already available). For the second one, it's important to realize that GPU development can take place now – hardware is not needed! Starting now would help decide which GPU is more suitable for the Mu2e-II needs. In the last decade, several HEP experi-

Software Trigger on GPUs

- In the last decade, several HEP experiments (like ATLAS and NA62) exploited the use of GPUs in their TDAQ systems

Radiation and R&D

- We are collecting last informations and will complete these paragraphs