

arXiv:2203.07620

# Innovations in trigger and data acquisition systems for next-generation physics facilities

Authors:

Rainer Bartoldus (SLAC),  
Catrin Bernius (SLAC),  
David W. Miller (Chicago)

SLACMass meeting,  
12 May 2022

**SLAC** NATIONAL  
ACCELERATOR  
LABORATORY



Submitted to the Proceedings of the US Community Study  
on the Future of Particle Physics (Snowmass 2021)

Innovations in trigger and data acquisition systems for  
next-generation physics facilities

A Snowmass 2022 White Paper

Rainer Bartoldus<sup>\*1</sup>, Catrin Bernius<sup>1</sup>, and David W. Miller<sup>12</sup>

<sup>1</sup>*SLAC National Accelerator Laboratory*  
<sup>2</sup>*Department of Physics, University of Chicago*

[arXiv:2203.07620](https://arxiv.org/abs/2203.07620)

- **TDAQ systems** are a critical interface between detector instrumentation and computing systems at current and future physics facilities
- Direct impact on the volume and quality of the datasets produced and the physics that can be extracted from them

⇒ **Innovative solutions to the new challenges being faced by these facilities will be essential**

# Challenges and Opportunities for TDAQ systems

- **Novel applications and uses of current generations of commodity and custom hardware, firmware, and software** for TDAQ systems
  - Known obstacles for TDAQ systems
  - Trend towards “heterogeneous” computing with accelerators
- Confronting state-of-the-art devices and TDAQ paradigms with the **specialized needs of future low-latency, high-throughput, and high-performance** physics facilities
  - Complex and extreme processing environments
  - Dramatically higher numbers of sensors and channels
- **Building, integrating, commissioning, and operating** large-scale, heterogeneous, and dynamic TDAQ systems at future facilities, including the **acquisition and retention of the required domain knowledge and technical expertise**
  - Bringing a novel idea to life
  - Growing and maintaining a technical knowledge base
  - Attracting and supporting talent

# Challenges faced by current & future facilities (1)

## Ever-evolving **operating conditions of particle accelerators**

- Higher energies, intensities, and data rates
- Signatures of interest become more and more challenging to analyze and select in the high particle density environment

### Example:

- **Fast** track reconstruction at the trigger level
  - Search for LLP (e.g. disappearing tracks) and exotic processes (e.g. R-parity violating SUSY)
  - Speed up with embedding algorithms in highly-parallelized and potentially heterogeneous computing architectures

⇒ Use of **heterogeneous digital and analog solutions** (CPUs, GPUs, FPGAs, ASICs, SOCs, ...) together with **advanced algorithms** (Artificial Intelligence (AI), Machine Learning (ML))

# Challenges faced by current & future facilities (2)

Upgraded experimental designs lead to **expansion** of both **fiducial detector volumes and channel multiplicities** by orders of magnitude

- Signal extraction in the **low-energy regime** challenging due to the enormous amount of additional, background-like contributions (e.g. DUNE Far Detector)
- **High data rates**, requires a scalable architecture of the DAQ system as well as specialized algorithms that are employed for signal identification and reconstruction (e.g. Project 8)
- **Computational approaches** can be limiting data analyses
  - Example: Multi-messenger astrophysics facilities (MMA)
  - Challenges include the characterization of an **extremely high number of data channels** to provide low-latency data quality evaluation, optimizing and fitting data in high-dimensional space, ...
  - Limiting factor for low latency estimates of gravitational wave signal significance to enable electromagnetic follow-up of gravitational wave sources

# Heterogeneous and dynamic systems

**Heterogeneous computing** can help address the needs of **high-throughput and high performance compute power** together with **machine learning algorithms** which are well-suited to run on these new computing architectures

- LHCb experiment already using fully GPU-based implementation of first level trigger
- Various design proposals with heterogeneous usage for upgrade of Mu2e experiment (Mu2e-II)
- Novel physics-inspired neural network image recognition techniques for differentiating signals, for example of Higgs bosons, from the background physics processes using MPSoC and FPGA devices
- Concept of self-driving trigger system to autonomously and continuously learn to select, filter and process data
- [FPGA-based artificial intelligence inference in triggered detectors](#), LoI submitted by SLAC collaborators (R. Herbst, A. Dragone, M. Kagan, L. J. Kaufman, B. Mong, and R. Teixeira De Lima)

# Hyper-dimensional DAQ and processing systems at scale



Some challenges may require specialised solutions that may or may not exist for more common industrial applications, e.g.:

- Extremely high rate applications and high bandwidth requirements
- Highly tailored or varying temporal data structures
- Data integrity and robustness
- Reproducibility requirements

⇒ **Further push the envelope of industry standard devices and systems** to confront and challenge the commodity solutions with the specific needs of **low-latency, high throughput and high performance** physics facilities

## Examples:

- Asynchronous first level trigger systems
- Edge computing devices for detectors
- Streaming DAQ

# Installation, commissioning, integration, and operations



Bringing the innovative detector-related research, development and design and construction to life

- The transition from final testing to their installation brings new challenges for new components
- Upgrades of today's facilities bear additional challenges of getting new components to work with existing ones

It is immensely important to **invest adequate effort and pay equal attention** to the installation, commissioning, integration and operation

Without these crucial efforts, **data-taking will be compromised and science goals will not be met.**

- Experience has shown that lack of attention can lead to substantial and costly delays, failure to reach design goals and insufficient performance, all at the cost of the physics output
- High-efficiency operations demand substantial level of effort and has to be maintained as long as the experiment is running



# Building and retaining domain knowledge & technical expertise

*"[Shifters and experts] wanted for hazardous journey. Small wages, bitter cold, long months of complete darkness, constant danger, safe return doubtful. Honour and recognition in case of success."*

– Alleged advertisement by Sir Ernest Shackleton for the voyage of the *Endurance* to Antarctica

Critical for our field to **recognize the importance of the highly skilled people** that are carrying out the deployment part of the innovation process

- Without the technical and operational skills contributed by a wide variety of personnel, it is not possible to deliver discovery science at physics facilities
- Availability and recruitment is problematic due to the challenging prospects of **career progression and promotion**
- **Long-term retention** is a serious problem for building a community of experts who are able to operate the experiment effectively
- Scientists who take such opportunities, e.g. to engage in operational efforts, must be considered an **integral part of the science programs and the facilities themselves**

# Summary: Three main themes

---



- Future TDAQ systems will have to leverage currently existing and available hardware, firmware, and software in new ways
- TDAQ developers will have to find ways to adopt and incorporate the unique technical requirements of physics facilities into new hardware concepts
- The TDAQ community must ensure that the people, knowledge, and experience required to build, commission, and operate new TDAQ systems are fully supported and retained in the field

# BACKUP

SLAC group meeting,  
25. Mar 2022

**SLAC** NATIONAL  
ACCELERATOR  
LABORATORY



# Example: Asynchronous first-level trigger systems

Increasingly **challenging to distribute** and synchronize stable, low-jitter, high-frequency **clocks over very large systems** of thousands of optical links

- Precision of timing to reach  $O(10 \text{ ps})$  and number of boards greatly increases with granularity

Rather:

- Tag data with **time marker only at detector front-end**
- Then **transport and process the data asynchronously** up to the first level trigger decision
  - Same as already done for higher-level triggers
- **Event builder runs at full rate from detector**
- First-level trigger moves from clock-driven to event-driven
- Some portions of LHC first-level triggers already asynchronous
  - Trigger optical links not multiple of machine frequency
  - Time de-multiplexing scheme that serves one bunch crossing worth of detector data to an individual trigger board
- **Blurs the lines between first-level (FPGAs) and higher levels (CPUs and now also GPUs)**

A set of FGPAs, CPUs and GPUs used to execute a mix of very fast and very complex algorithms

# Example: Streaming DAQ (1)

Traditional TDAQ architectures, for decades, based on these fundamental assumptions:

- It is **impossible to read out** the detector **at full rate**
- Even if it were possible, one **could not process** the data **at full rate**
- Even if that could be done, it **would create unsustainably large datasets** for offline computing

Now: continued advances in computing, networking and storage technologies have challenged these assumptions

Most of today's TDAQ systems have **pipelined and triggered detector readout**

- Detector signals are digitized at fixed clock rates, written to **circular buffers**
- A trigger system reads part of these data, often over separate links, processes them, and decides whether to keep or discard the buffer
- If the buffer is to be kept, full readout is initiated over the DAQ path and sent to the next stage

Hugely successful in working around rate and bandwidth limitations, but known shortcomings

# Example: Streaming DAQ (2)

## Pipelined mode limitations:

- Latency (processing time) limited by depth of front-end buffers (electronics)
- Trigger can become very complex to achieve rate reduction
- Readout data frames arrive randomly due to trigger latency (and jitter)
- Downstream event building has to collect and arrange all contributions **while data are in motion**

## Known downsides:

- Trigger decision introduces **selection biases**
- Doesn't deal well with **event pileup** (what is an “event” becomes a trigger artifact)
- Performance ultimately constrained by latency and the need to execute within **deadtime** path
- Data reduction comes at a **loss of information**

## Alternative: Streaming mode

- Read out many parallel continuous streams of data, encoded with their original time and location
- Flow of data can be reduced by applying thresholds and zero-suppression **locally**

# Example: Streaming DAQ (3)

Streaming mode advantages:

- Many streams can be **aggregated** into single link or transport channel
  - Reduction can be substantial because detectors with high number of channels typically have low occupancy
  - Aggregation is time-based and agnostic to data payload: can use generic components
- Streams can be **translated** with one type in and another type out
  - E.g. clusters from hits on the fly
- Event building now takes place **when the data are at rest**
- Further event processing, filtering (or High-Level “triggering”) can be done asynchronously
  - Not limited by the front-end electronics, but the depth of the intermediate storage
- **Unbiased raw data samples** for calibration and monitoring without the need for special runs
- Leverages high-performance computing for a **much simpler architecture**
  - no complex custom-hardware trigger, no separate trigger path, but a **deterministic time-ordered data** transport system with tiered storage and parallel processing

Concept can be taken to the next level (HLT) for data scouting and Trigger-Level Analysis

# Outlook and next steps

---

- We appear to be on a technology threshold where innovations are beginning to shift long held paradigms of TDAQ systems. Those innovations will be essential for the success of future physics facilities.
- Current experiments may have pushed scales to practical limits. When previous leaps were to big data and massive CPU farms near the control rooms, we now see an explosion of real-time processing and heterogeneous, accelerated computing
- These innovations can only be brought to life if we acquire new domain knowledge, and if the next-generation facilities attract and retain highly skilled people and support their careers.