Triggering and Reco in HEP: a comparison across experiments (we're not as different as we think!)

Reco & SW Trigger WG Kick-Off Meeting
26th April 2023

HEP Software Foundation

Intro

- We started thinking how our experiment’s complex SW trigger systems came to be as they are today
- We realised that while physics targets drives most of the differences, the overall experiment design (including the accelerator) should not be forgotten
- Despite large differences in type of physics, accelerator system and detector design, it was easy to find similarities across the different experiments
- We would like to highlight the similarities that we found…
- and how we, the SW-HEP community can move towards **common solutions to common problems**

Disclaimer: the following “comparison” exercise is heavily biased from our own experiments, and is in no way an exhaustive view of all SW trigger systems in HEP… so we welcome input and ideas from all of you!

FASER

Search for decay of long-lived particles far from proton-proton collisions @ the LHC.

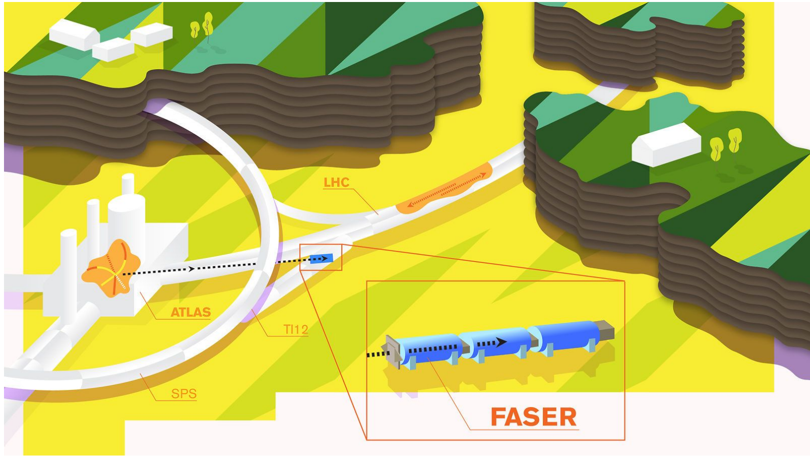

An LHC experiment without the DAQ challenges

- 1-level HW trigger, DAQ output BW:
  32 MB/s (1.2 kHz peak event rate)
- All inclusive triggering with redundancy for long-lived particle searches
  Thresholds below MIP, all seen signals collected
- Uses common DAQ framework: DAQling from EP-DT DI CERN division.
- All reconstruction offline: Using *Calypso*, derivative of ATLAS Athena (Gaudi-based) software framework.

But with own “small experiment” challenges:

- Small team of people & shifterless: Largely runs by itself - alert system and monitoring for quick diagnosis important.
- Limited funding + limited access to experiment site; Yet looks for very rare signal:
  - Keep as simple as possible, avoid complexity online, monitor everything.

ATLAS

Large general-purpose detector @ LHC
proton-proton (or heavy ion) interaction point.

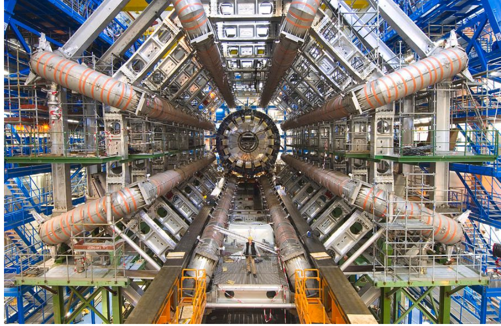

Great complexity & resource challenges online & offline

- 2-tier trigger: L1 HW + HLT SW trigger
- DAQ output BW: ~3 GB/s at ~7 (~3) kHz event rate for all (full event physics) data.
- Highly selective triggering w/ strong channel dependence
  (~50% central Z —> ee, ~50% central HH->4b)
  High trigger thresholds for hadronic jets -> focus on new physics at high masses.
- Large resource usage both online and offline
  - Online reconstructs partial or full event to offline-like quality
    CPU and readout limited at high beam intensities.
  - Offline reconstructs event from raw detector input again.
    with exception of data scouting stream but full offline reco default for analysis.
  - Production of simulated events, using same offline resource & with same statistics as triggered data
    Analyses largely use simulation-driven background estimates.
- Software: Use of multithreading & multiprocessing; No accelerated computing for Run 3, but lots of R&D for HL-LHC.

Belle II

High-precision measurements and search for new physics signature at e+e- collisions @ SuperKEKB

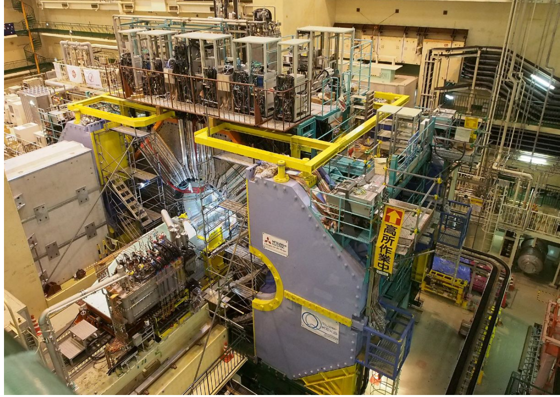

Challenges at the high-luminosity frontier @ e+e-

- 2-tier trigger: L1 HW + HLT SW trigger
- DAQ output BW: ~3 GB/s at 10 kHz event rate for all data (300 kB/evt)
- Very-high-efficient triggering (100% on e+e- → qqbar)
- Large resource usage both online and offline
  - Online sw reconstructs full event (same as offline, pixels not available)
    a challenge at high lumi, mostly because of machine backgrounds (injection, single beam, lumi)
  - Offline reconstructs event from raw detector input again.
    with the addition of pixels, and the best calibration (critical for high precision measurements!)
  - Production of simulated events, using same offline resource & with same statistics as triggered data Analyses largely use simulation-based background estimates.
- Ramping up toward (even) higher luminosity

LHCb

A forward arm spectrometer for heavy flavour physics @ LHC
proton-proton, PbPb, and fixed target with gas injection

Online reconstruction & purely software-based trigger in high intensity conditions

- Two software stages: HLT1 & HLT2 , splitting enables online alignment & calibration
- First stage (HLT1) partial event reconstruction at **30 MHz** (5 TB/s) on Nvidia **GPUs** -> **output rate 1 MHz** (O(100) GB/s)
- Second stage (HLT2) full reconstruction with offline quality on CPU farm -> **output BW 10 GB/s**
- Highly efficient triggering on beauty and charm signatures, now expanding
- Real Time Analysis paradigm - majority of processing online
    - Selective persistency (TURBO) for ~70% of events to reduce storage needs
    - Full raw information & offline re-processing for 30% of events - majority of data volume
    - Offline resources heavily used for simulation - exploring additional resources & fast simulation to cope with increasing simulation needs

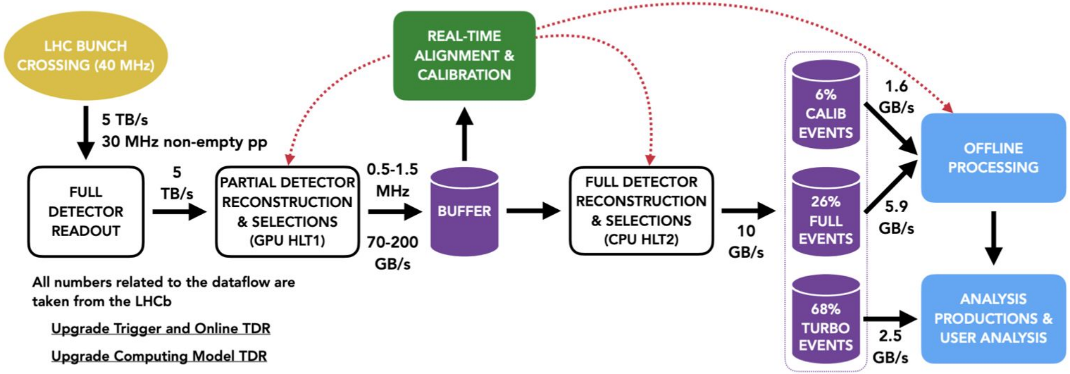

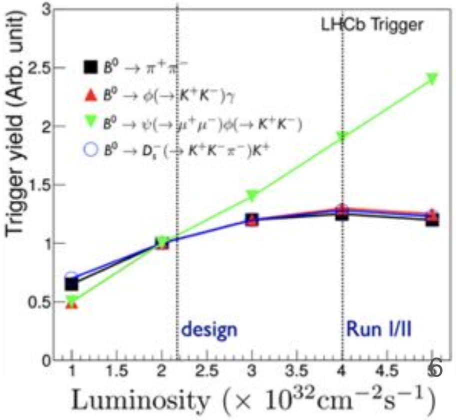

Interplay with Accelerator (examples)

- Collisions seen by Detectors:
  - Belle II sees basically continuous collisions (~ uniform in time)
  - experiments @ LHC see each collision in a separate trigger
- Beam Lifetimes:
  - much smaller @ e+e-, needs continuous injection -> high(er) beam backgrounds, noise in the detector from particles not belonging to the triggered collision
  - LHC, very stable beam -> uniform data quality vs machine timescales

experiment	RF frequency	min bunch separ	beam lifetimes
LHC	400 MHz	40 MHz / 25 ns	~ 10 hrs
SuperKEKB	~509 MHz	256 MHz / 4 ns	order of minutes

Comparison of some triggers parameters (examples)

- Rich variety, depending on beam/collision properties and physics goals

experiment	trigger hw/sw	trg throuput	trg jitter	trg efficiency
Atlas	L1 hw + HLT sw	~3-5 GBs/s	synched to LHC clk	signature dependent
LHCb	sw	70-200 GB/s (HLT1) - 10 GB/s (HLT2)	synched to LHC clk	signature-dependent, > 80% for key rare B channels
FASER	hw	~37 MB/s	synched to LHC clk	target 100% eff for LLP decay
Belle II	L1 hw + HLT sw	~3 GB/s	o(10 ns), varies with event type	goal is 100% on e+e- -> qqbar

Reconstruction requirements (examples)

- Online reconstruction needed to (further) reduce data rate
  - More important for experiments with high throughput
  - Quality depends on resources & online calibration / alignment
- Offline reconstruction not always needed but often desirable
  - usually offline calibration allows to improve the data quality
  - can profit from new algorithms when raw data are available
  - but computationally expensive, extra resources needed
- Persisted information
  - Persisting low-level information allows for offline re-processing, easier efficiency estimation, error spotting & handling but, storage intensive
  - Selective persistency models more storage efficient, but care needed to ensure inclusiveness (physics that will not be explicitly looked for will be physics discarded)

experiment	need online reco?	record all event?	bulk of processing?	calibrations
ATLAS	yes	yes (though special streams exist)	online & offline	online & offline
LHCb	yes	yes for HLT1, depends for HLT2	online	yes, mostly online
FASER	no	yes	offline	offline
Belle II	yes	yes and store raw data	online & offline	yes, offline

Example common challenges/solutions

- Biased triggering & missed physics
  - CMS, LHCb, and ATLAS all reworked their trigger system to do more scouting/turbo/trigger-level-analysis
  - How do these selective persistency models affect trigger “inclusivity”?
  - Strategies to mitigate biases?
  - What is the right balance between inclusive / exclusive triggers?

- Moving towards more sustainable software
  - Increasing computing requirements doesn’t come for free
  - Evaluate cost not only in financial terms, but also in terms of environmental impact
    - What are useful key indicators to judge the energy consumption of an online/offline data processing system?
    - What are good practices for programmers, analysers, experimentalists designing/upgrading experiments?
    - What about power consumption of the SW reco/trigger infrastructure (storage, server cooling etc)?
    - Power profile comparison of hardware used for SW triggers (CPUs, GPUs, FPGAs etc)

Example common challenges/solutions

- Reconstruction in high luminosity conditions:
  - Pushing towards higher luminosities -> Pileup becomes an increasing problem
  - Most LHC experiments are now including **precision timing** in their reconstruction strategy to distinguish vertices and suppress PU-induced backgrounds
  - What are the common 4-D reconstruction/trigger challenges? Could common solutions be envisioned ?
  - Interface between hardware / firmware / software projects?
  - Could other experiments benefit from all the developments?

- Heterogeneous architectures for reconstruction and SW triggering
  - High throughput experiments start departing from “classic” pure-CPU architectures for their SW processing
  - CMS, ALICE & LHCb already ported parts of their SW trigger to GPUs and FPGAs for Run 3 of LHC - more experiments investigating similar architectures
  - **Lessons learned:** Were there/are unexpected challenges? How was the interplay between hardware and software commissioning? Reproducibility of performance offline / in simulation ?
  - How will the heterogeneous resources be best used outside of data-taking?
  - How will the systems scale with future increases in luminosity?