LHCb Computing

Computing for the LHCb Upgrade

LHCb Upgrade: goal and timescale

- **LHCb upgrade will be operational after LS2 (~2020)**

- Increase significantly the statistics collected by the experiment, keeping the present excellent performance
  - Raise operational luminosity to a levelled 2×10^{33} cm^{-2}s^{-1}
    - Necessitates redesign of several sub-detectors
    - Does not require HL-LHC
  - Full software trigger at 40 MHz bunch crossing rate
    - Allows effective operation at higher luminosity
    - Improved efficiency in hadronic modes
    - Necessitates upgrade of the DAQ and HLT

- The gain is a huge increase in precision, in many cases to the theoretical limit, and the ability to perform studies beyond the reach of the current experiment
  - Flexible trigger and unique acceptance also opens up opportunities for other topics apart from flavour

LHCb Upgrade Computing Timeline

- 2020 is tomorrow
  - No time (or effort) for major changes in technology
  - Focused R&D based on existing experience
  - Possibility to use Run 2 as a test bed for new ideas
- Roadmap document for TDR by Q4 2015
  - Specify R&D required for informed decisions in TDR
- All R&D reports ready Q1 2017
- TDR scheduled for Q3 2017
  - Baseline technology choices made
- Computing model finalized Q4 2018

Trigger at the upgrade

- Trigger-less readout at full LHC crossing rate
  - No hardware (L0) trigger
- First and second level software trigger (HLT1/2) running on Event Filter Farm
  - Full deferral as in Run 2
    - Offline quality detector calibration
- Event Size:
  - 100 KB maximum (constraint from readout system)
    - 10-20 times smaller for channels going to TURBO stream

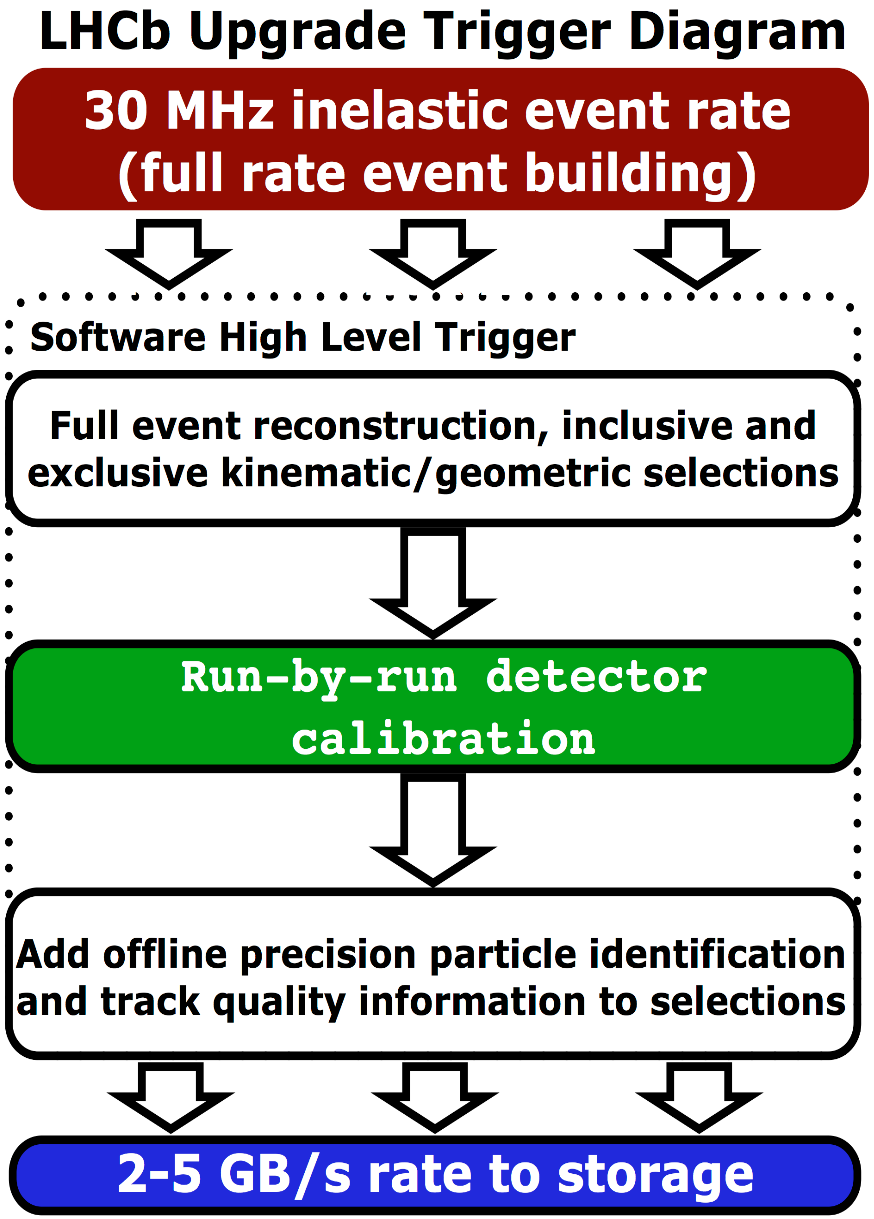

Trigger now and at the upgrade

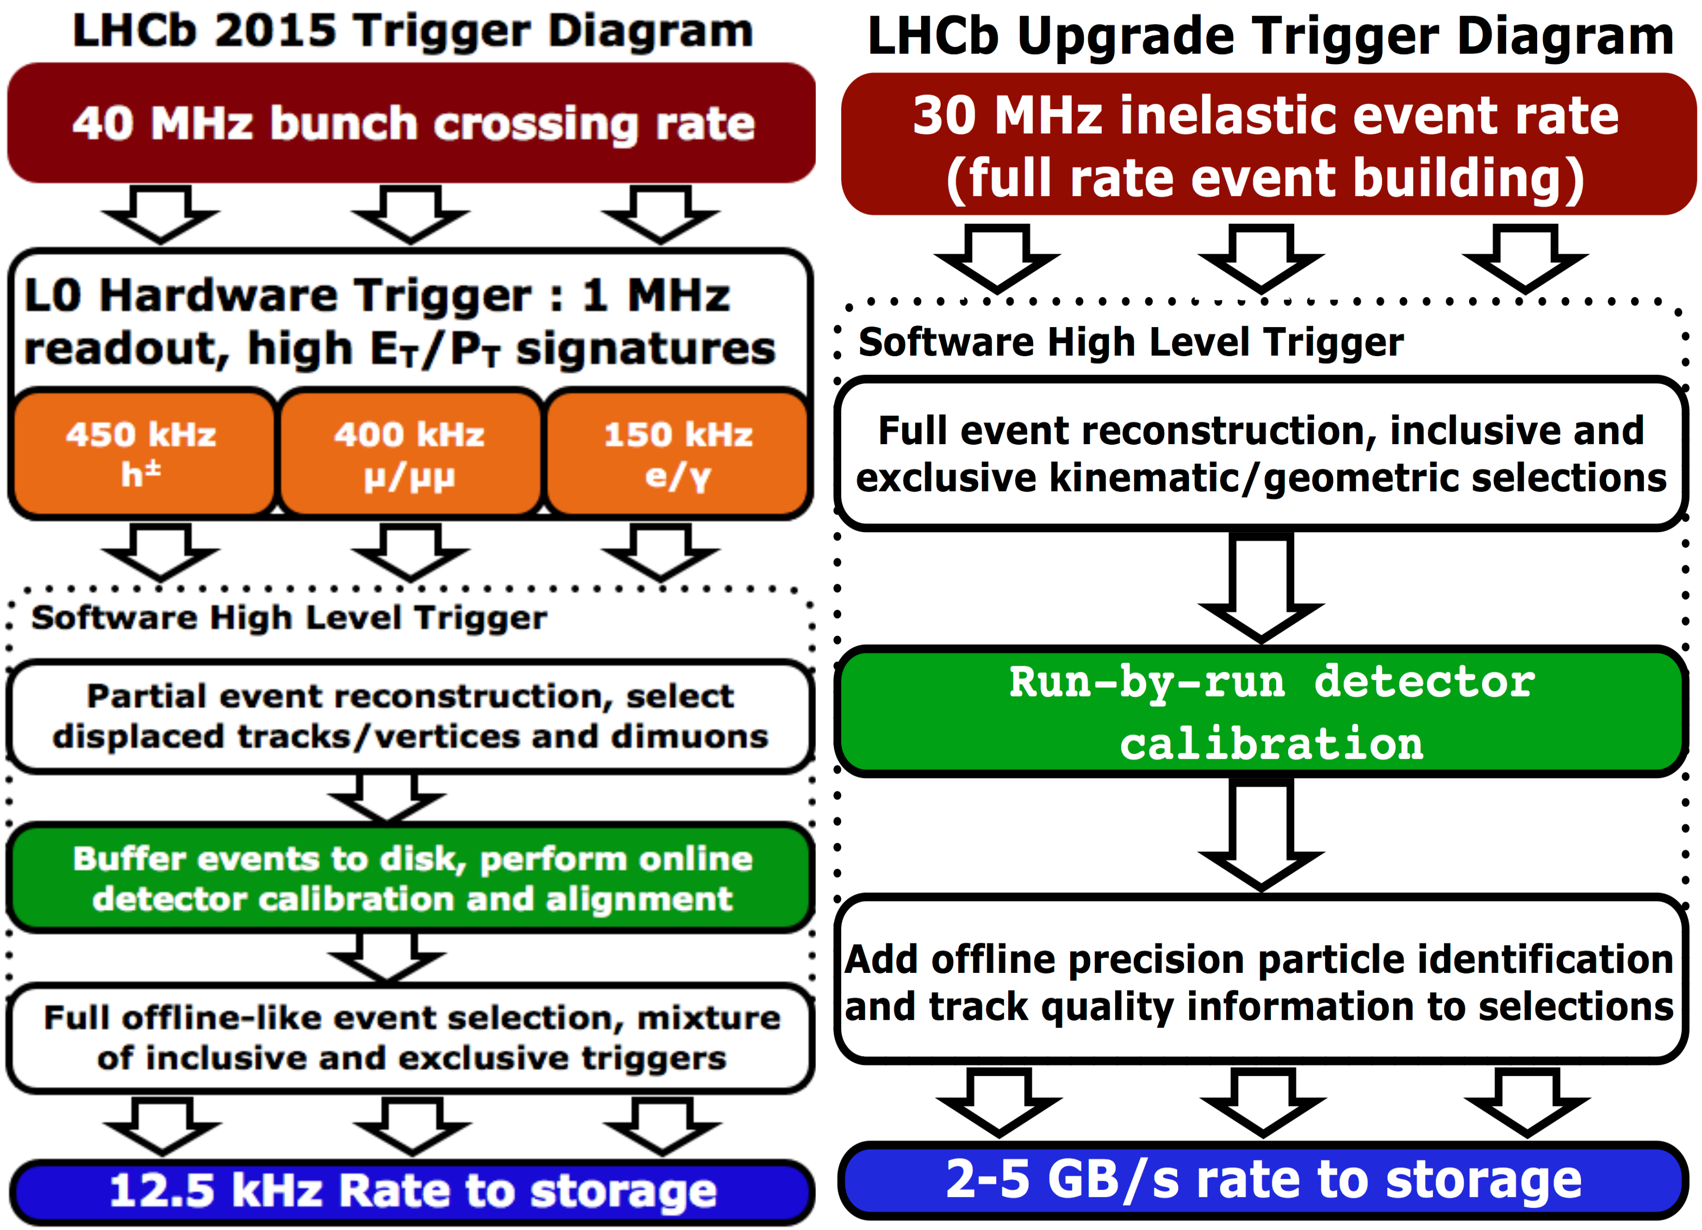

HLT 1 reconstruction strategy

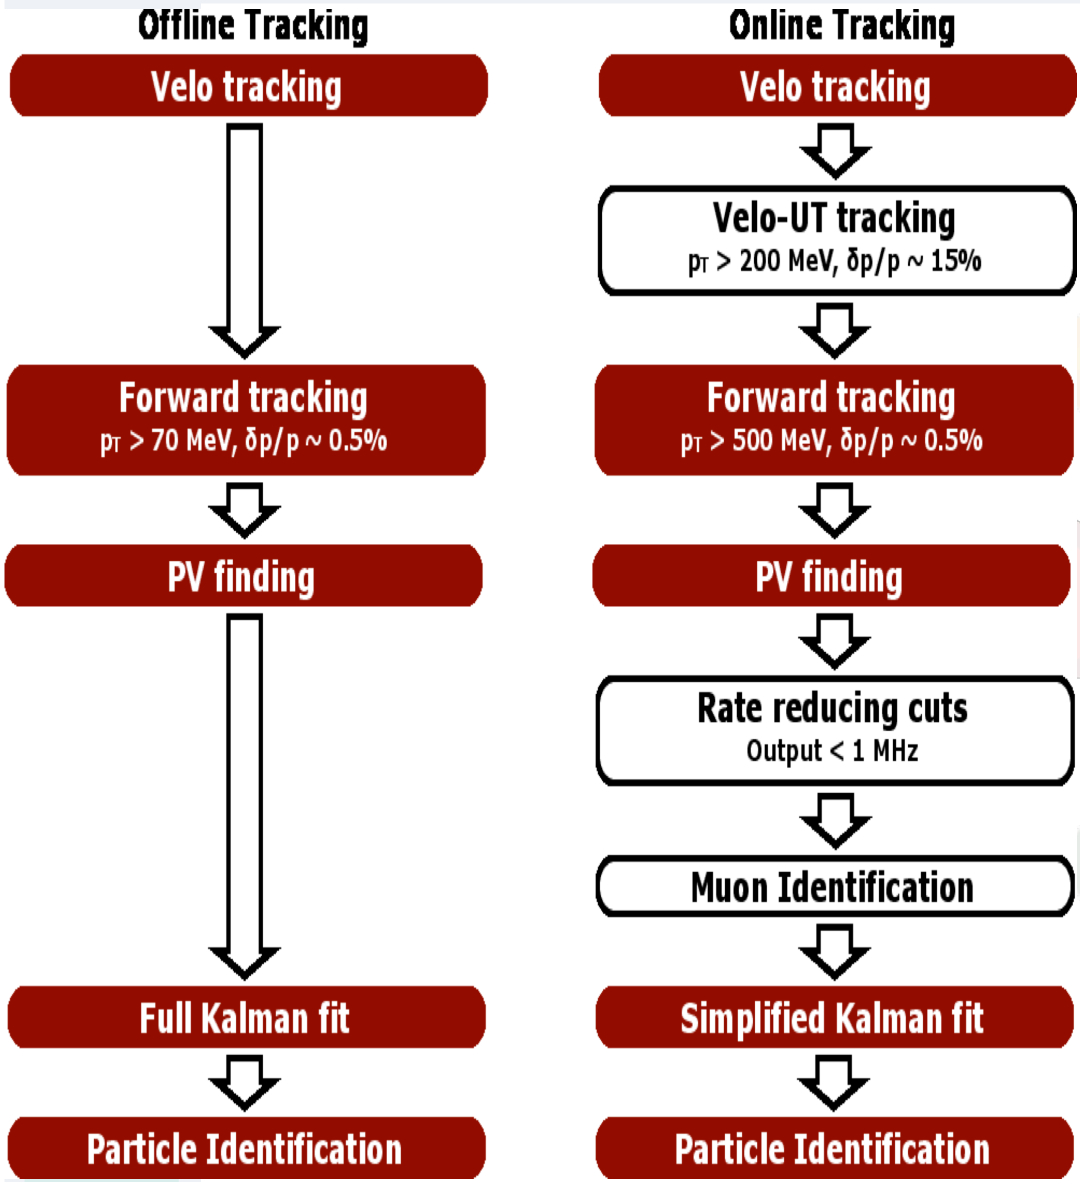

- With upgrade tracking detectors
  - Total HLT1 tracking time ~6 ms
- Offline quality tracking at 30 MHz is possible in software!
  - Concept already in use (at 1MHz) in Run 2
- LHCb first collider experiment to operate an all-software trigger at full event rate

Computing Model Brainstorming: Reconstruction

- After HLT 1, data buffered on local disk
  - Deferred processing after offline quality calibration has been obtained
    - Already planned for Run 2
- Aim to run full offline reconstruction in HLT 2
  - If reconstruction is identical to offline, is there a need to run it again offline?
  - And if reconstruction is not run offline, can we redefine RAW data to be ONLY the reconstruction output?
- CPU budget for reconstruction is crucial
  - Optimise for the HLT farm hardware
  - R&D for alternative architectures
    - Since reconstruction runs only in one place (HLT farm), can optimise hardware and software together
      - x86, but also GPU, Xeon Phi, ARM, …
    - Metric is events reconstructed per CHF
      - But remember code must also run on standard CPU (e.g. for offline simulation)

LHCb

Event selection: The Game Has Changed (© Tim Head)

- In the upgrade area there are no "boring" events, HLT is about classifying signal events!

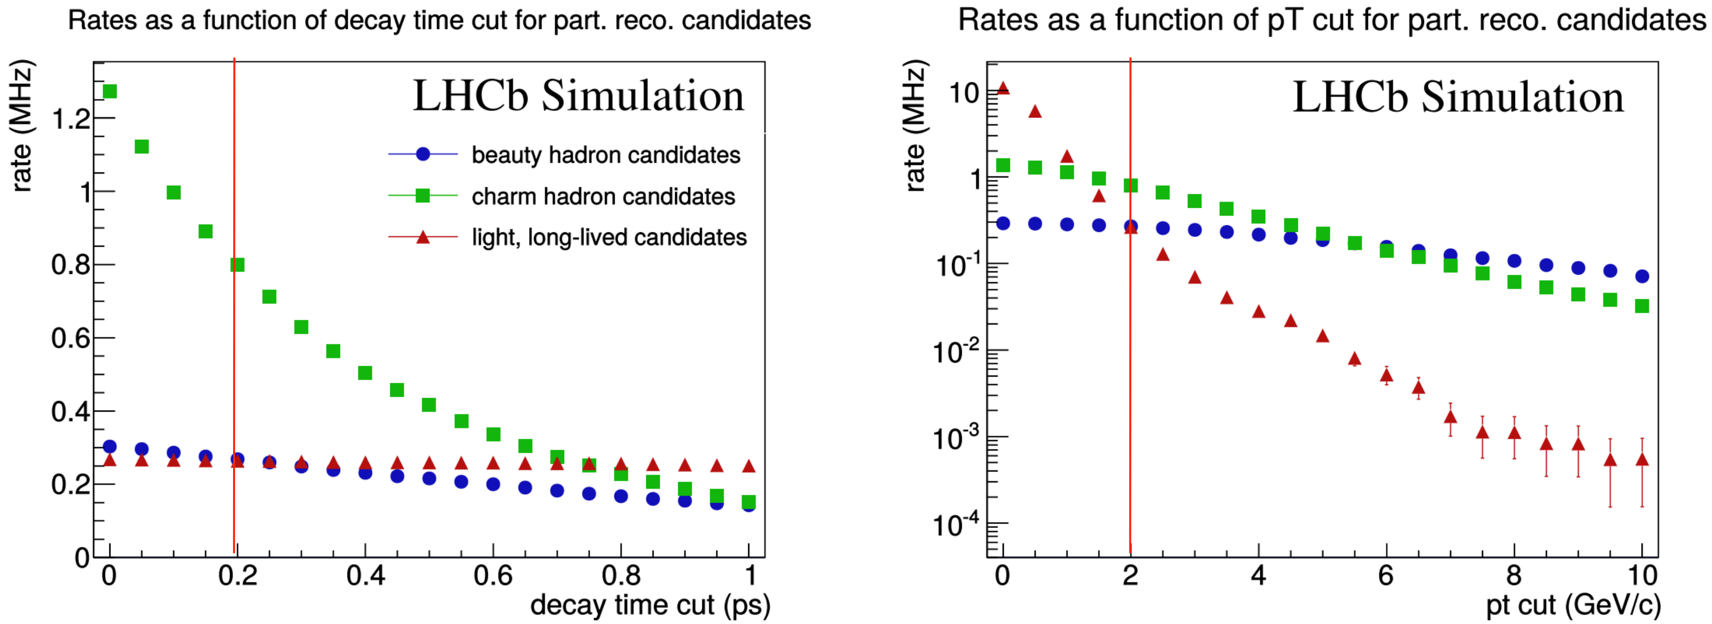

- 800 kHz of reconstructible charm hadrons, 270 kHz of reconstructible beauty hadrons
  - c.f. 2-5 GB/s allowed to offline storage

The Game Has Changed

Triggers today

Triggers in the future

83

(© Vava Gligorov)

Brainstorming: The Game Has Changed

- Trigger is no longer selecting events, but classifying them
  - Write out what bandwidth and offline resources allow, but everything written out will be analysed

- In many exclusive analyses, interested only in the decay tree of the triggering signal
  - So write out only the interesting part of the event, not the whole event
  - Turbo Stream idea, being tested in Run 2
    - 2.5kHz of signal, ~5kB/event, going directly to analysis
    - x10 reduction in event size

- If all events are interesting offline, current model of stripping no longer applies
  - Streaming more relevant, but how many streams?
  - Is direct access to individual events more relevant?
    - Needs event index, and R&D on efficient access to single events

Brainstorming: output event rate

- Design range for offline storage is HLT output rate of 2-5 GB/s
  - This represents huge range of event rates, depending on mix of Full events and Turbo events:
    - 2 GB/s of Full events (100kB) -> 20kHz
    - 5 GB/s of Turbo events (5kB) -> 1MHz
  - Reality will be somewhere in between, physics optimisation of bandwidth to be done
- Main implication for offline resources is CPU for simulation
  - Factor 50 in events to be simulated between two extremes above
    - This will be important ingredient of physics optimisation
  - Clear that major development effort in Fast MC techniques is needed
    - Already started for Run 2 physics, focusing on reducing needs for full simulation
      - Parametric approaches, partial event simulation

Brainstorming: storage model

- Three classes of storage
  - Disk for active data analysis
    - Real data, frequently accessed simulation
  - Active Tape
    - Less frequently accessed simulation
    - Migration between disk and tape based on popularity predictions
  - Archive Tape
    - Only for data, simulation and analysis preservation
      - No need for large disk cache
      - No I/O latency constraint, can be outsourced?
- A few sites for disk, even fewer for tape
  - ~3 sites with active tape sufficient
  - Sufficient disk sites to provide low latency and high availability for analysis jobs
    - No technical need for many small disk pools
      - (but recognise it as important funding/sociological issue)
- CPU can continue to be anywhere
  - Current WLCG distributed model
  - Leverage on opportunistic resources for simulation

Conclusions

- Three areas of research:

- Software optimisation
  - Use cases: reconstruction on HLT farm, fast simulation

- I/O optimisation
  - Use case: sparse event access

- Storage optimisation
  - Use cases: data popularity, archiving