From collisions to offline analysis
Stage 2 – Real Time Analysis

Marian Stahl

CERN

LHCb UK seminar

January 26, 2023

[LHCB-TDR-016]

Partially reconstructed signals

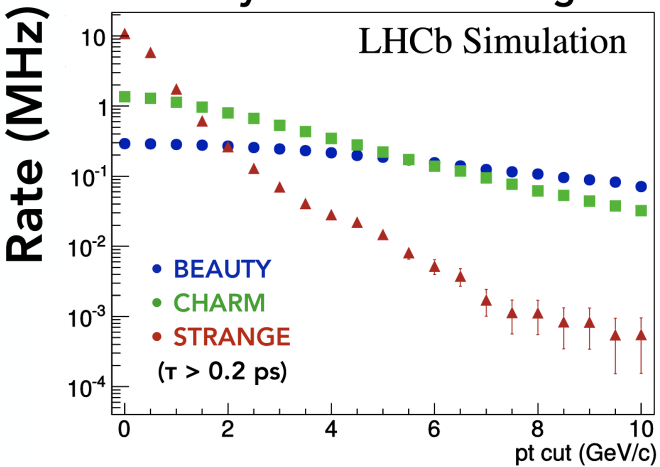

- There are $\mathcal{O}$ MHz of signal in the LHCb acceptance.
- We can only store $\sim$ 50 kHz of that if we store the raw event.
- Need flexible trigger, also because LHCb's physics program has vastly expanded.

LHCb's solution

[LHCb-DP-2021-003]

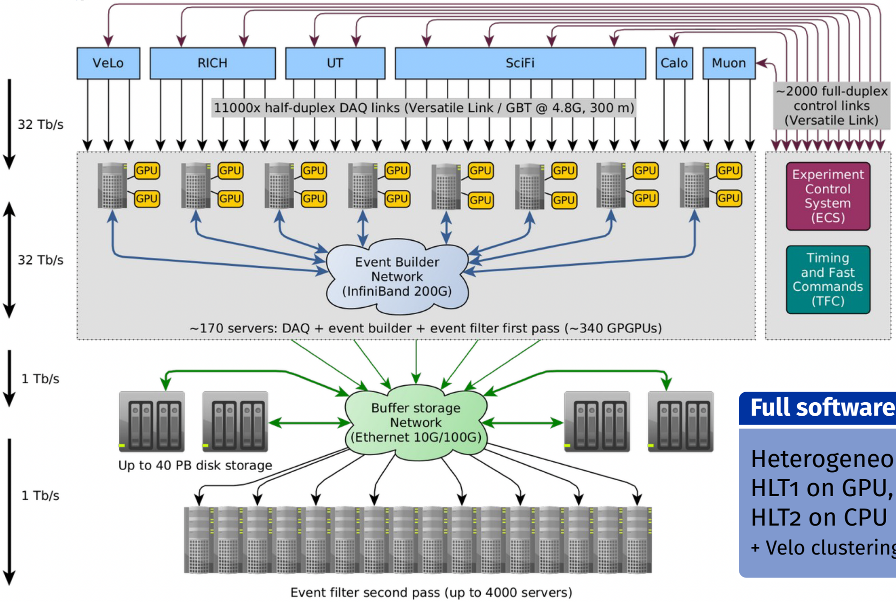

Full software trigger

Heterogeneous architecture
HLT1 on GPU,
HLT2 on CPU
+ Velo clustering on FPGA

[LHCb-FIGURE-2020-016]

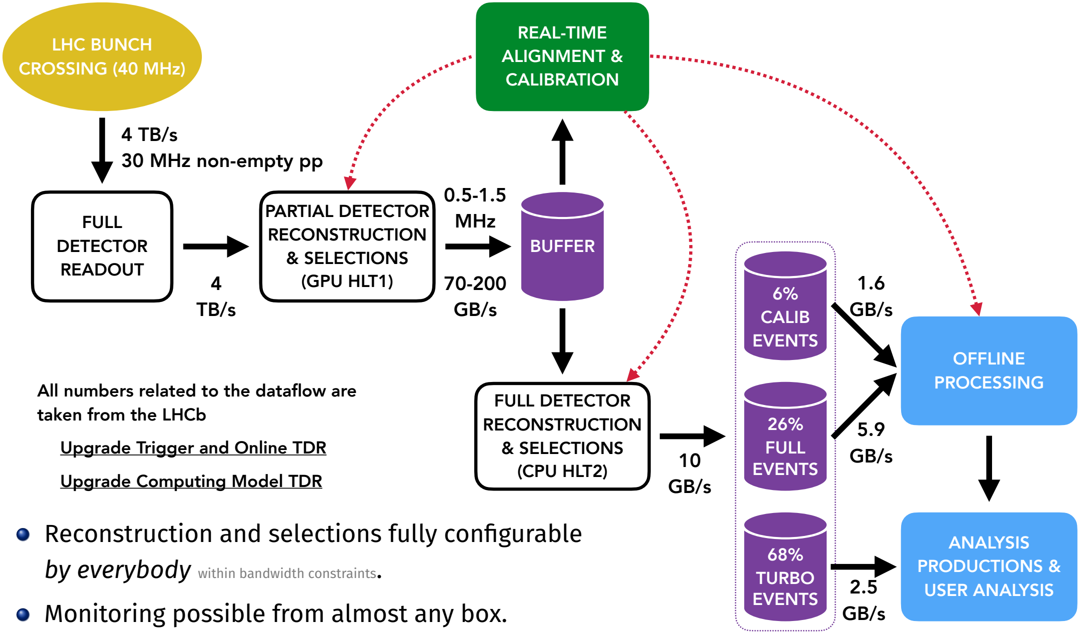

All numbers related to the dataflow are taken from the LHCb

Upgrade Trigger and Online TDR

Upgrade Computing Model TDR

- Reconstruction and selections fully configurable *by everybody* within bandwidth constraints.
- Monitoring possible from almost any box.

Offline Dataflow

[LHCb-FIGURE-2020-016]

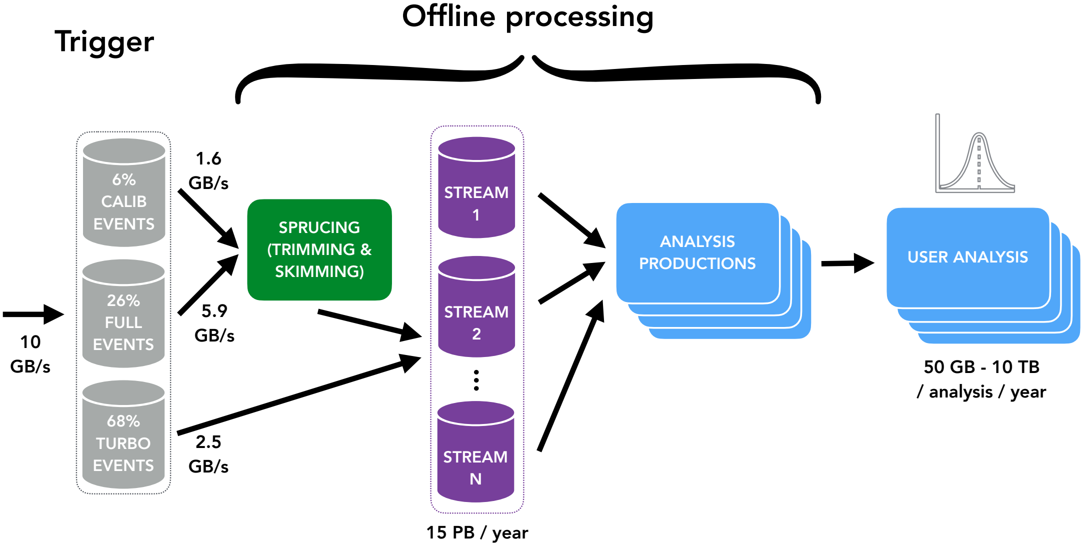

A lot of knobs to tune on this side too!

Dataflow details: HLT1

- The HLT1 application is called **Allen**, after Frances Allen.
- Allen runs standalone or within the LHCb software stack GaudiAllen; single- or multithreaded; compiled for GPU and CPU.
- Online, GPUs in the event builder (EB) nodes are running Allen standalone with Multi-Event-Packet (MEP) as input [EDMS 2100937].
    - NB: afraid we're missing "From collisions to offline analysis data - Stage 1.5" with everything happening between detector frontends and HLT1.
- Allen *transposes* MEPs to a detector-consolidated event-by-event format: MDF "Markus Data Format".

```
|14:00:31|n8190401|mstahl|~|$ pcie40_mdfreader /hlt2/objects/LHCb/0000256047/Run_0000256047_HLT22738_20221222-065830-220.mdf | less -FX -p ODIN
 Bank: 0x0002 (subsystem: 0 'ODIN', number: 2)
  Size: 48B (40B payload)
  Type: 16 'ODIN'
  Version: 7
   0x0000 | 2F E8 03 00 06 00 20 00 80 5F 96 2E 62 EE 05 00 00 00 00 00 9C 1F 00 00 0C 39 40 03 22 53 02 00
   0x0020 | 1E 9C 89 0A 00 00 00 00 |
  ODIN run_number: 256047 (0x3E82F)
  ODIN step_number: 32 (0x20)
  ODIN event_type: 6 (0x6)
  ODIN gps_time: Sat Nov 26 17:32:19 2022
  ODIN tck: 0x00000000
  ODIN partition_id: 8092 (0x1F9C)
  ODIN bunch_id: 2316 (0x90C)
  ODIN bx_type: 3 (0x3)
  ODIN nzs_mode: 0
  ODIN tae_central: 0
  ODIN tae_window: 0 (0x0)
  ODIN step_run_enable: 1
  ODIN trigger_type: 6 (0x6)
  ODIN tae_first: 0
  ODIN calib_type: 0 (0x0)
  ODIN orbit_id: 152354 (0x25322)
```

- The HLT1 application is called **Allen**, after Frances Allen.
- Allen runs standalone or within the LHCb software stack GaudiAllen; single- or multithreaded; compiled for GPU and CPU.
- Online, GPUs in the event builder (EB) nodes are running Allen standalone with Multi-Event-Packet (MEP) as input [EDMS 2100937].
  - NB: afraid we're missing "From collisions to offline analysis - Stage 1.5" with everything happening between detector frontends and HLT1.
- Allen *transposes* MEPs to a detector-consolidated event-by-event format: MDF "Markus Data Format".
- Minimal information is added as **Raw Banks**: Routing Bits, DecReports and SelReports.
  - Routing bits define how the data will be processed (physics, monitoring, calibration, errors ...).
  - DecReports encode *trigger line* decisions yes/no on selection `xyz`.
  - SelReports store (at least) *LHCbIDs* *e.g.* tracker hits, calo cellIDs... for each selected object.

[JINST 14 (2019) 04, P04006]

- HLT2 steered from **Moore**, hosting the configuration.
- Online: Run asynchronously on buffered data wait for alignment, free resources...
- HLT2 limited by *Bandwidth* = processing rate × event size.
  - "More physics" with faster code and/or less persisted information.
  - Eventually limited by disk/tape capacity.
- HLT2 adds `DstData` Raw Bank for reconstructed objects
  and updates/adds routing bits, Dec-/Sel-Reports.
- Offline, **Sprucing** passes through Turbo data, and filters streams written to tape; using same framework as HLT2.

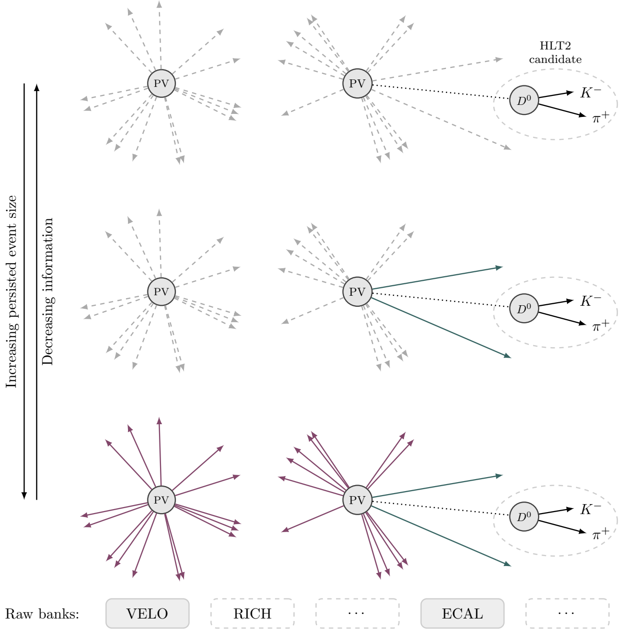

Intermezzo: Organization

https://twiki.cern.ch/twiki/bin/view/LHCb/RealTimeAnalysis

Connection

- Computing attaché.e
- Simulation attaché.e
- Online attaché.e
- Upgrade 2 attaché.e
- Run 1+2 WG attaché.e
- PWG Liaisons

Coordination

- PL & Deputy/ies
- IB chair (ex-officio)
- Work package coordinators
  - WP1 Data Structures
  - WP2 Reconstruction
  - WP3 Selections
  - WP4 Align & Calib
  - WP5 QA
  - WP6 R&D

Institutional Board

- IB chair
- PL & Deputy/ies (ex-officio)
- Institute representatives

Implementation

- WP deliverable responsibles
- Piquets
- Release shifters
- Voluntary developers and PWG line authors

	RTA/EMTF (former Run3 Perf)	RTA/DPA
QCD, Electroweak and Exotica	Kiera Farmer	Ross Hunter
B hadrons and Quarkonia	Jialu Wang	Yajing Wei
Charm physics	Dylan White (Manchester), Giulia Tuci (UCAS), Federica Oliva (Edinburgh)	Lorenzo Pica (Pisa), Ryunosuke O'Neil (Edinburgh), Felipe Luan Souza De Almeida (UFRJ Rio)
Rare decays	Felicia Volle (RTA), Hanae Tilquin (RTA)	Pere Gironella (DPA) Jiahui Zhuo
B decays to Charmonia	Xiaofan Hu	Ozlem Ozcelik
Charmless b-hadron decays	Juan Leite	Zewen Chen
B decays to Open Charm	Outgoing Fidan Suljik , Fionn Bishop Incoming Maciej Giza , Jonathan Davies	Aidan Wiederhold (DPA), Dong Ao (DPA), Alessandro Bertolin (RTA), Shunan Zhang (RTA)
Semileptonic decays	Tamaki Holly McGrath	Michael K. Wilkinson, Julián Garcia Pardinas, Suzanne Klaver
Ions and Fixed Target	Pasquale Di Nezza (EMTF)	Saverio Mariani (SMOG), Oscar Boente Garcia(PbPb), Samuel Belin(DPA)

Management	
Institutional Board Chair	Monica Pepe-Altarelli
Project Leaders	Carla Marin, Michel de Cian
WP Coordinators	
WP1 (Data Structures)	Chris Jones & Daniel Campora
WP2 (Reconstruction)	Dorothea vom Bruch & Maarten Van Veghel
WP3 (Selections)	Mika Vesterinen & Anton Poluektov
WP4 (Alignment & Calibration)	Sneha Malde & Marianna Fontana & Silvia Borghi
WP5 (Quality Assurance)	Rosen Matev & Roel Aaij
WP6 (Accelerators)	Conor Fitzpatrick
Attaché.e.s	
Computing	Ben Couturier
Simulation	Michal Kreps
Online	Patrick Robbe
Upgrade 2	Yasmine Amhis
Monitoring	Suzanne Klaver
Run1+2	Michael Alexander
Talks & Technical Paper Contacts	
Murilo Rangel	
Xavier Vilasis Cardona	

LHCb Software projects

- Structure of LHCb software **stack**. Not all projects shown. Framework, Component library and Applications would say that Allen, Gauss, Boole are mixtures.
- Use GitLab service to host projects; deployed on `cvmfs`.
- Most developers use lb-stack-setup; quick developments with `lb-run`, `lb-dev`.

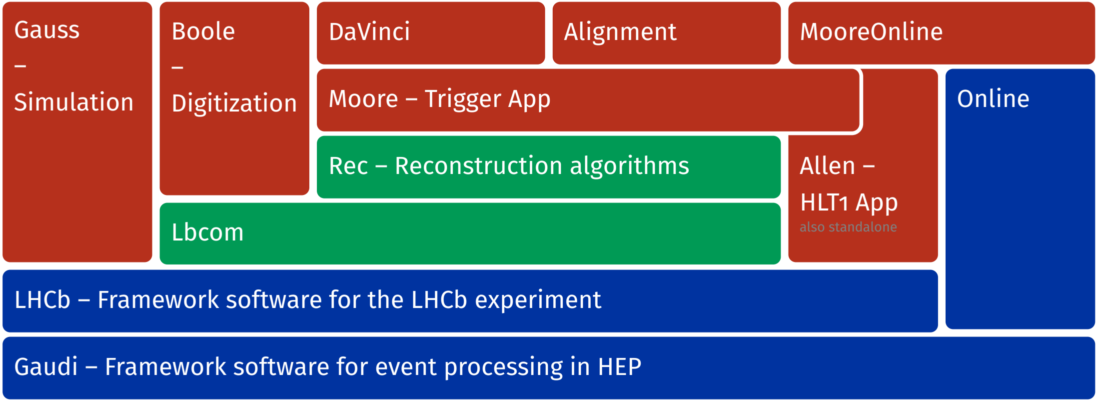

Track types

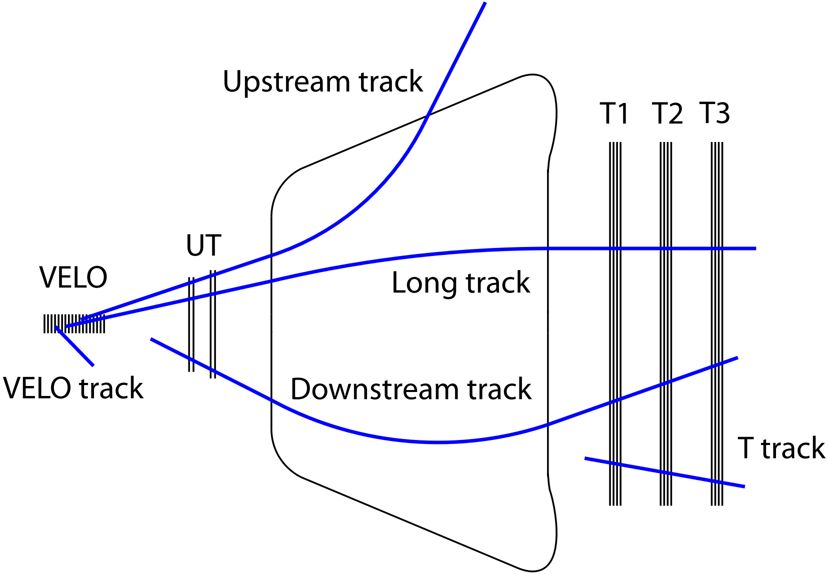
Upstream track
T1
T2
T3
UT
VELO
Long track
VELO track
Downstream track
T track

[LHCB-TDR-021]

- HLT1 selections based on tracking and lepton ID.
  - By now also ECAL reco.
  - 2nd GPU to be installed
    ⇝ remove GEC, downstream-tracking, loosen reconstruction...?
- HLT2: offline reconstruction quality, dedicated reconstruction algorithms diverse selections
  i.e. lots of branching.
  Still, mostly vectorized (SIMD).

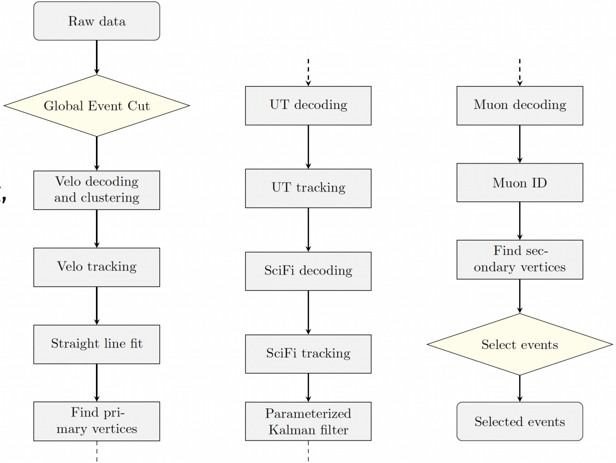

Tip

Use IDEs, like VSCode, and follow the configuration; e.g. HLT1, HLT2 to understand which algorithms run under what conditions. When developing, it's often helpful to search existing tests and follow them.

- Limited bandwidth $\rightsquigarrow$ compact data format.
  - Data contains **encoded** information of hits, the event, and other information.
  - Further reduction by optional zero-suppression and/or clustering before HLT1.
- Example: UT hit encoding. 64 bit header defines *payload*. Here: hits grouped in lanes, each hit encodes (local) stripID in 11 bits, ADC value in 5 bits. For comparison: `int` has at least 16 bits, usually 32.

one event
one TELL40

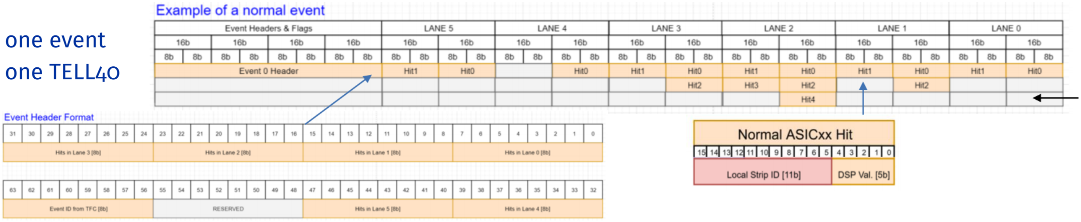

- Decoding translates local information to global reference frame using raw-bank data, **detector description** (geometry), and **conditions** (alignment).
- Each detector with specialized format, but similar concept.

Velo tracking

- Challenge: combinatorics with $\mathcal{O}(100)$ hits per plane (P); $\exists$ 26 planes, 2 modules each.
- GPU and CPU algorithms similar. Operates on a 4 plane sliding window.
- Seeding step: build hit triplets (use φ proximity).
- Extend step: Search hit aligned with previous two hits. Stop if no hit found in 2 consecutive planes.

[J.Comput.Sci. 54 (2021) 101422]

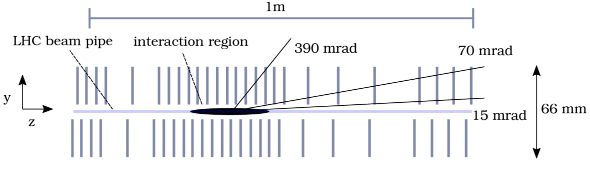

VELO detector composed of 52 modules

[talk by Arthur Hennequin]

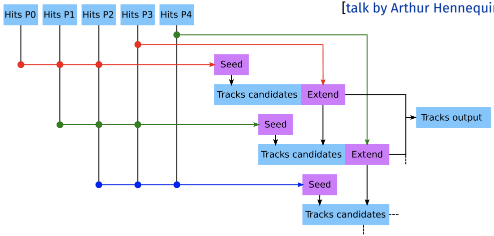

[JINST 15 (2020) 06, P06018]

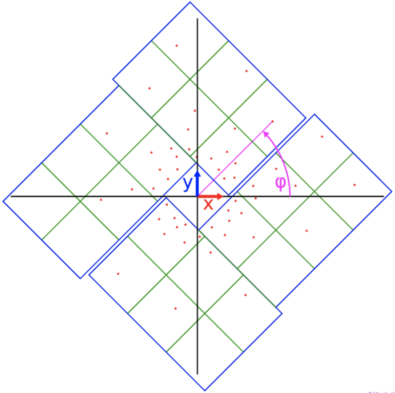

- Straight line fit of output tracks.

[LHCb-FIGURE-2020-005]

- Use fitted Velo tracks (forward and backward).
- Fill 1D histogram with track density along beamline.
- Peak finding gives PV seed position.
- Vertex fit gets actual position.

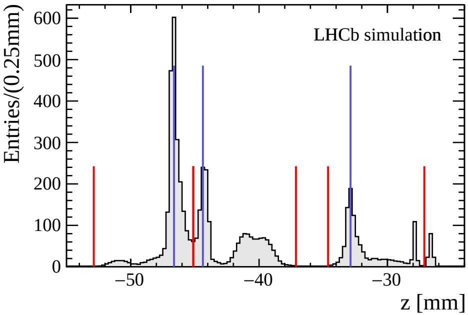

- Will nominally run at 5 times higher PV multiplicity in Run 3. PV mis-association a problem?
- Have 3D PV finder as backup.

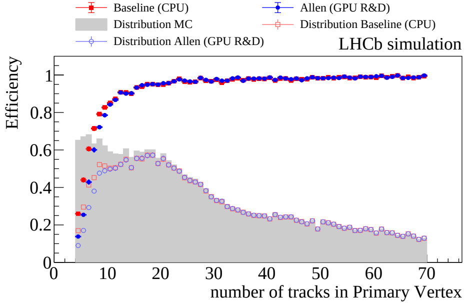

Long tracks

- Two algorithms: Forward and Seeding & Matching.
- Either or in HLT1; HLT2: first Seeding & Matching, then Forward on residual Velo tracks and SciFi hits.
  - Similar combination also under study for HLT1.
  - Can easily switch between configurations; Flexibility for running with and without UT.

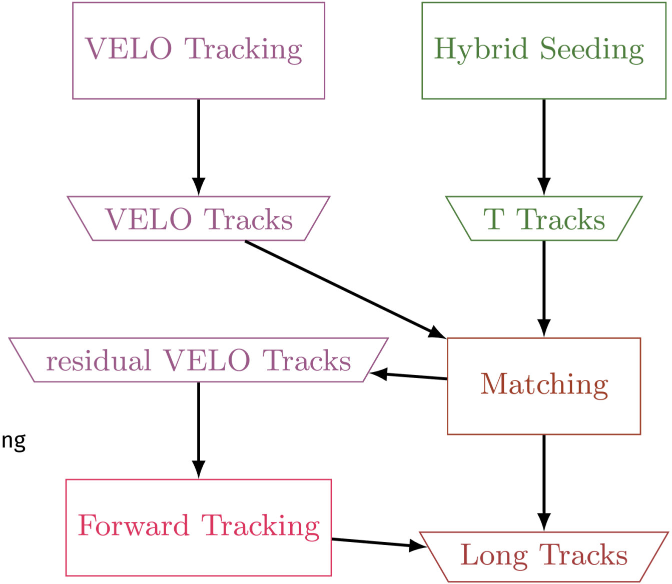

Forward tracking

- Start from Velo(-UT) seed;
  define hit search window.
- Treat magnet as optical lens
  to simplify track and hit projection.
- Hough-like transform:
  project hits to reference plane;
  count SciFi layers in histogram.
- Polynomial track fit;
  get momentum estimate.

[LHCb-TALK-2022-191]

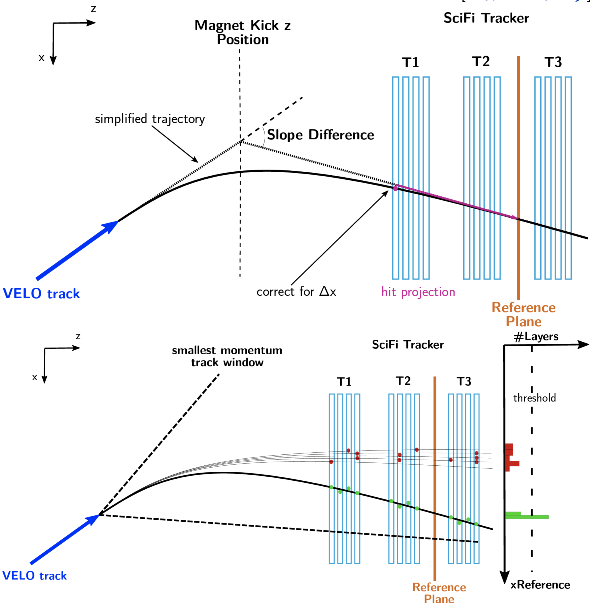

[LHCb-TALK-2022-206] see also [Comput.Phys.Commun. 260 (2021) 107713]

- Search hit doublets in two of four *x*-layers in SciFi's T1 and T3, assuming minimal momentum and track origin at $(0, 0, 0)$.
- Look for T2 *x* hits using momentum estimate from the 2 hits and $(0, 0, 0)$-constraint.

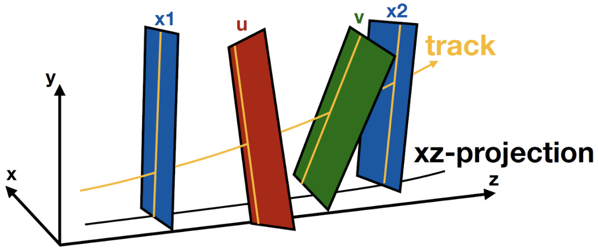

one of three T-stations

- Get remaining two *x* hits from parabolic fit; remove outliers; flag hits used for next iteration.
- Run three iterations relaxing momentum search window and alternating *x*-layers.
- Hits in SciFi stereo layers added by search for Hough clusters.
- Final recovery step for downstream-, low-momentum tracks; fitting and clone removal.
- Matching to Velo- or upstream-track parametric (HLT1) or assisted by neural network (HLT2).
- Downstream tracks from extrapolating T-tracks to UT.

- **Cherenkov light** emitted along particle trajectory, if particle faster than speed of light in traversed medium.
- Mirror geometry in RICH such that Cherenkov light forms **ring**.
- For each track state and type hypothesis, the RICH reconstruction generates & traces Cherenkov photons and simulates their detector response.
- Start with pion hypothesis compute global likelihood; change hypothesis iteratively.

[talk by Giovanni Cavallero on 2022 commissioning]

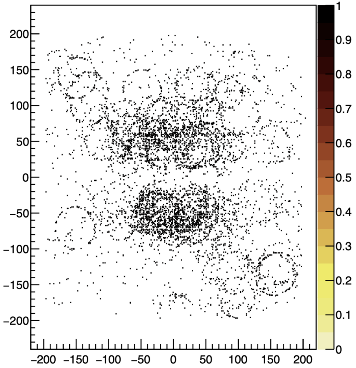

⟶ Finally, compute likelihood differences for each particle.

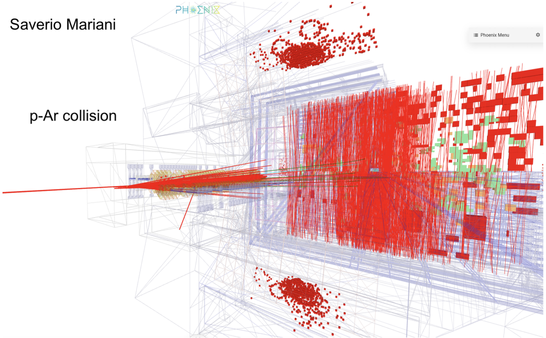

ECAL reconstruction

- Photons and electrons shower in ECAL. Their E_{deposit}/p is close to 1; unlike hadrons.
- Electrons are IDed as tracks pointing to ECAL cluster. Recover Bremsstrahlung.
- Neutral reconstruction by clustering energy deposits (**digits**) in calorimeter cells.
- Start with seed cell: local maximum E_{deposit} above a configurable threshold in 3×3 cell window.
- Represent digits as graph: Seed is sink, bi-directional nodes are shared between clusters.
- Cluster-shape used to identify merged $\pi^0/\eta\ldots$ and isolated γ.

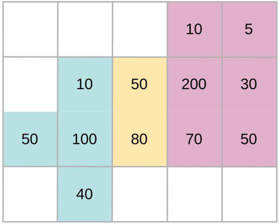

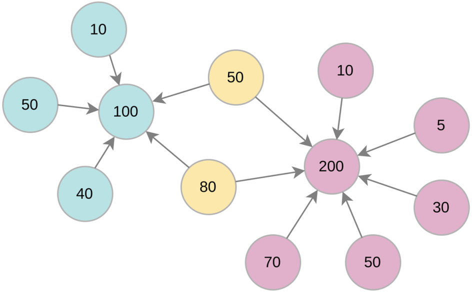

- Define IsMuon as coincidence of hits in M2 to M5 in x-y field around track extrapolation.
  - Require 2, 3, 4 hits depending on track momentum 3 to 6, 6 to 10, above 10 GeV.
- New Muon ID variable: correlated sum of spatial residuals of hits $\delta\boldsymbol{x}$, $\delta\boldsymbol{y}$ w.r.t. extrapolation

$$\chi^2_{\text{corr}} = \delta\boldsymbol{x}^T \boldsymbol{V}_x \delta\boldsymbol{x} + \delta\boldsymbol{y}^T \boldsymbol{V}_y \delta\boldsymbol{y}.$$

  - Multiple scattering in calorimeters and absorbers, as well as hit correlation taken into account in off-diagonal elements of $\boldsymbol{V}$.

[LHCb-DP-2020-002]

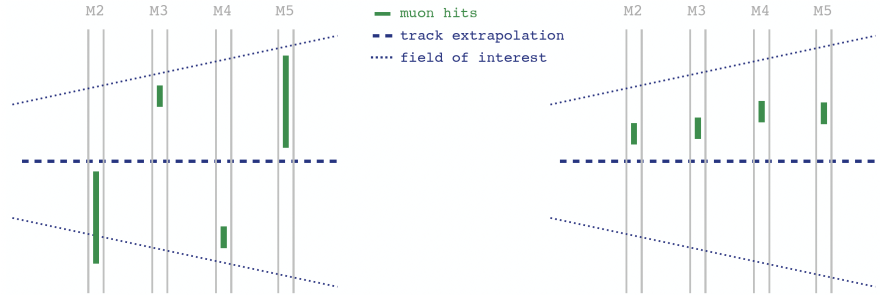

- Determine most accurate estimates of track parameters and covariances.
- Track as collection of lines tangent to trajectory: **State** $\vec{x} = (x, y, t_x, t_y, q/p)^T$ and its covariance.
- State + measurement in detector: node k used in **Kalman filter**.
- Measurements added node-by-node; Minimize χ^2 of measurements on track in 3 steps:
  - **Predict** state in next node $\vec{x}_k^{k-1}$ from *filtered* state $\vec{x}_{k-1}$; add noise Q_k to covariance.
  - **Filter** step adds information of measurement m_k.
  - **Smoother** gives best estimates of states at all previous nodes. Result is track χ^2.

[CERN-THESIS-2005-040]

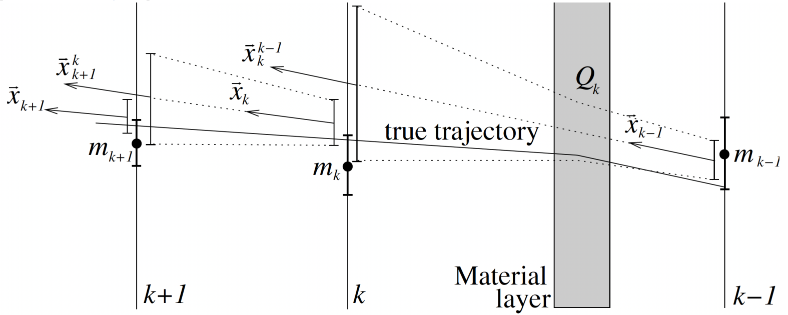

[LHCb-TALK-2022-299]

- Detectors not *exactly* where expected by description database. Can also be rotated/tilted or moving with time.
- **Data-driven** alignment corrects for that in quasi-real-time.
- Tracker Velo, UT, SciFi, Muon alignment finds at most 3 constants for translations $T_{x,y,z}$, 3 for rotations $R_{x,y,z}$ for each "alignable" element; exported to database.
- Insert alignment constants in track fit; iterate smoothing step with global covariance of track states [NIMA 600, 471]; use constraints [NIMA 712 48]

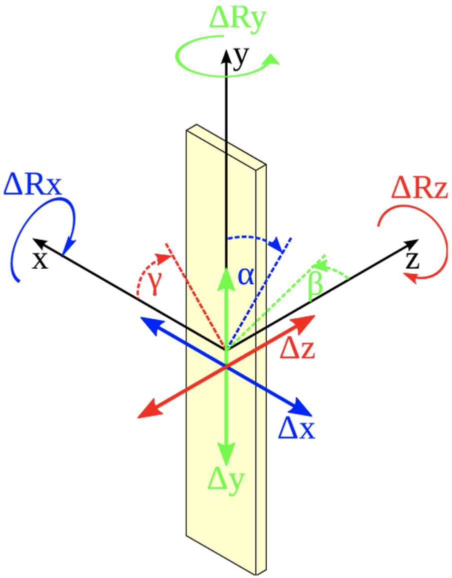

Alignment and calibration – Velo

- Velo closes once LHC declares STABLE BEAM.
- Align translation constants of Velo halves.
- Check with PVs reconstructed independently in both halves.

[LHCb-FIGURE-2019-003]

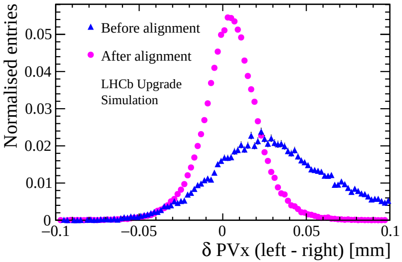

https://lbcomet.cern.ch/Online

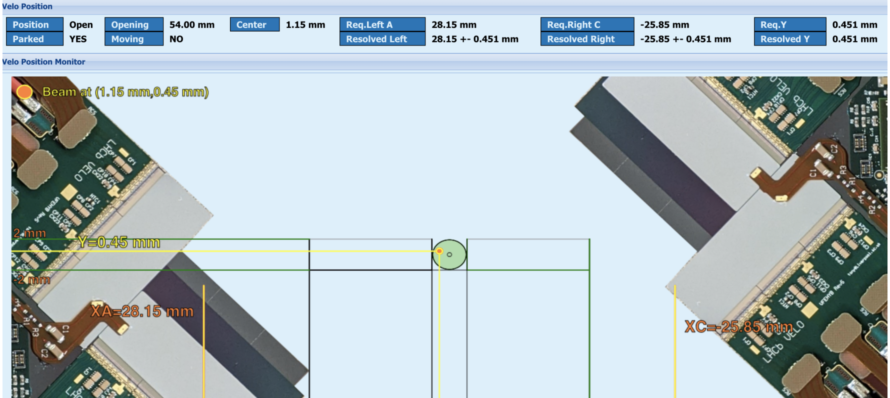

- Different detector calibrations/alignments happen at different timescales.
- Dedicated streaming and resources online.

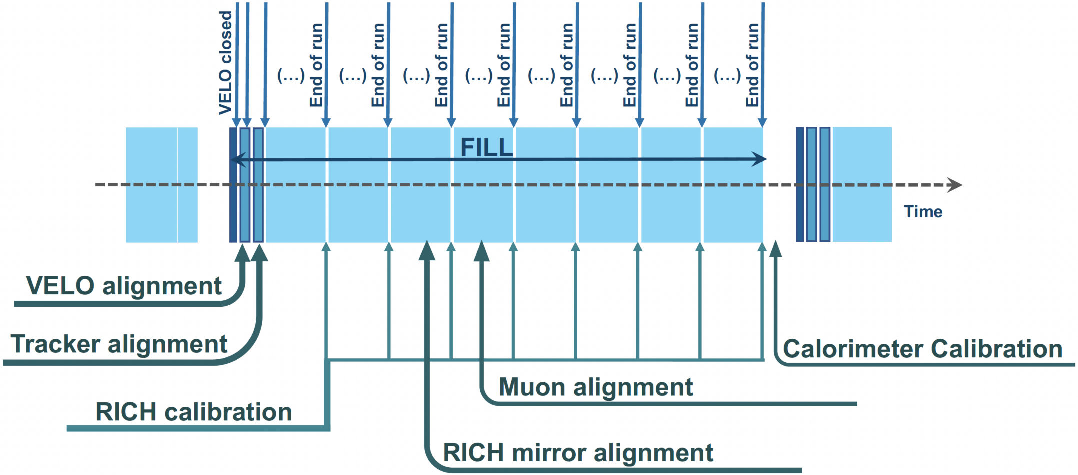

Selections

- Bandwidth constrains impose tighter and more exclusive selections $\rightsquigarrow$ Flexibility is key.
- Much branching in selections $\rightsquigarrow$ especially challenging on GPUs.
  - Allen builds all combinations, filters later; Moore usually faster when breaking early.
- Currently $\sim$ 40 HLT1 physics lines, $\sim$ 1500 in HLT2.
  - Much of *b* physics covered by `Hlt1(Two)TrackMVA` and Topo lines;
    They use novel monotonic neural net [arXiv:2112.00038]
- Throughput constraints don't allow track and vertex fit in HLT1;
  HLT2 uses `PrKalmanFilter` and `ParticleVertexFit` [Wouter's note].
- HLT1 selections CUDA-based; HLT2 uses python interface to underlying SIMD (ThOr-) Functors and Combiners in C++.
- Consult Allen, Moore docs, current hackathon, starterkit etc. for writing your selection.
- Technically reduce bandwidth by packing and optimized streaming.
- However, most optimization clean selection, custom persistency has to come from analysts discussion at A&S week.
  - Great way to contribute to 2023 commissioning.

- Upgraded SMOG2 System for Measuring the Overlap with Gas will allow LHCb to record fixed-target collisions concurrently with $p - p$ physics.
- LHCb holds record as highest-energy fixed-target experiment.

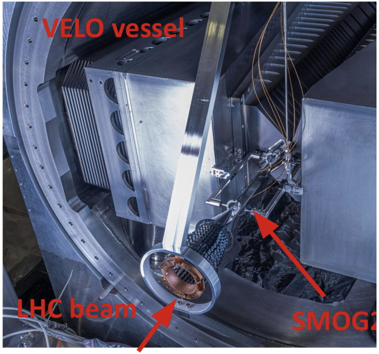

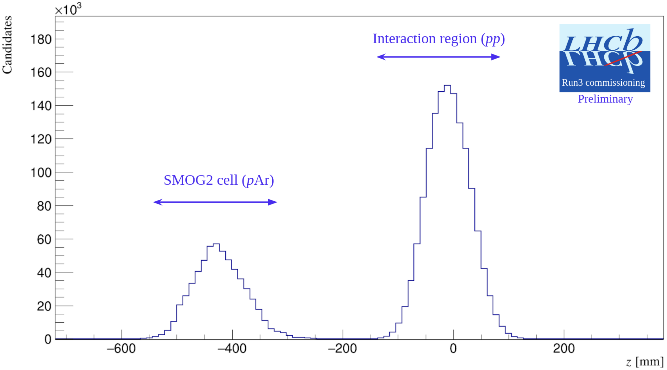

- LHCb participates in LHC's heavy ion program: *p*Pb, Pb*p* and PbPb configurations.
- Requires dedicated configuration of the software.
- High occupancy in the detector; Only 2 fills in 2022. Much more expected this year.

[talk by Benjamin Audurier]

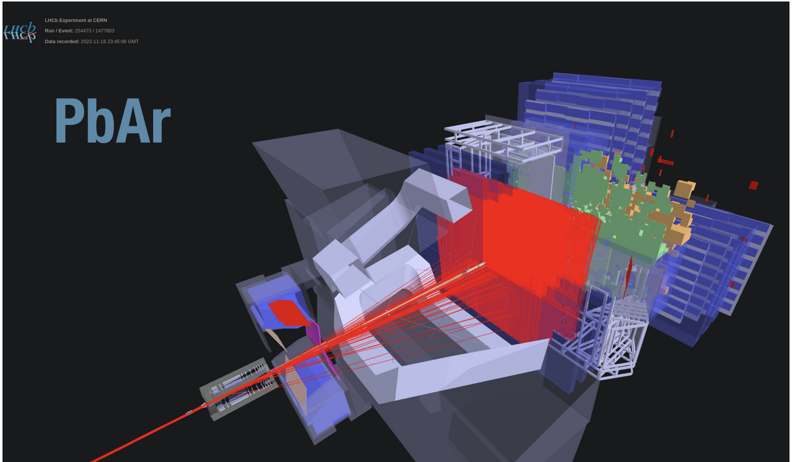

QMTests and nightly builds

- Basis of LHCb’s test driven development are Q(uality)M(anagment)Tests, *e.g.* run PV reco on reference simulation sample and compare (efficiency) counters.
- Each merge request tested against stack; On demand webhooks in GitLab or in nightlies.

Marian Stahl @mstahl · 1 week ago — Developer

/ci-test --merge Rec!1869

1

Lhcb Software @lhcbsoft · 1 week ago

Started reference and integration test builds. Once done, check the comparison of build and test results.

master vs. master+this MR+Rec!1869

lhcb-master-ref - build: 297 (2020-01-10) --> lhcb-master-mr - build: 335 (2020-01-10)

Project	Version	x86_64-centos7-gcc9-opt		x86_64-centos7-gcc9-dbg		x86_64-centos7-clang8-opt		x86_64-centos7-clang8-dbg	
DBASE	None	OK		OK		OK		OK	
PARAM	None	OK		OK		OK		OK	
Gaudi	master	OK	249	OK	249	OK	249 (4)	OK	249 (4)
Online	master	OK	68 (2)	OK	68 (2)	15	68 (2) +1	37	68 (3)
LHCb	master	OK	247	OK	247	OK	247	OK	247
Lbcom	master	OK	1	OK	1	OK	1	OK	1
Boole	master	OK	14	OK	14	OK	14	OK	14
Rec	master	OK	33 +1	OK	33	OK	33	OK	33

https://lhcb-nightlies.web.cern.ch/nightly

lhcb-gaudi-head - build: 2508 (2020-01-20)

head of everything against Gaudi/HEAD and LCG_96b

available on: **cvmfs**

Compare with previous build | Compare with other slot | Browse files

Project	Version	x86_64-centos7-gcc9-opt (Completed at 05:01:49)		x86_64-centos7-gcc9-dbg (Completed at 07:24:59)		x86_64-centos7-clang8-opt (Completed at 04:56:28)		x86_64-centos7-clang8-dbg (Completed at 05:56:44)		x86_64-centos7-gcc9-do0 (Completed at 05:16:20)		x86_64+avx2+fma-centos7-gcc9-opt (Completed at 05:05:34)		skylake_avx512+vecwid256-centos7-gcc9-opt (Completed at 02:46:14)	
PARAM	None	build		build		build		build		build		build		build	
DBASE	None	build		build		build		build		build		build		build	
LCG	96b														
Gaudi	HEAD	build	tests	build	tests	build	tests (4)	build	tests (4)	build	tests	build	tests (3)	build	tests (3)
Online	HEAD	build	tests	build	tests	build	tests (1)	build	tests	build	tests	build	tests (1)	build	tests (2)
LHCb	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests (6)	build	tests (7)
Lbcom	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
Boole	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests (6)	build	tests (6)
Rec	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests (2)
Brunel	HEAD	build	tests (1)	build	tests	build	tests	build	tests	build	tests (10)	build	tests (13)	build	tests (13)
Phys	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests
Moore	HEAD	build	tests (2)	build	tests (2)	build	tests (2)	build	tests (3)	build	tests (6)	build	tests (4)	build	tests (14)
Analysis	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests (1)	build	tests	build	tests
Stripping	HEAD	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)	build	tests (1)
DaVinci	HEAD	build	tests	build	tests	build	tests	build	tests (2)	build	tests (2)	build	tests (1)	build	tests (1)
Panoramix	HEAD	build	tests	build	tests	build (6)	tests (0)	build (6)	tests (0)	build	tests	build	tests	build	tests
Bender	HEAD	build	tests (3)	build	tests (3)	build	tests (3)	build	tests (4)	build	tests (6)	build	tests (3)	build	[illegible]
MooreOnline	HEAD	build	tests (1)	build	tests (1)	build (2)	tests (1)	build (2)	tests (1)	build	tests (1)	build	tests (1)	build	[illegible]
Panoptes	HEAD	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests	build	tests

- Several flavors of nightly builds platform, git branches, compiler like clang sanitizer etc..
- Periodic weekly **Performance- and regression** (PR) tests on larger input samples.

Periodic tests

https://lhcb-nightlies.web.cern.ch/periodic/, details: EPJ WoC 214, 05014, EPJ WoC 214, 05042

- Scheduled or on demand tests in controlled conditions.
- Have throughput, rate/bandwidth or physics performance tests very close to online plots – Monet

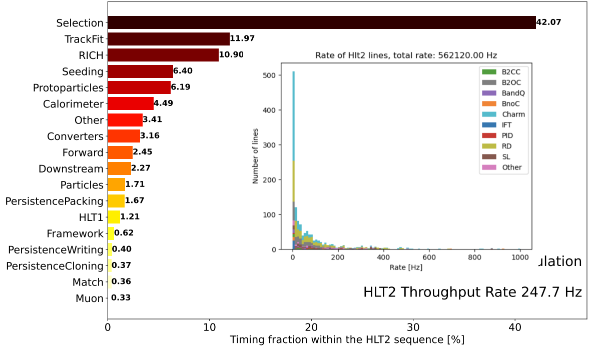

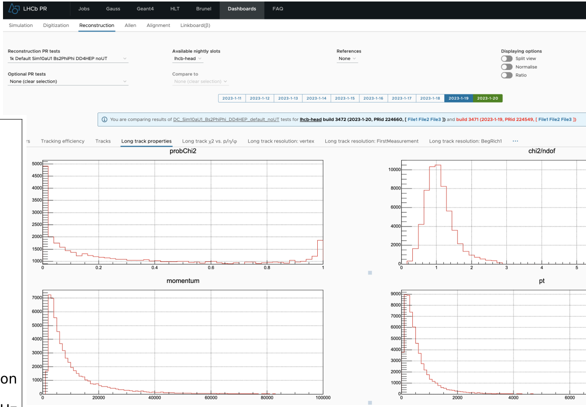

Online monitoring

http://lbwebmonet.cern.ch/

- Dedicated monitoring farm. Tasks run their own sequence even their own stack.
- Usually ∼ 10 kHz of raw data; Histograms saved every 10 mins in root files (Savesets).
- Flexibility run dedicated stream of events select what you want to monitor.

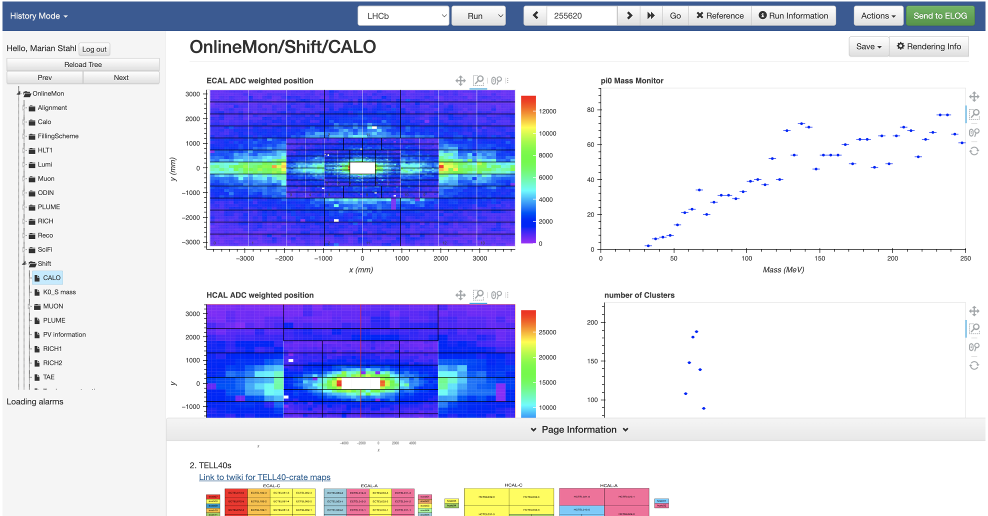

Summary

- With the upgraded detector, LHCb faces exceptional technical challenge.
- RTA responds with full software trigger on heterogeneous architectures to provide flexibility needed to cover increasingly broad physics program.
- There is a big overlap with detector, simulation, computing, online and offline projects.
- This year's commissioning will be a fun and rewarding challenge.

[talk by Vava Gligorov]

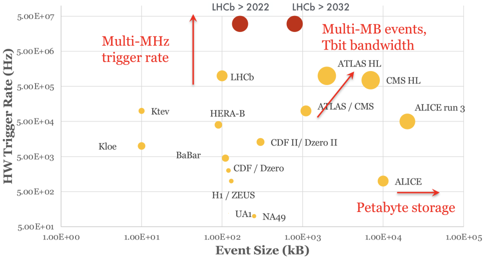

I look forward to meeting many of you again in the control room!

Support material

LHCb detector (upgrade)