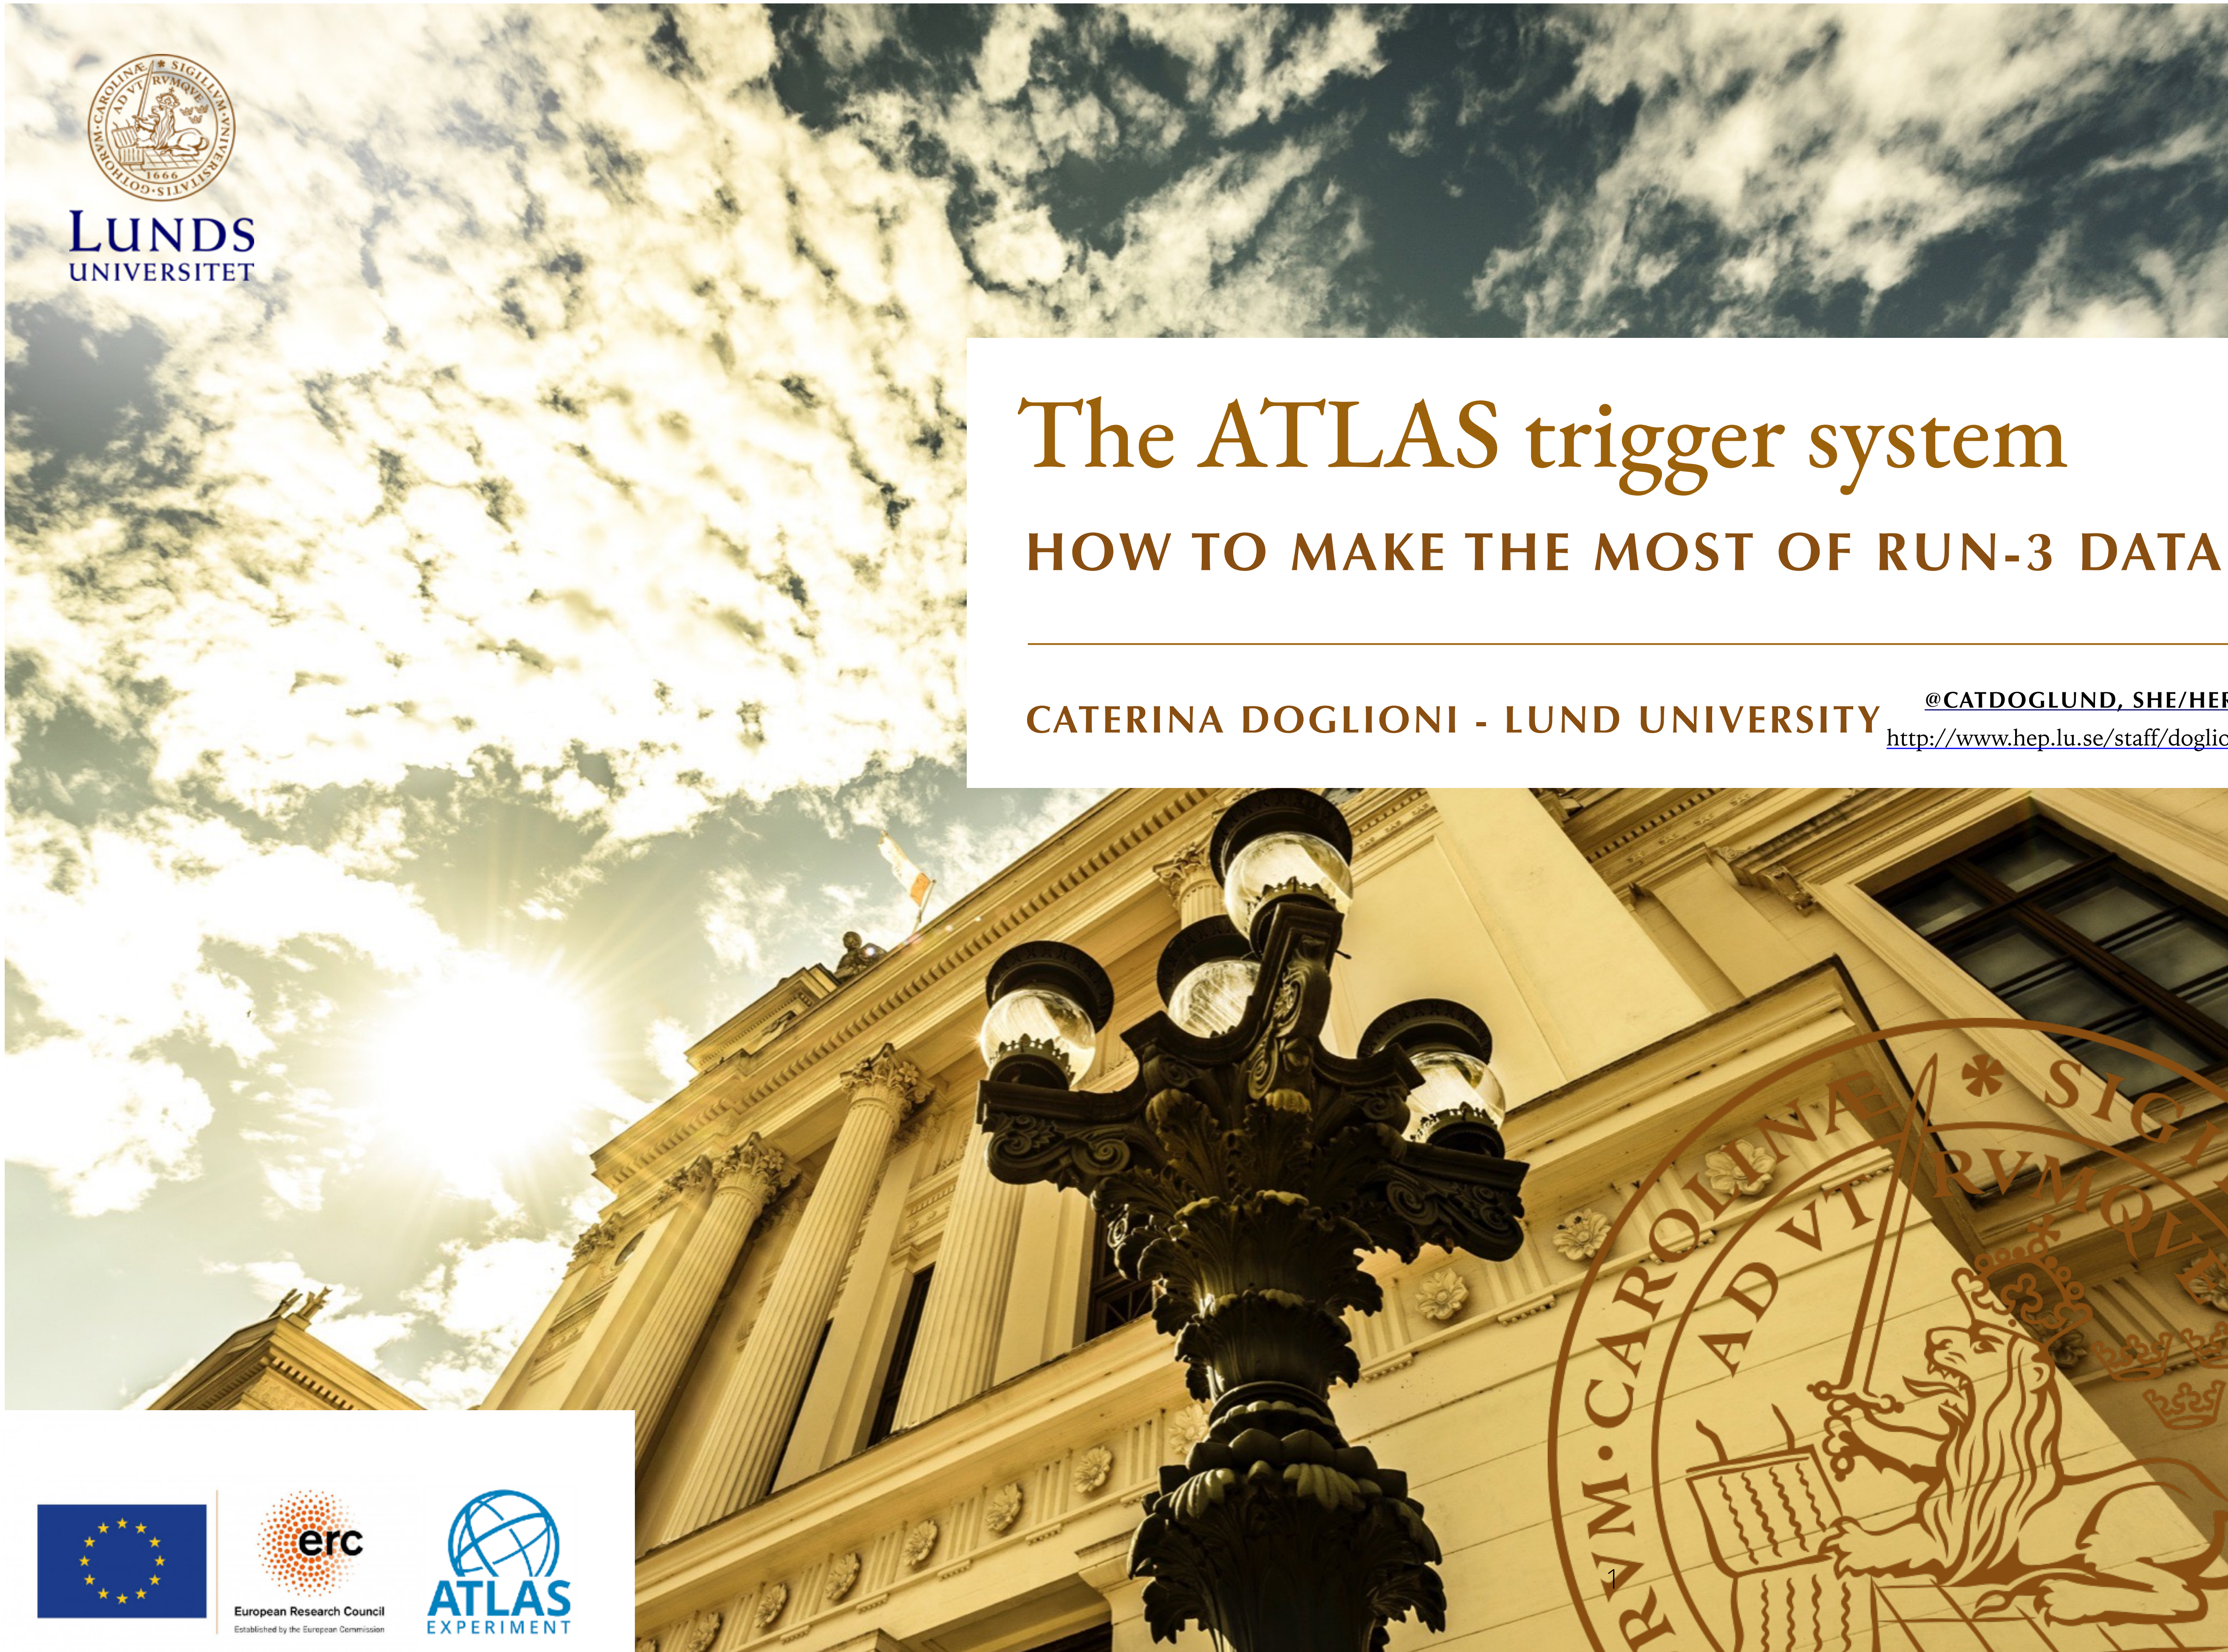

LUNDS UNIVERSITET

The ATLAS trigger system

HOW TO MAKE THE MOST OF RUN-3 DATA

CATERINA DOGLIONI - LUND UNIVERSITY

@CATDOGLUND, SHE/HER

http://www.hep.lu.se/staff/doglioni/

European Research Council
Established by the European Commission

ATLAS EXPERIMENT

Outline

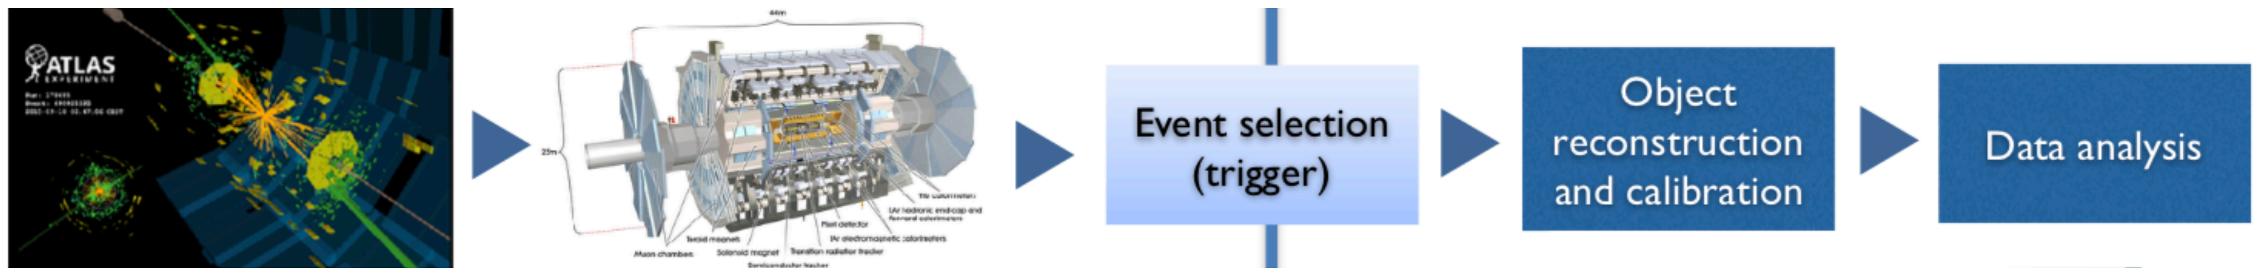

- A quick intro to triggering
- Trigger hardware improvements for Run-3
- Software trigger improvements
- Non-standard workflows

This talk: trigger tools to get more physics out of Run-3
Later talks: what we can measure/discover with Run-3 trigger system

Triggering: a rapid intro

Online ◀ ▶ Offline

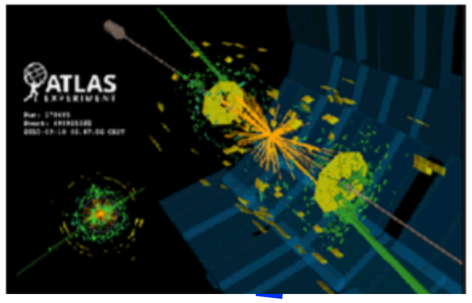

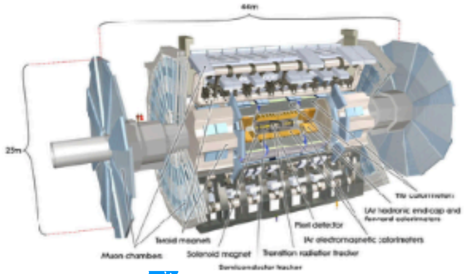

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Data I like

Data I would have used to like if I had been a data hipster

Data you like

...maybe discoveries?

Data produced by the LHC (multiplied by large number)

What we can record in full after the trigger

Data ~nobody likes

This works for a number of LHC measurements (& searches…)

https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2020-010/

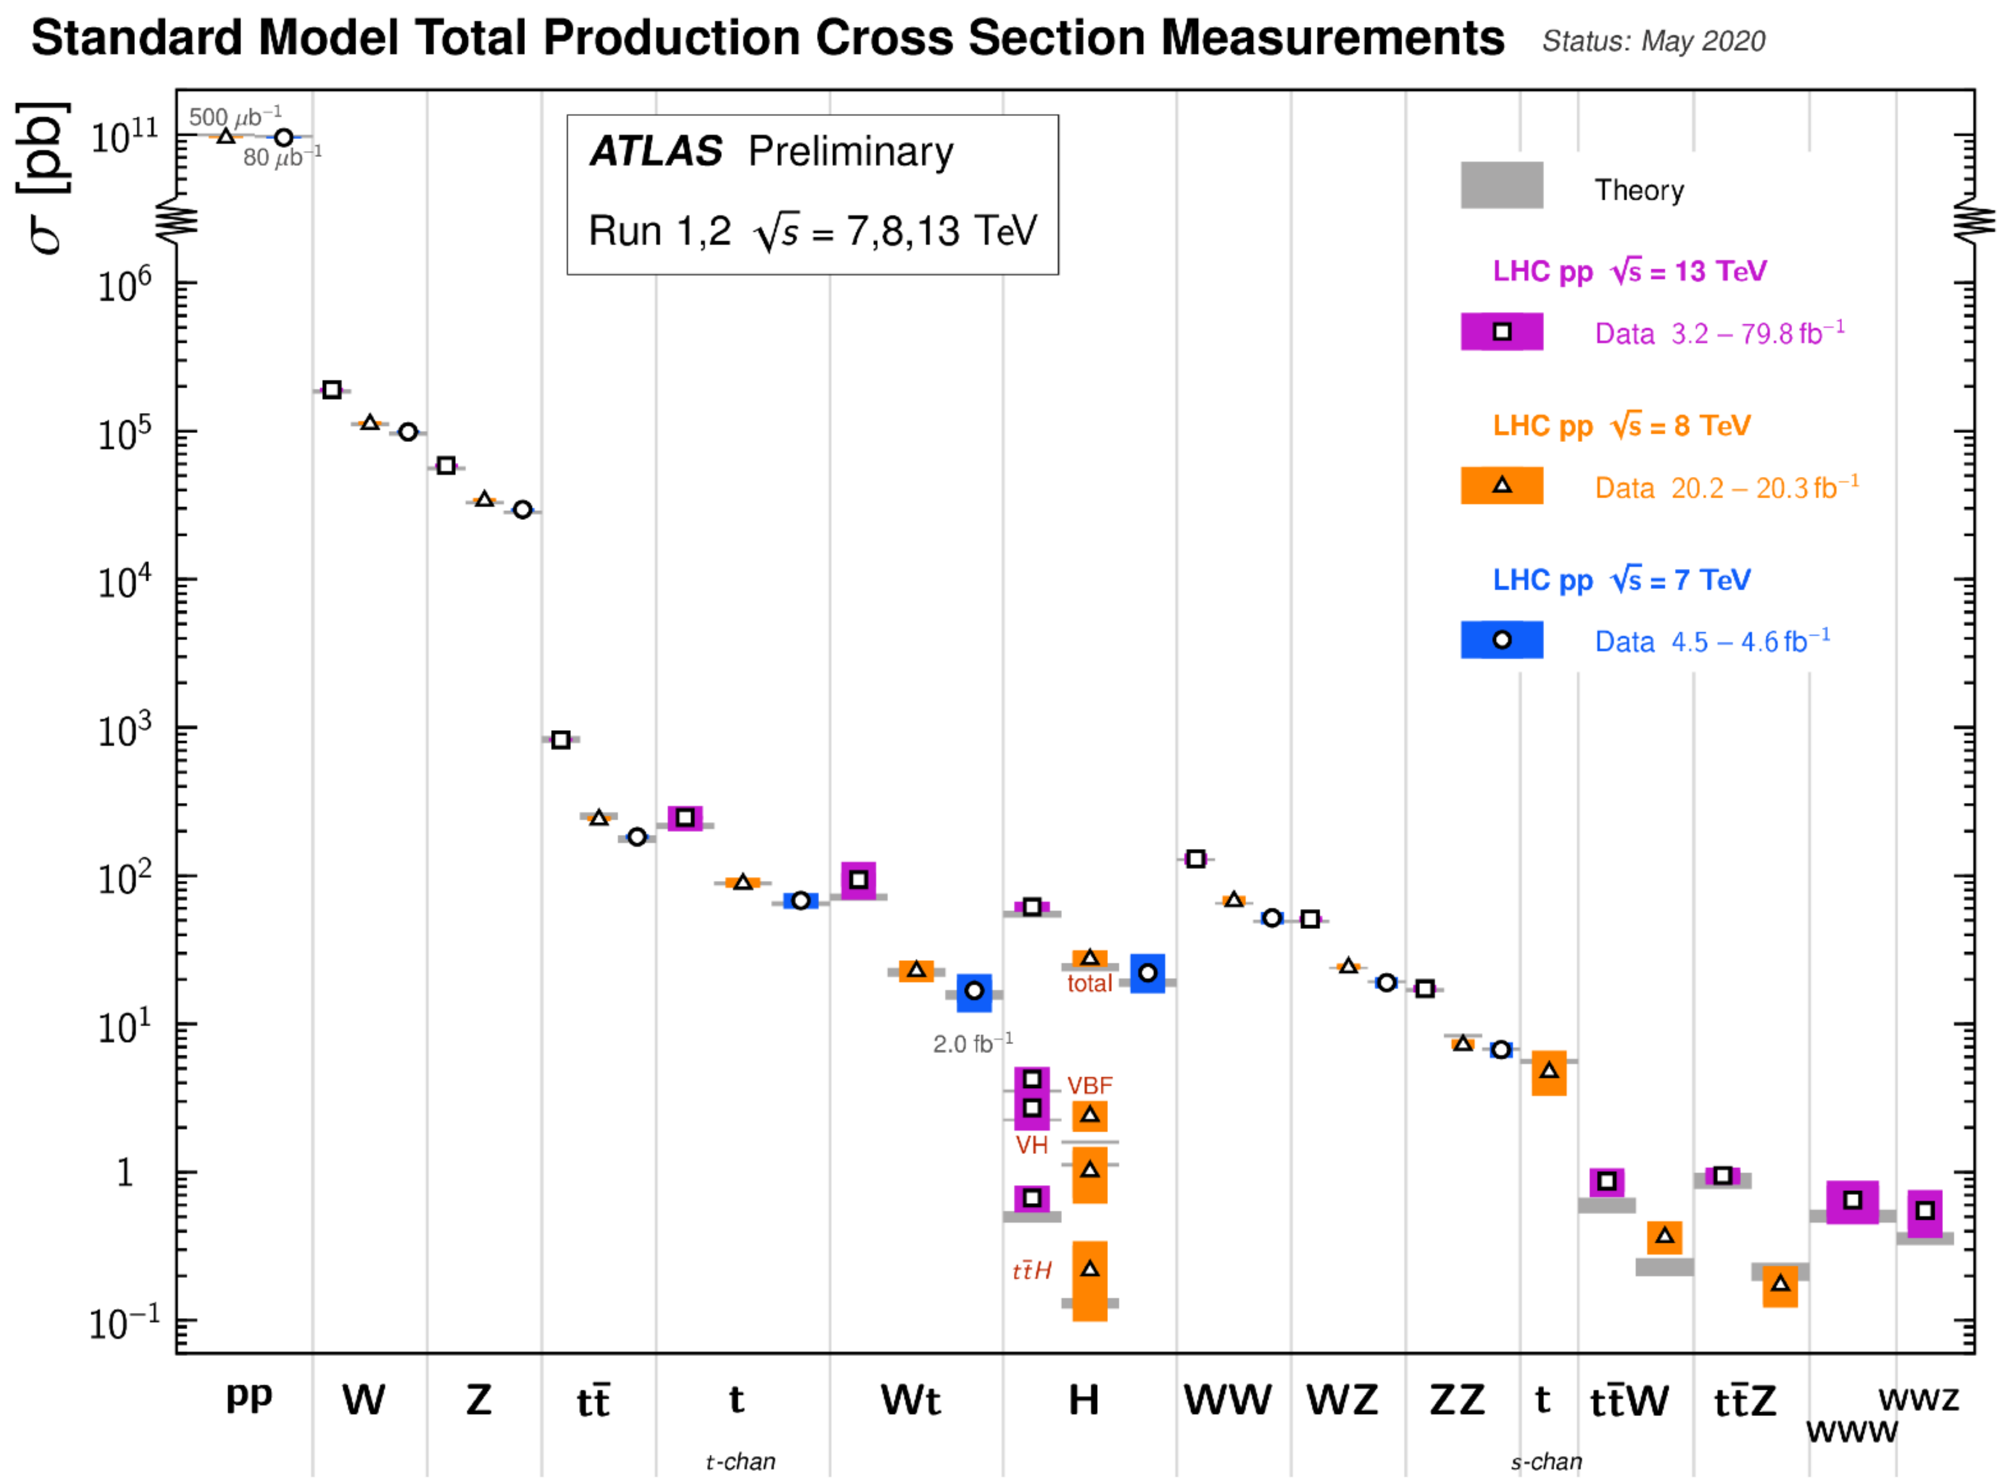

...but are we missing something?

What is interesting at the LHC?

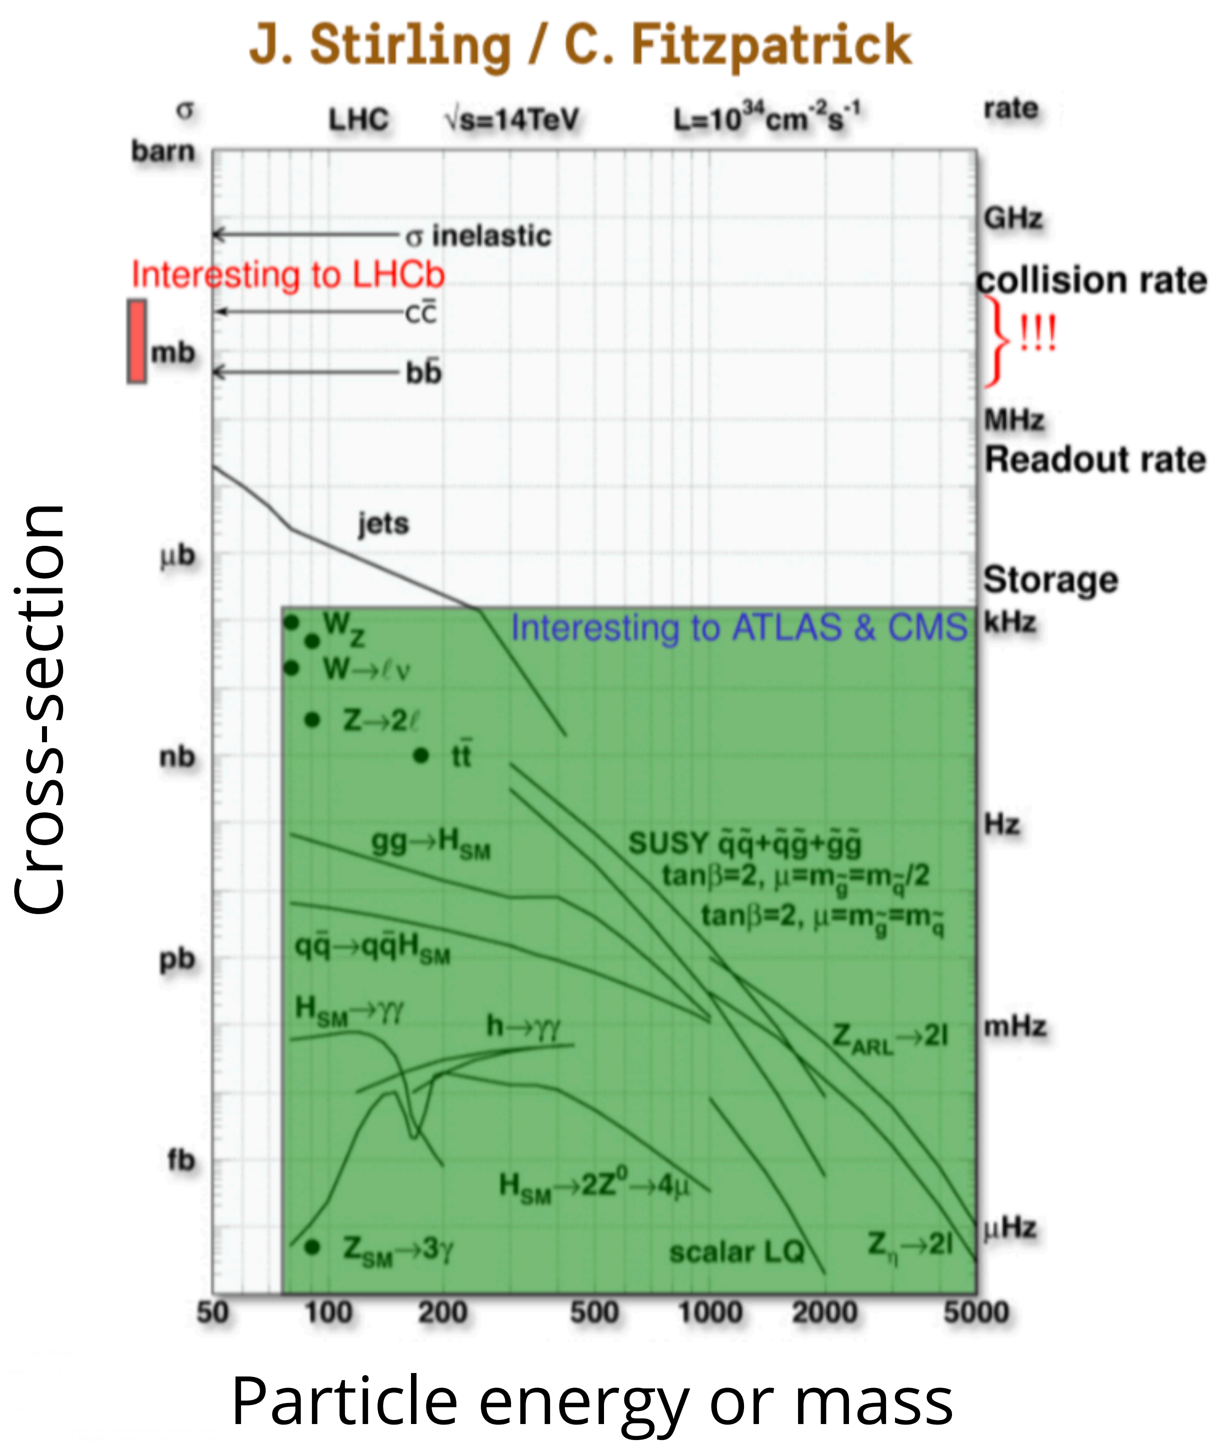

Cross-section * Luminosity
= number of events produced

Challenges:

The **definition of "interesting"** changes experiment by experiment

Rare signal processes that are buried in **high-rate backgrounds** have to be discarded

The ATLAS trigger menu in Run-2

Trigger menu decided in advance of data taking period: example for 2018

Trigger	Typical offline selection	Trigger Selection		L1 Peak Rate [kHz]	HLT Peak Rate [Hz]
		L1 [GeV]	HLT [GeV]	$L=2.0\times10^{34}$ $cm^{-2}s^{-1}$	
Single leptons	Single isolated μ, $p_T > 27$ GeV	20	26 (i)	16	218
	Single isolated tight e, $p_T > 27$ GeV	22 (i)	26 (i)	31	195
	Single μ, $p_T > 52$ GeV	20	50	16	70
	Single e, $p_T > 61$ GeV	22 (i)	60	28	20
	Single τ, $p_T > 170$ GeV	100	160	1.4	42
Single photon	One loose γ, $p_T > 145$ GeV	24 (i)	140	24	47
Single jet	Jet ($R = 0.4$), $p_T > 435$ GeV	100	420	3.7	35
	Jet ($R = 1.0$), $p_T > 480$ GeV	111 (topo: $R = 1.0$)	460	2.6	42
	Jet ($R = 1.0$), $p_T > 450$ GeV, $m_{jet} > 45$ GeV	111 (topo: $R = 1.0$)	420, $m_{jet} > 35$	2.6	36

TRIG-2019-04, ATL-DAQ-PUB-2019-001

- More or less flexible to adjustments (changes need very good reasons!)
  - Follows priorities dictated by experiment's physics strategy
  - Algorithms for object identification/selection also make use of machine learning

Where are the limitations to record (more) data?

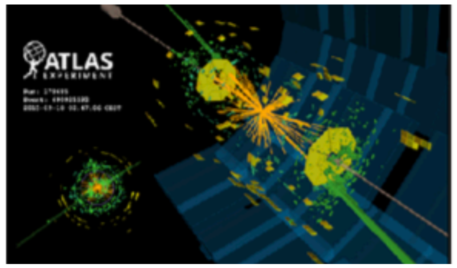

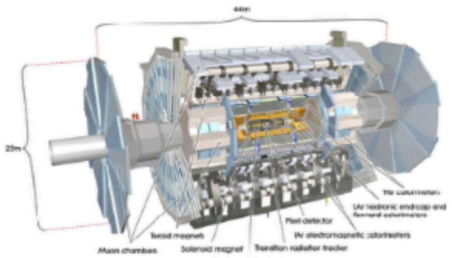

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Detector readout to hardware trigger

L1 hardware trigger
reconstruction & identification of physics object in hardware → rates

CPU for processing events
both within the software trigger (HLT) and in the offline farm

Disk/tape to store events

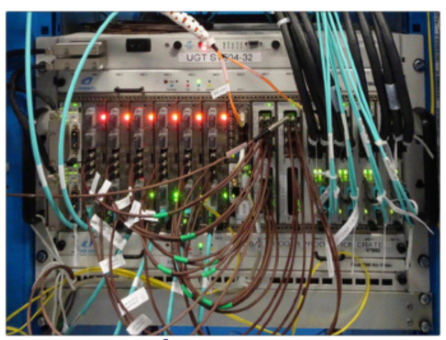

Image from CMS HEPHY

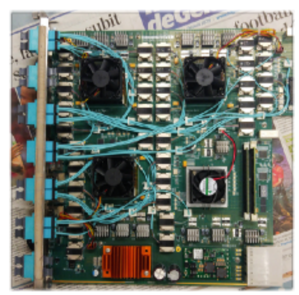

Image from UChicago website

Image from C. Bernius's talk

Image from CERN

Generally triggering on what we know about / what we expect

Main take-home point for this talk:

If we overcome these limitations, we can gain improvements beyond sqrt(N) for the Run-3 dataset

Where are the limitations to record (more) data?

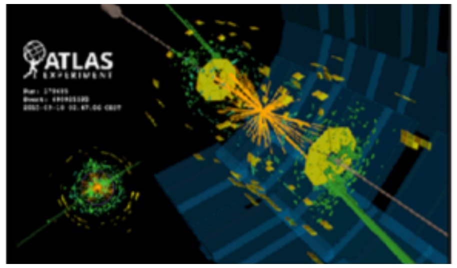

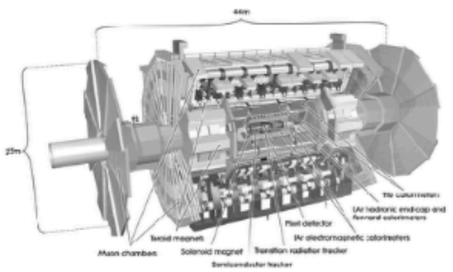

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Detector readout to hardware trigger

L1 hardware trigger

CPU for processing events

Disk/tape to store events

Generally triggering on what we know about / what we expect

For Run-3, see e.g. LHCb and ALICE

For ideas, see e.g.
this talk by M. Pierini
related: Ben Nachman's talk later
LHC Olympics paper

LHCB-TDR-016

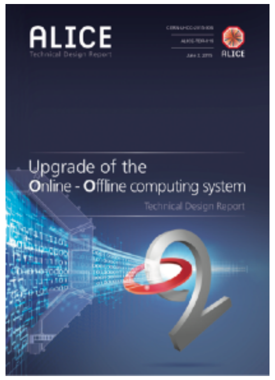

ALICE-TDR-019

Deferred to HL-LHC and future experiment trigger talks…

[some ideas from the Snowmass Instrumentation Frontier]

Where are the limitations to record (more) data?

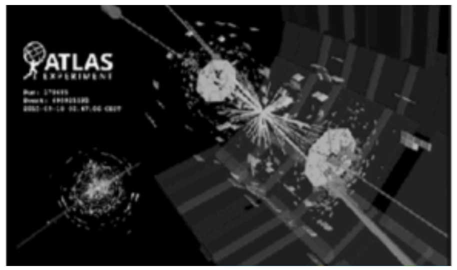

Detector readout
to hardware trigger

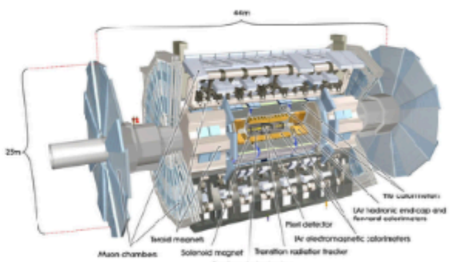

L1 hardware trigger
reconstruction & identification of
physics object in hardware → rates

Event selection (trigger) → Object reconstruction and calibration → Data analysis

CPU for processing events

Disk/tape
to store events

The importance of
trigger turn-on curves:

Naively: the more events we record,
the more we can analyse

Pragmatically:
- we want to reduce rates from events that we are not going to analyse (e.g. fakes, below turn-on…)
- *but also:* if we limit ourselves to analysing events where the trigger is fully efficient, we are missing out (especially in steeply falling distributions)

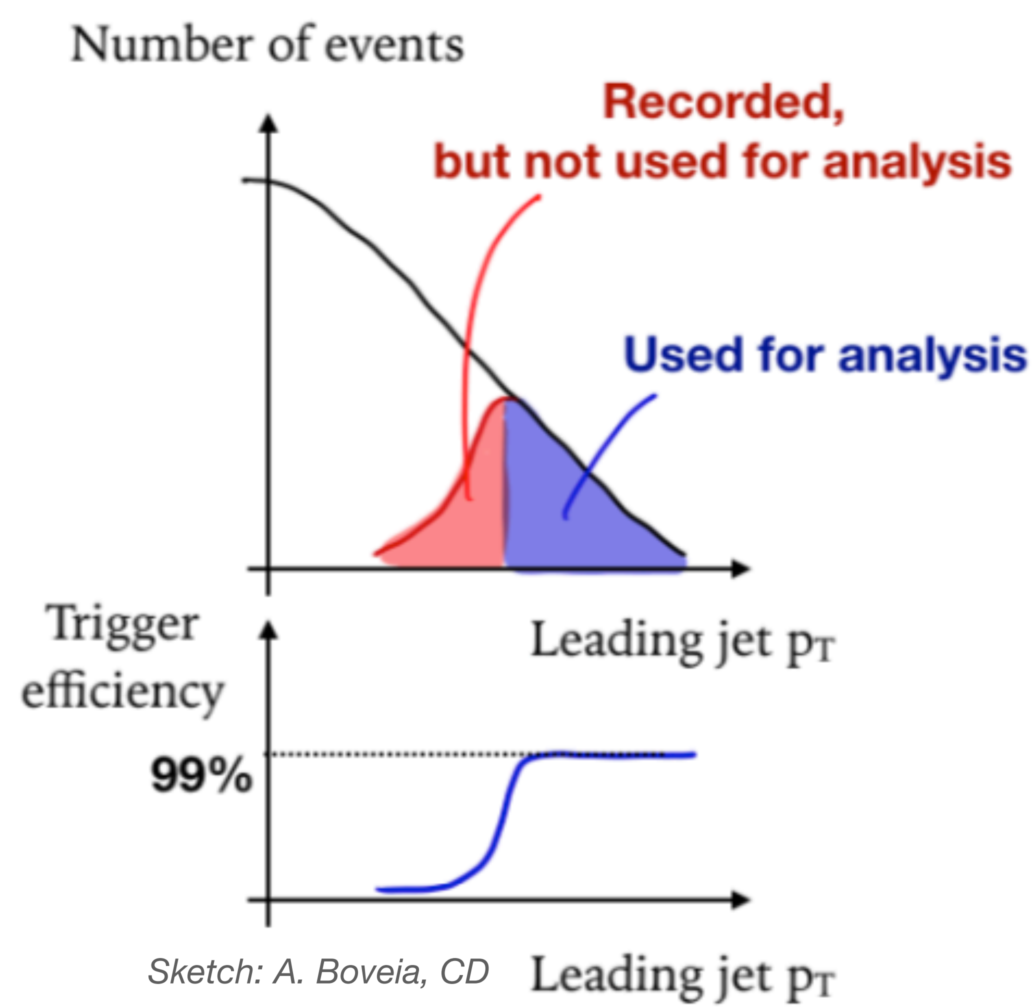

Sketch: A. Boveia, CD

Where are the limitations to record (more) data?

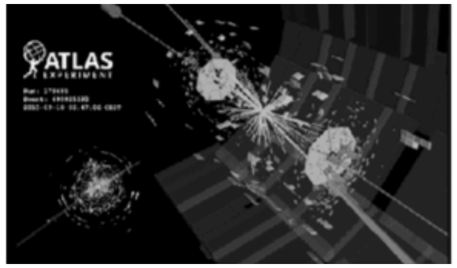

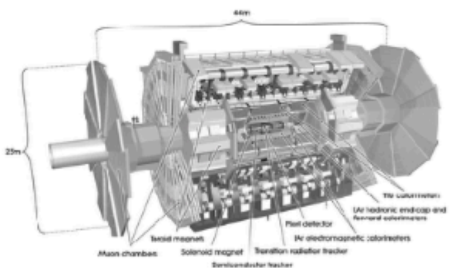

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Detector readout to hardware trigger

L1 hardware trigger

CPU for processing events
both within the software trigger (HLT) and in the offline farm

Disk/tape to store events

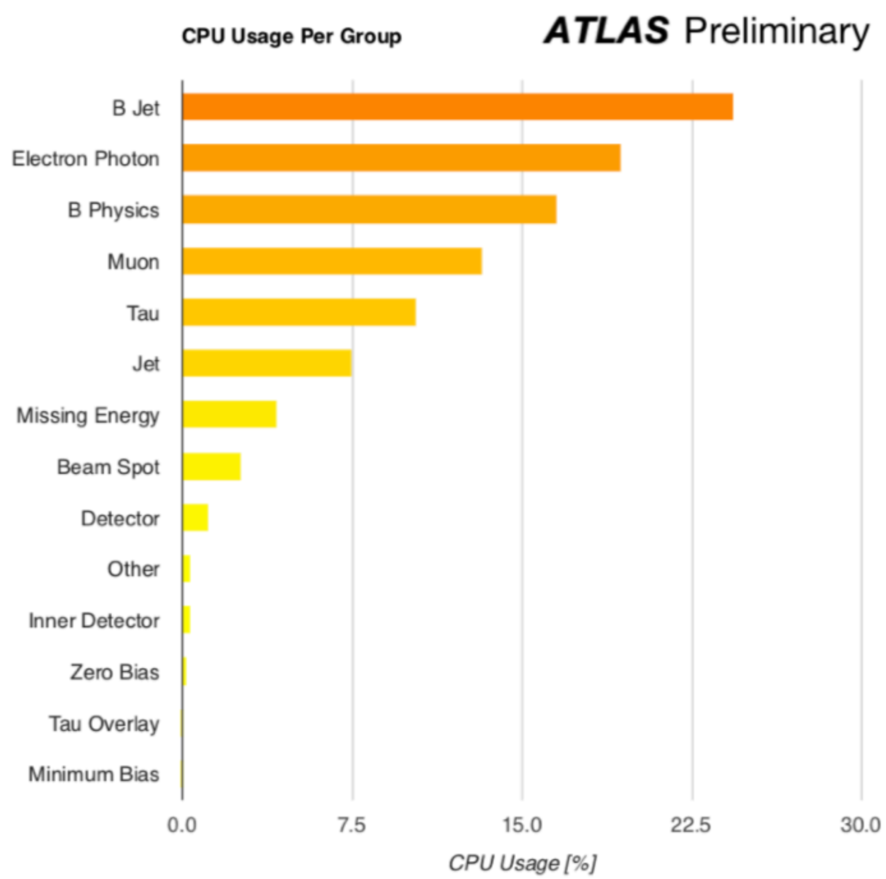

Both online and offline (+ grid) **CPU farms** have a large but **limited processing power**

→ this affects **how much data we can process** promptly / offline and therefore the **trigger thresholds**

to overcome this limitation:

Optimize online/offline reconstruction code
+
Aim for **similar performance** for online and offline physics objects

Refine **trigger algorithms and selections** to get more of the data we want

Use **non-standard workflows** that reduce (immediate) CPU use

Where are the limitations to record (more) data?

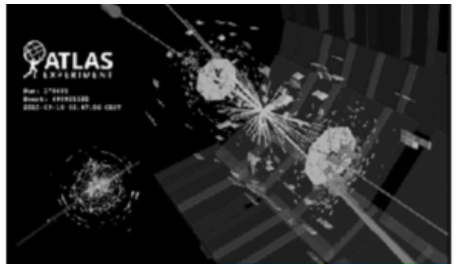

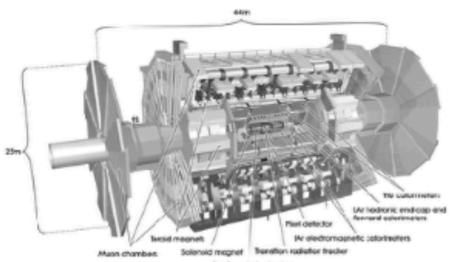

Event selection (trigger) → Object reconstruction and calibration → Data analysis

Detector readout to hardware trigger | **L1** hardware trigger | **CPU** for processing events | **Disk/tape** to store events

- A naive estimate of LHC data sizes:
  - up to 30 million collisions/second (MHz)
  - 1-1.5 MB/data per collision
    - Assuming no improvements, this increases with pile-up
  - 30 MHz * 1 MB = 30 TB/s
  - 30 TB/s * 10e+6 s/year (day & night) ~ 0.05 ZB/year
- [Facebook 2014]: 600 TB/day ~ 200 PB/year
  - **"There's always a bigger fish"** [C. Tully's talk @ siRTDM18]
  - But bigger fish also have bigger money…

A less-naive estimate (dominated by MC needs):

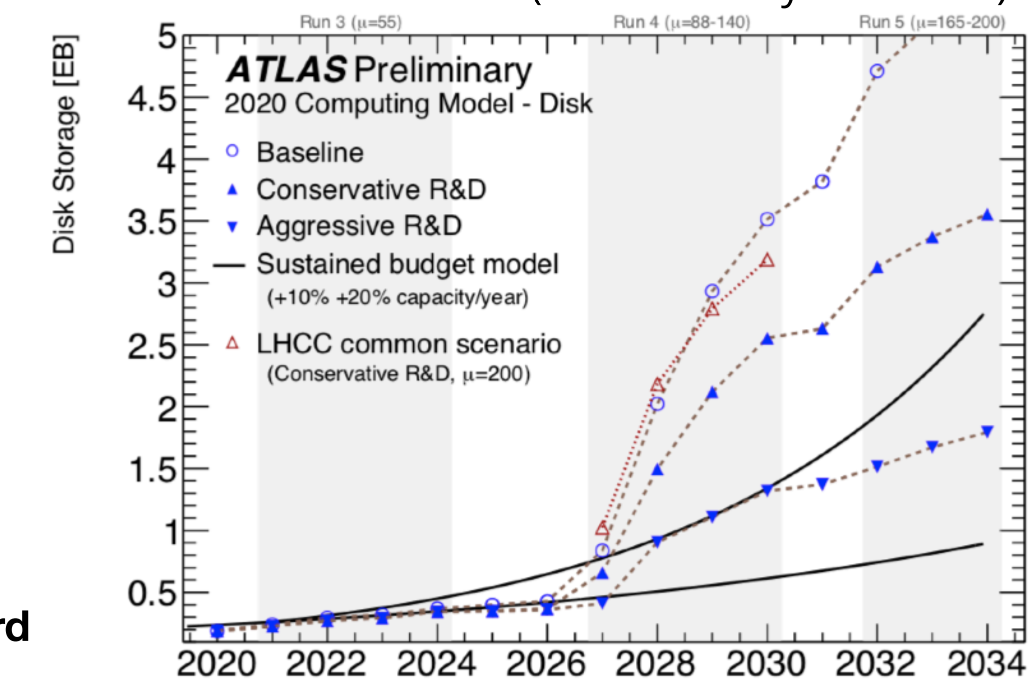

[ComputingandSoftwarePublicResults]

to overcome this limitation:

- Use **non-standard workflows** that reduce disk use
- Refine **trigger algorithms and selections** to get more of the data we want

Hardware trigger improvements for Run-3

The ATLAS Run-3 trigger system

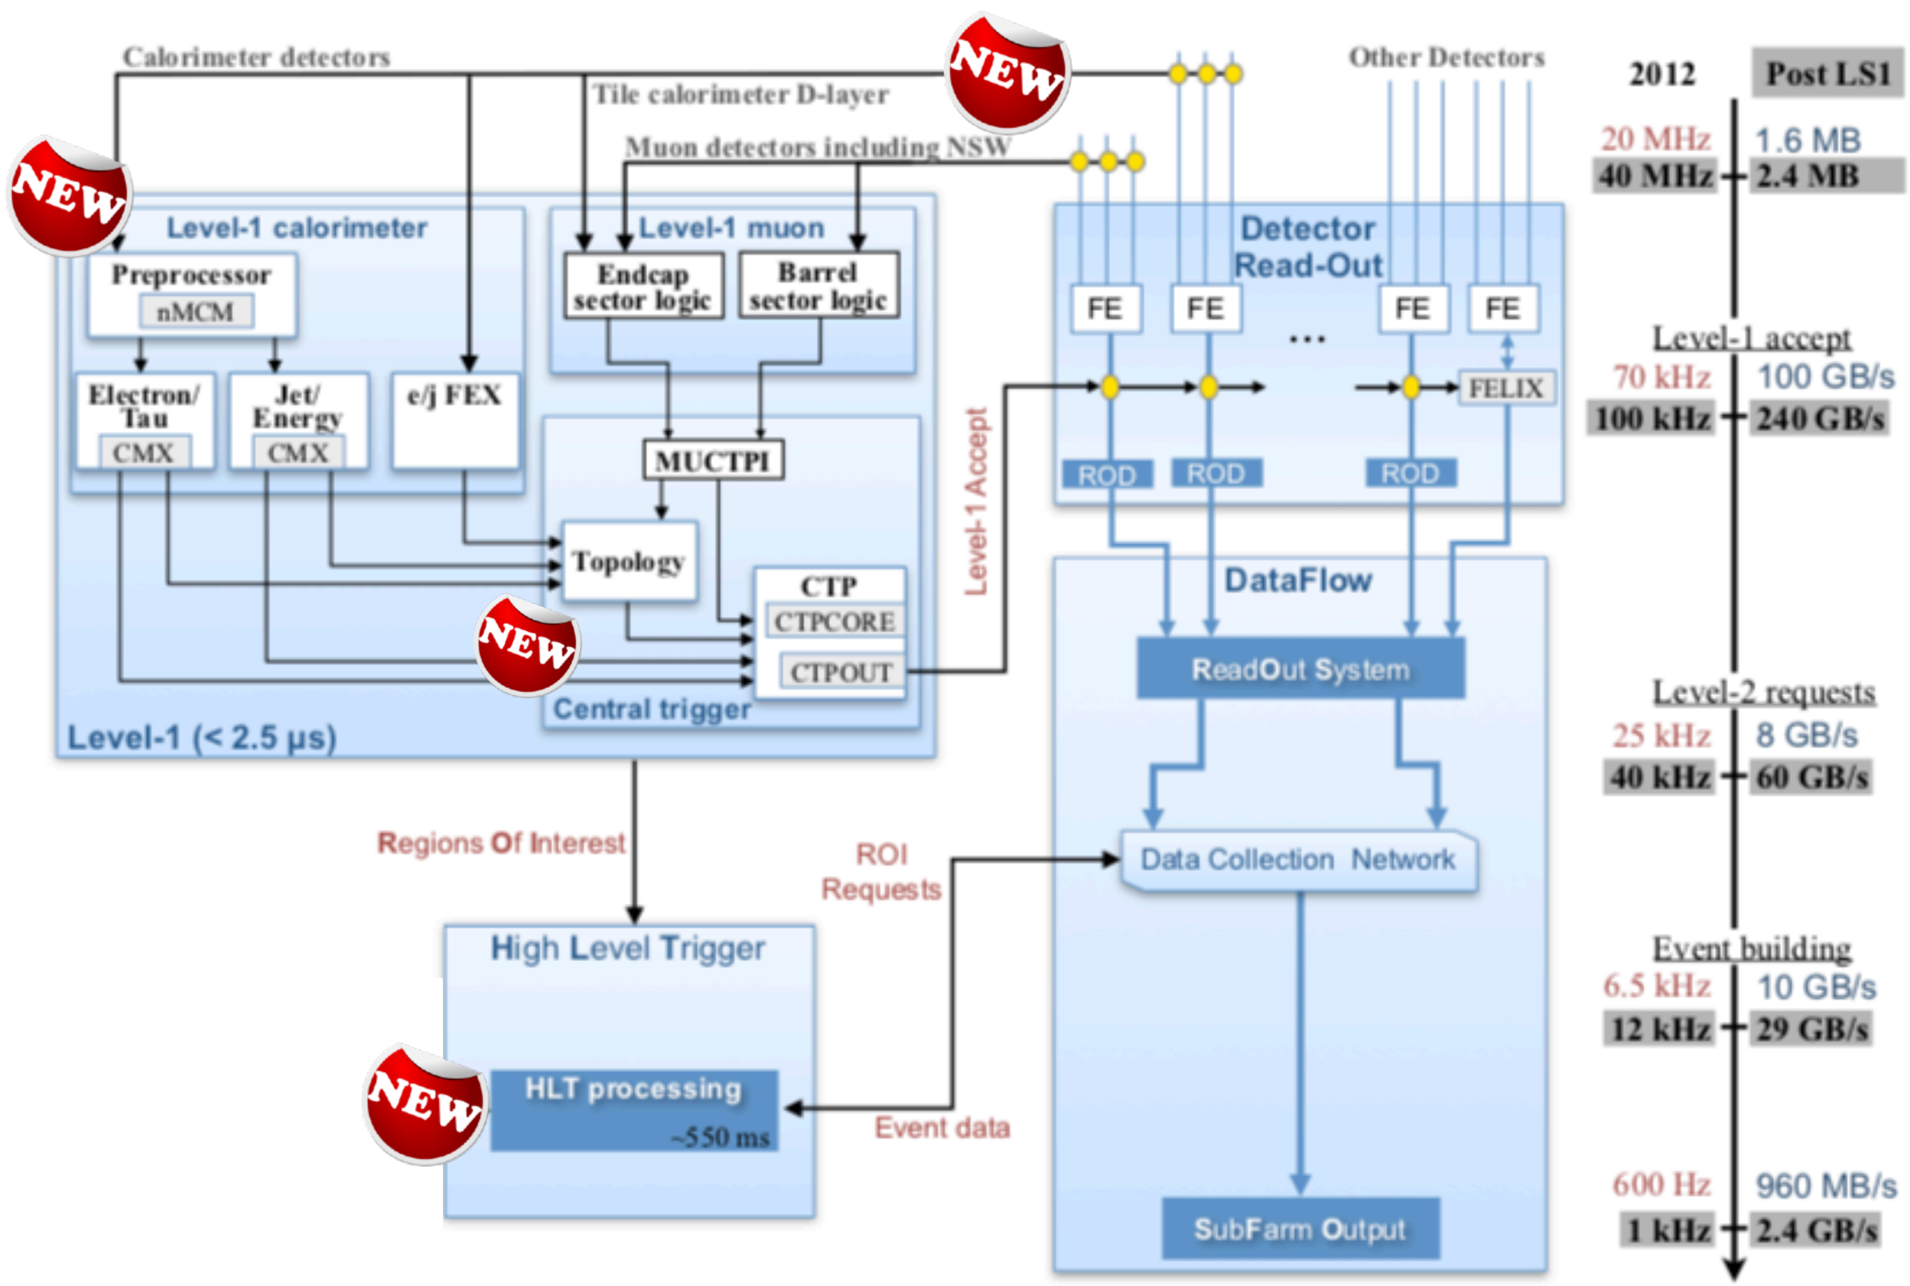

L1Calo for electrons and photons: *eFEX*

ATL-DAQ-SLIDE-2020-310

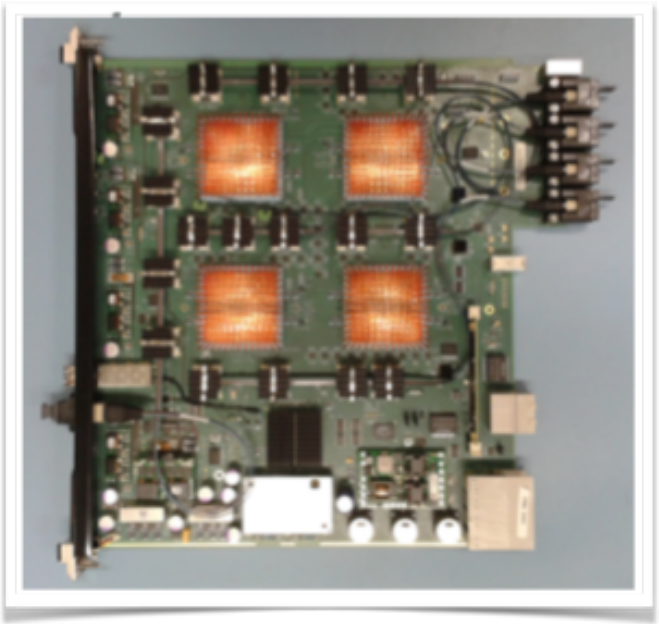

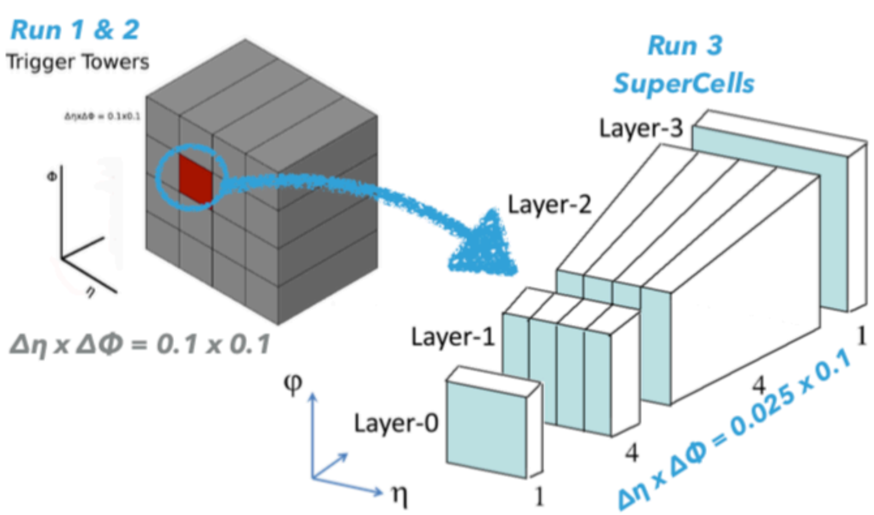

- More granular input for electron and photon identification in Run-3
  - Can be used for more sophisticated algorithms

ATL-COM-DAQ-2018-019

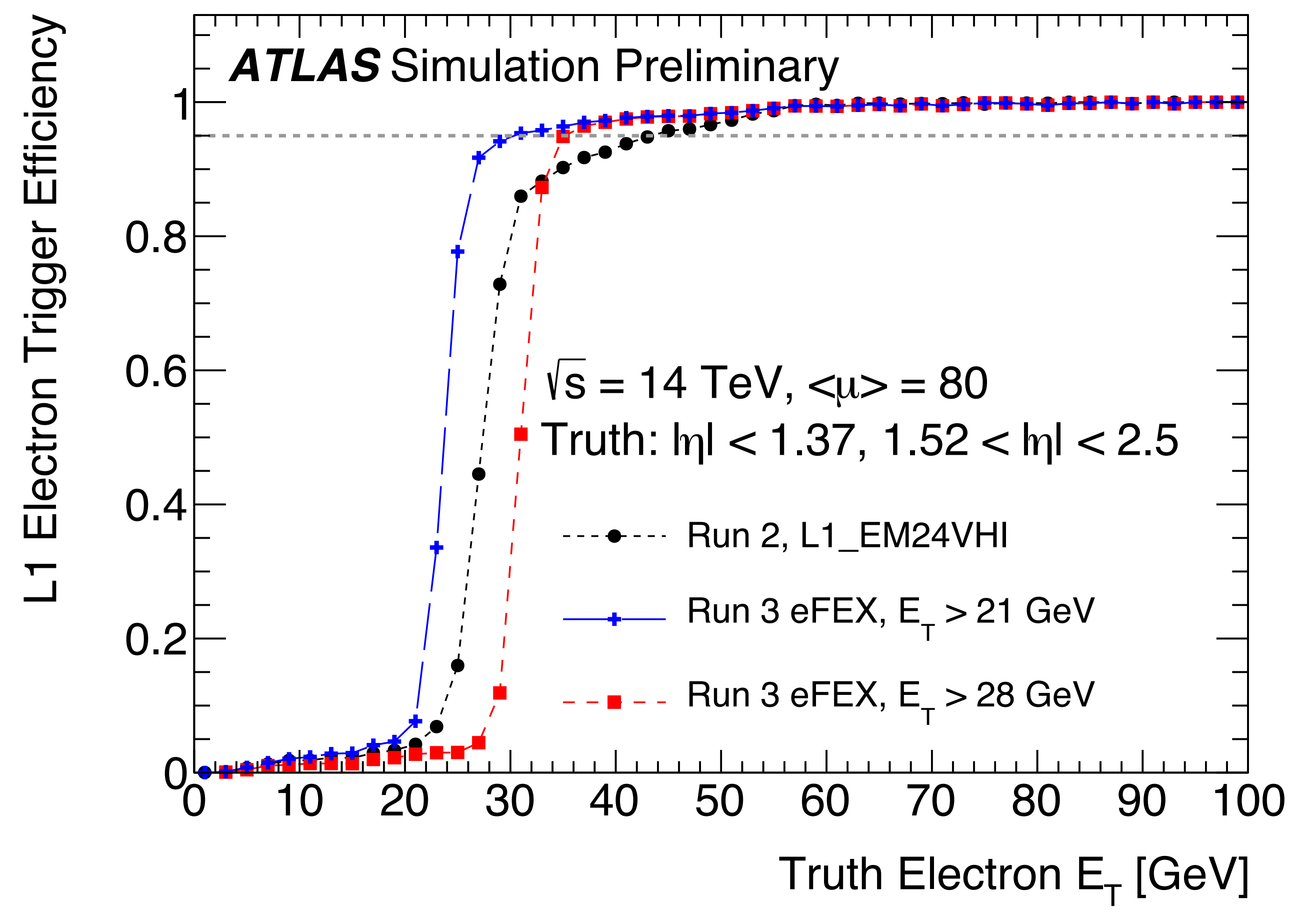

- Much "steeper" turn-ons for Run-3
  - Improves the rate of useful events
- Trigger rate depends on threshold
  - Run-3 L1 21 GeV threshold leads to same event rate as 24 GeV Run-2 L1 threshold
  - ET>28 GeV has half the rate as ET>28 GeV

L1Calo for jets, MET and taus: *jFEX*

ATL-DAQ-SLIDE-2020-135

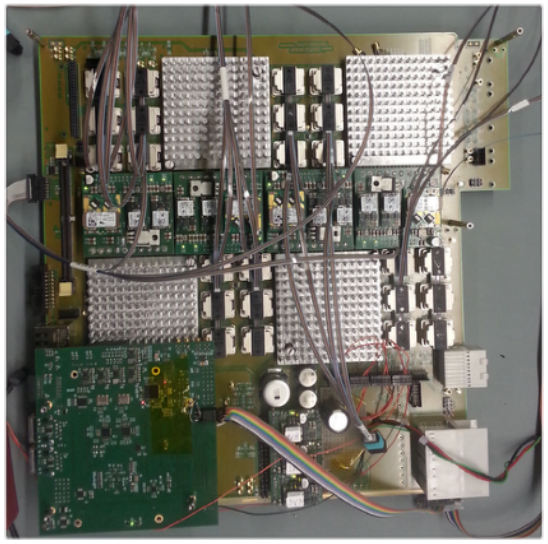

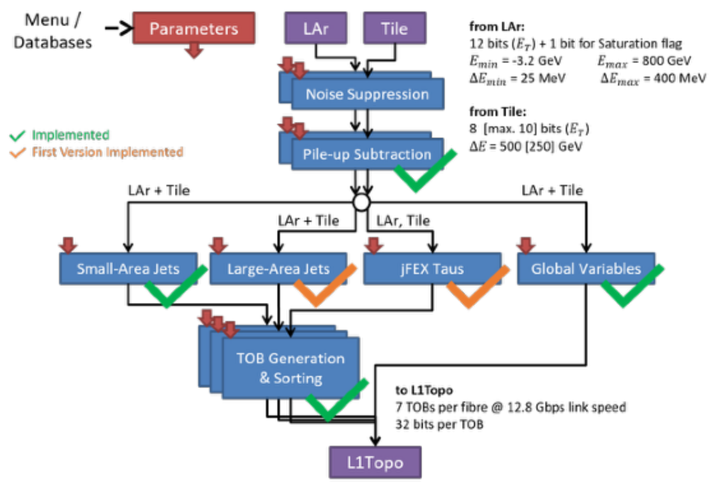

Board prototype, from arXiv:1806.09207

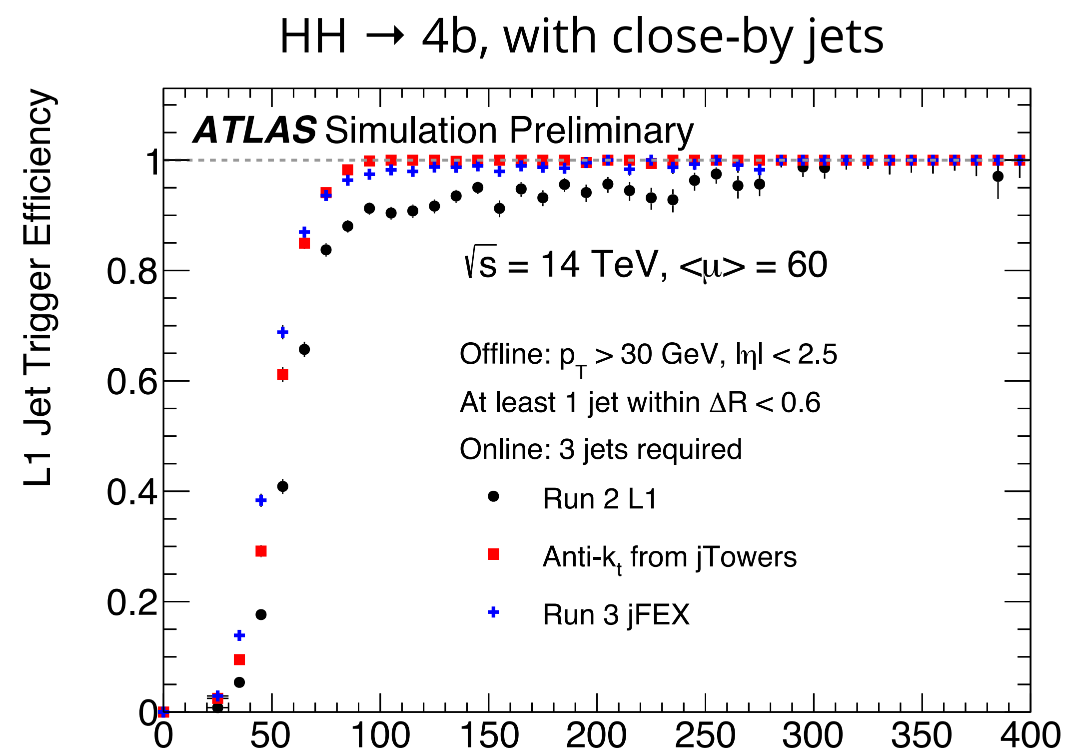

ATL-COM-DAQ-2018-019

- Used to trigger on jets, MET and hadronic taus
  - Inputs: calorimeter towers
- Improvements with respect to Run-2: more refined algorithms, e.g.
  - square jets (Run-2) →rounder jets (Run-3)

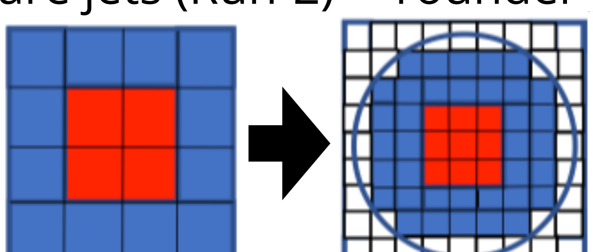

  - improved pile-up mitigation
    - use custom noise thresholds on inputs
    - MET calculated after average energy subtraction

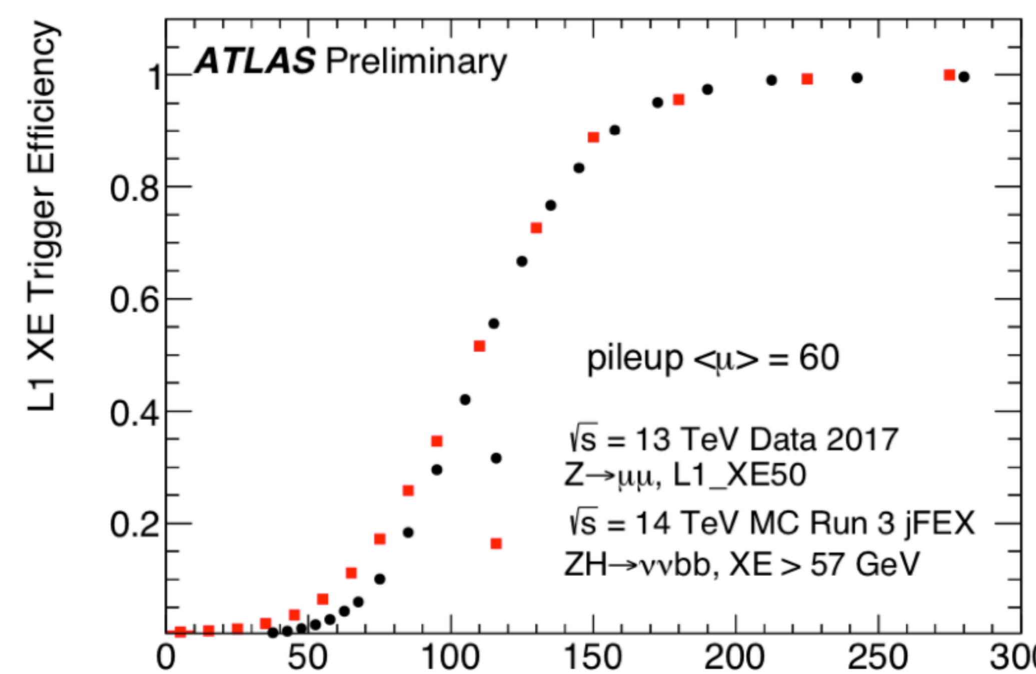

Using run-2 as reference

L1Calo for large-R jets and MET: *gFEX*

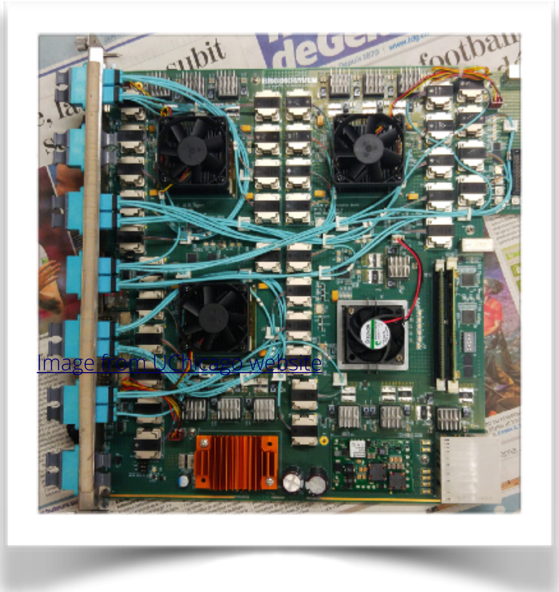

- Inputs to the board: coarse towers from **entire calorimeter**
  - Ideal for large-R jet identification

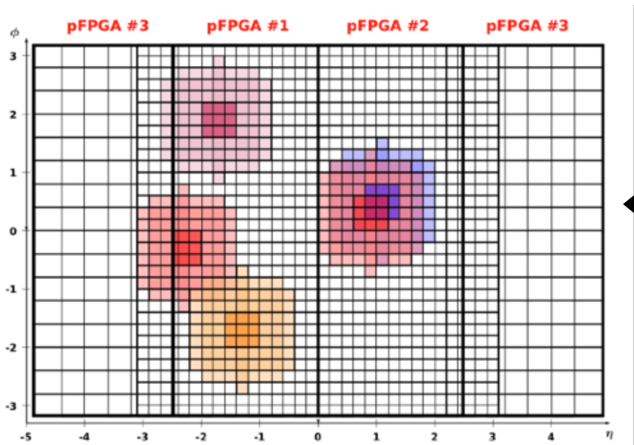

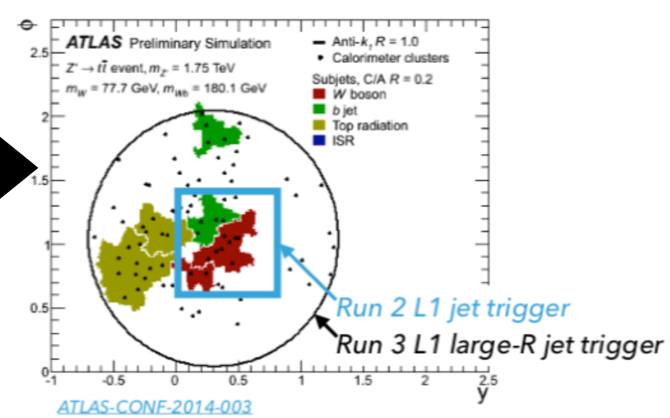

  - Full-scan algorithms can be used for event-level quantities (e.g. pile-up density)

Boosted top simulation

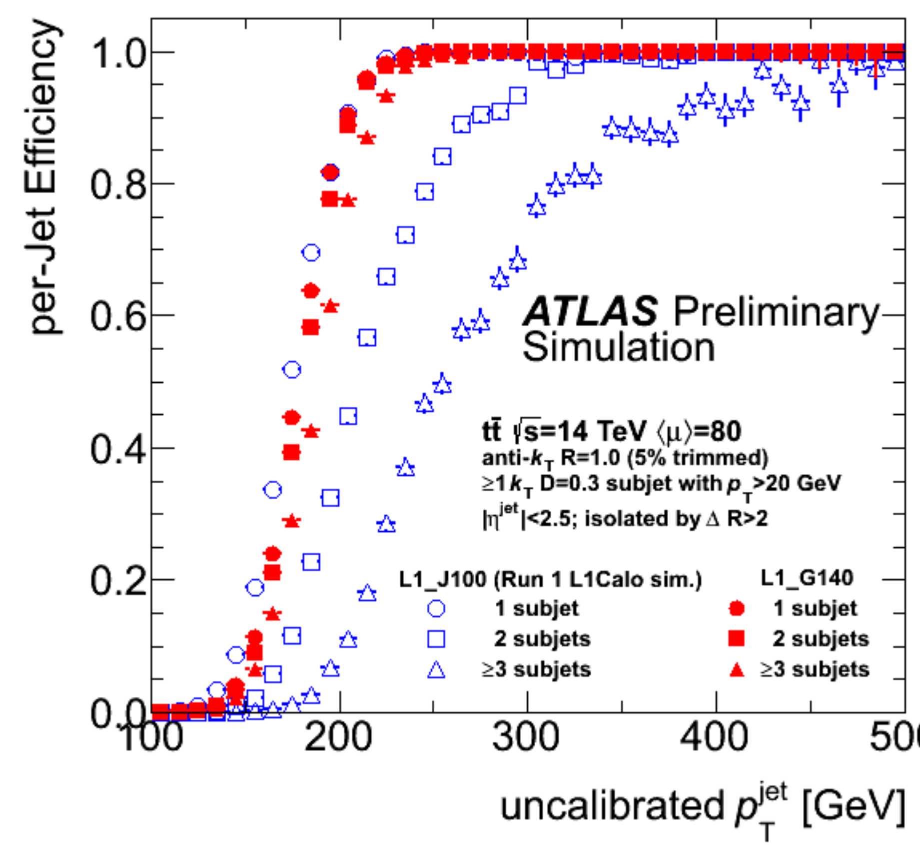

More efficient triggering on large-R jets-with-subjets (with gFEX) than Run-2 (standard square jets)

Muon triggers and the New Small Wheels (NSW)

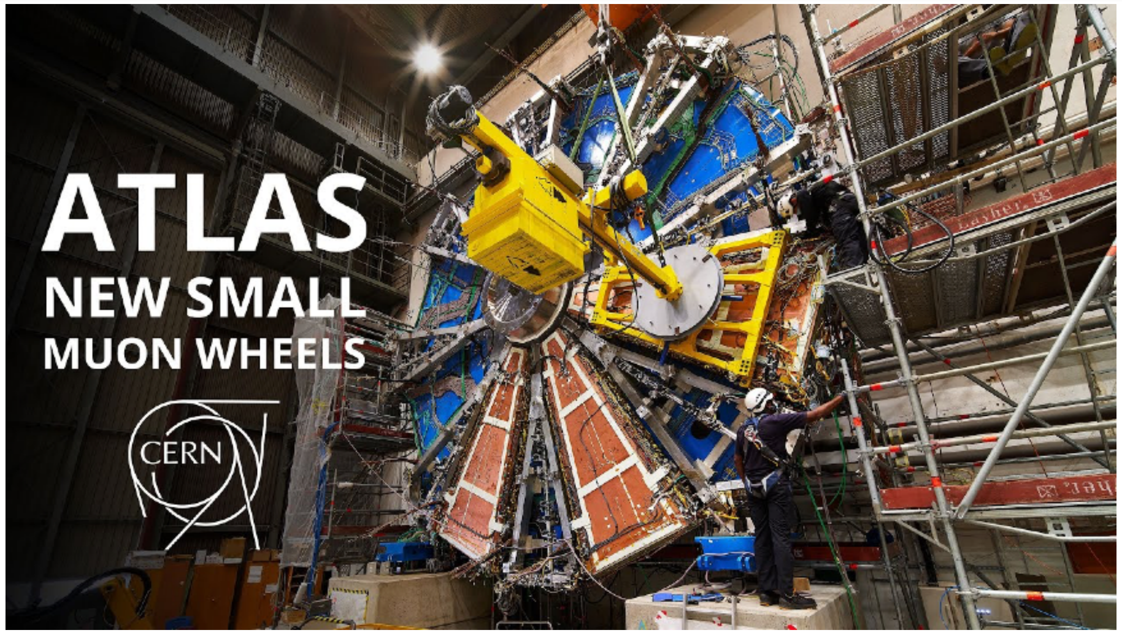

- Significant trigger rate from endcap muon detector
- Replace forward muon detectors with improved New Small Wheels
- NSW playing significant role in Run-3 triggers

Better identification of "real" low-pT muons using coincidences
L1MuonTriggerPublicResults

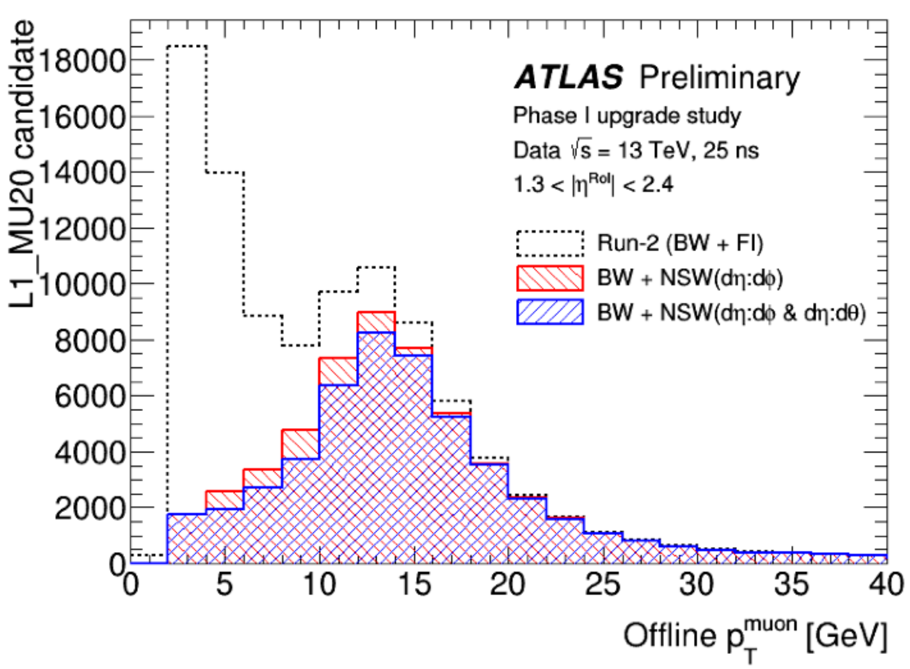

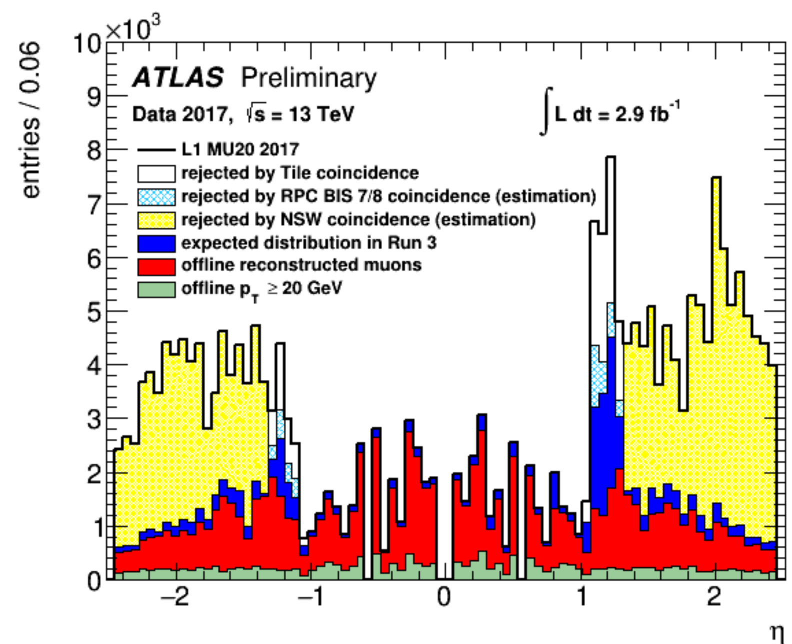

Lower rate → lower thresholds
MuonTriggerPublicResults

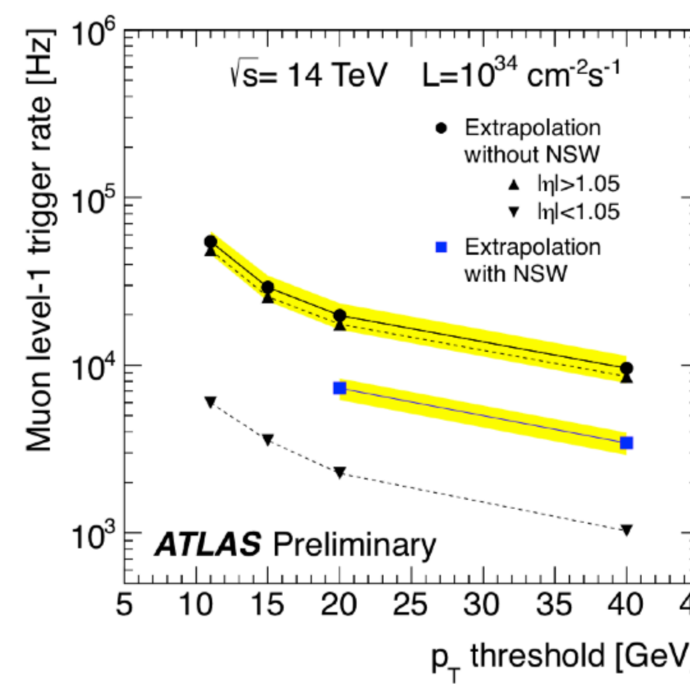

L1 trigger combinations: L1Topo

L1Topo: trigger board to combine L1 objects in multi-object and topological algorithms

B. Carlson @ PITT-PACC 2017

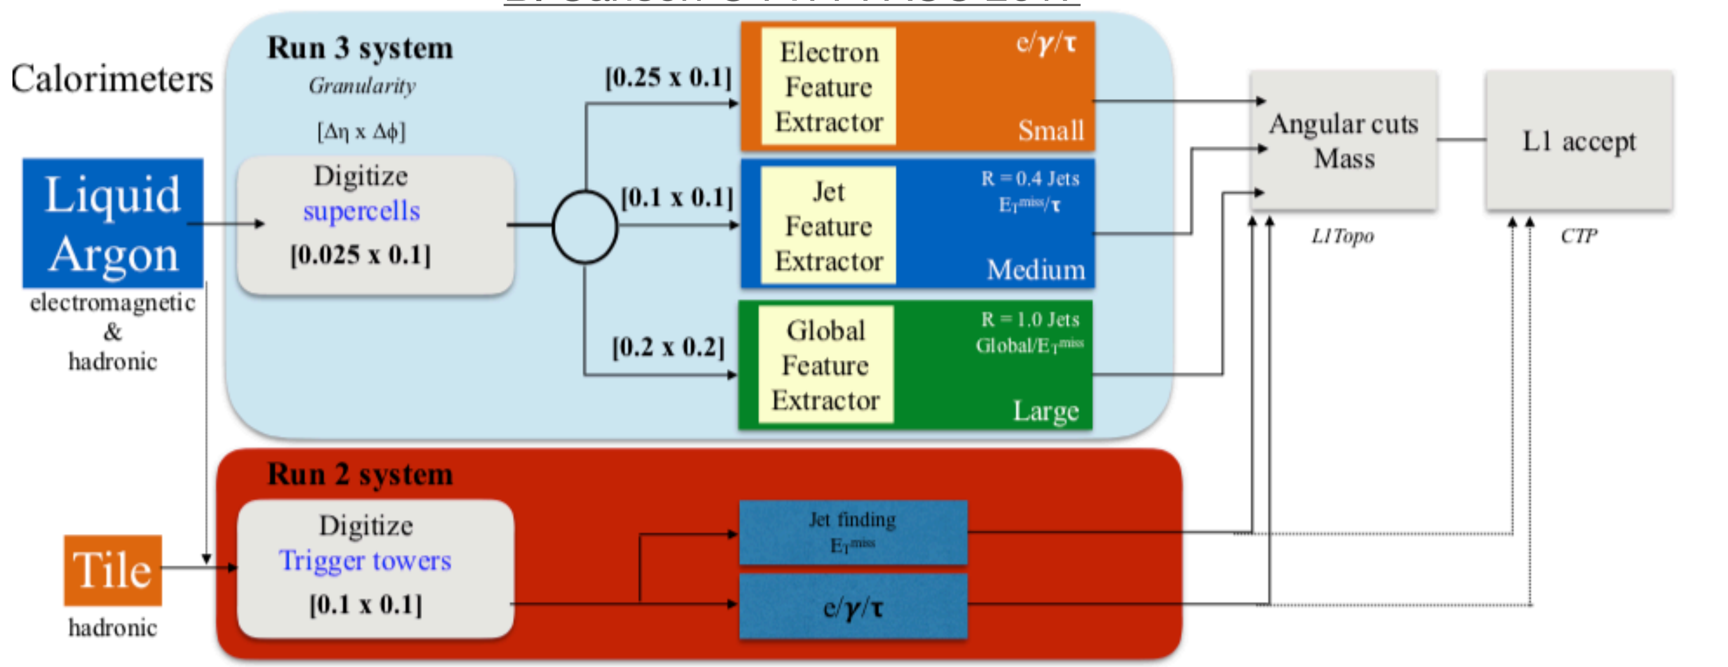

*muons not shown

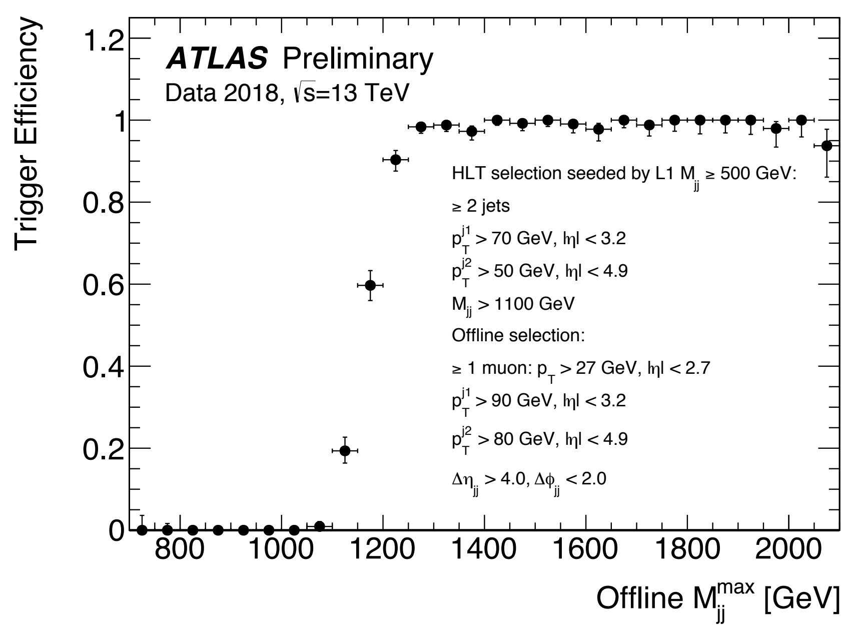

- Case of **VBF topologies**: can select large-dijet- mass events already at the hardware level
- Can also be tailored towards e.g. measurements / searches in the forward region

TriggerOperationPublicResults

Software trigger improvements for Run-3

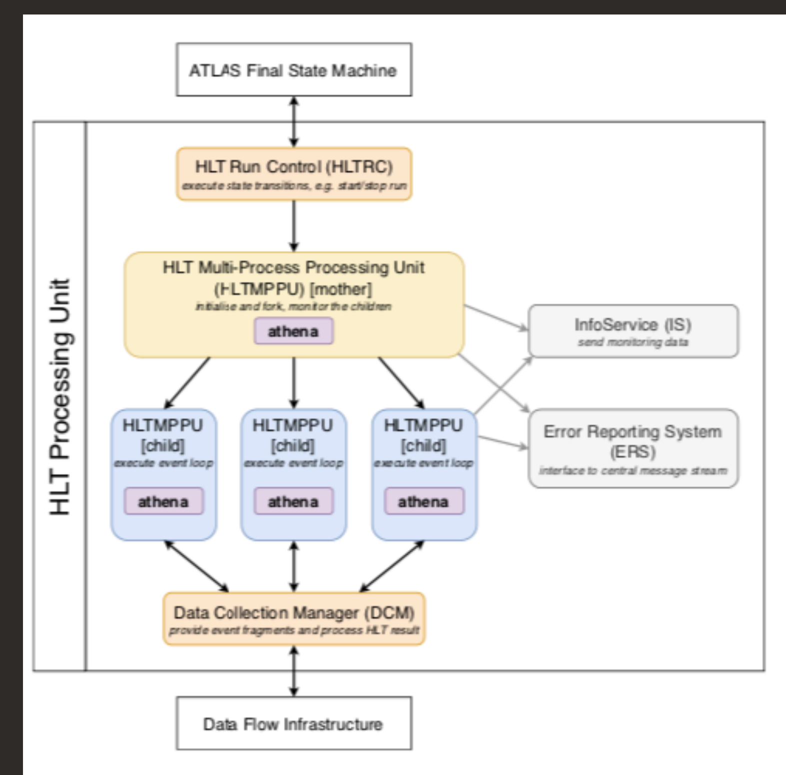

Note: ATLAS is rewriting its trigger software using multithreading (AthenaMT)
Not covered here, for more information, see these proceedings by R. Bielski

Pile-up (and tracking)

Increase in dataset size means increase in instantaneous luminosity
→ increased number of simultaneous interactions / beam crossing → pile-up → higher trigger rates

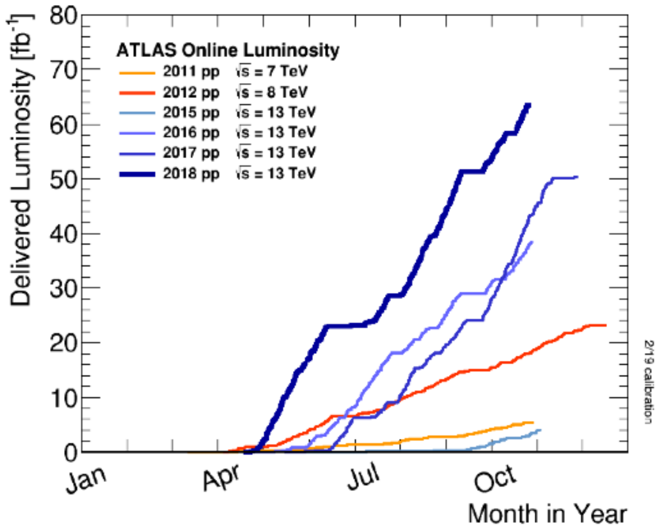

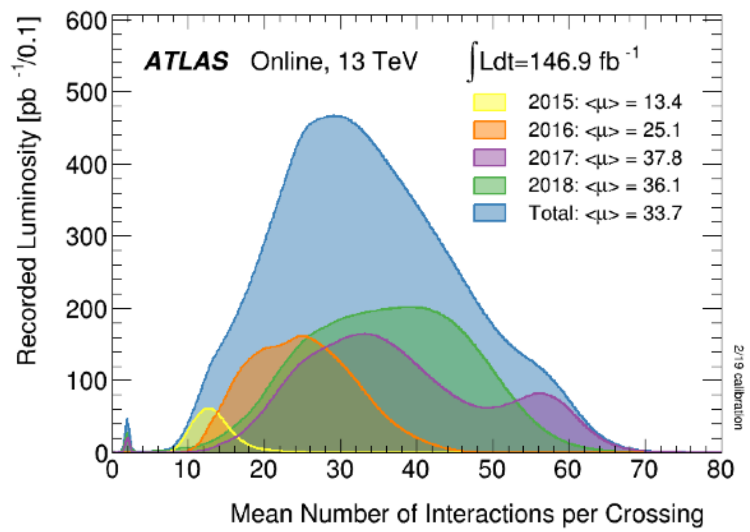

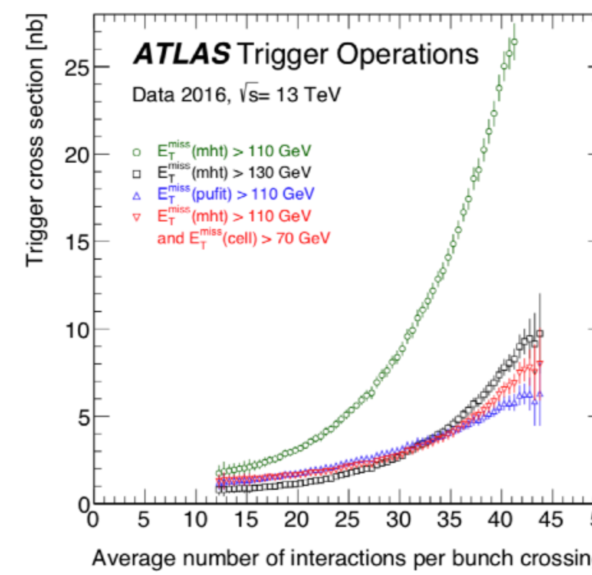

Trigger and data acquisition systems are designed to be as robust as possible to increased pile-up

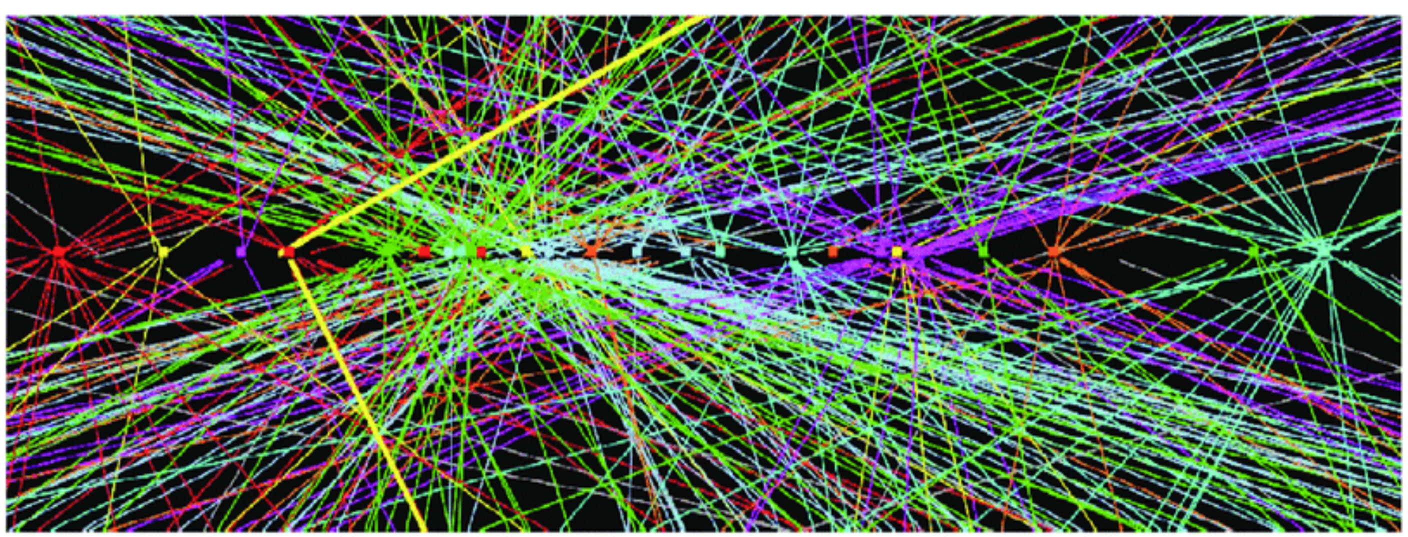

How to meet the pile-up challenge:

- **Software tracking**
  - Challenge: computationally expensive

    FTK paper: arXiv:2101.05078 accepted by JINST
- **Detector timing**
  - Challenge: precision / simulation
  - (Not covered in this talk)

Design and optimization of software tracking

ATL-COM-DAQ-2020-104

(only prepare data once)

Fast tracking step

can reject early

input to

Precision tracking step

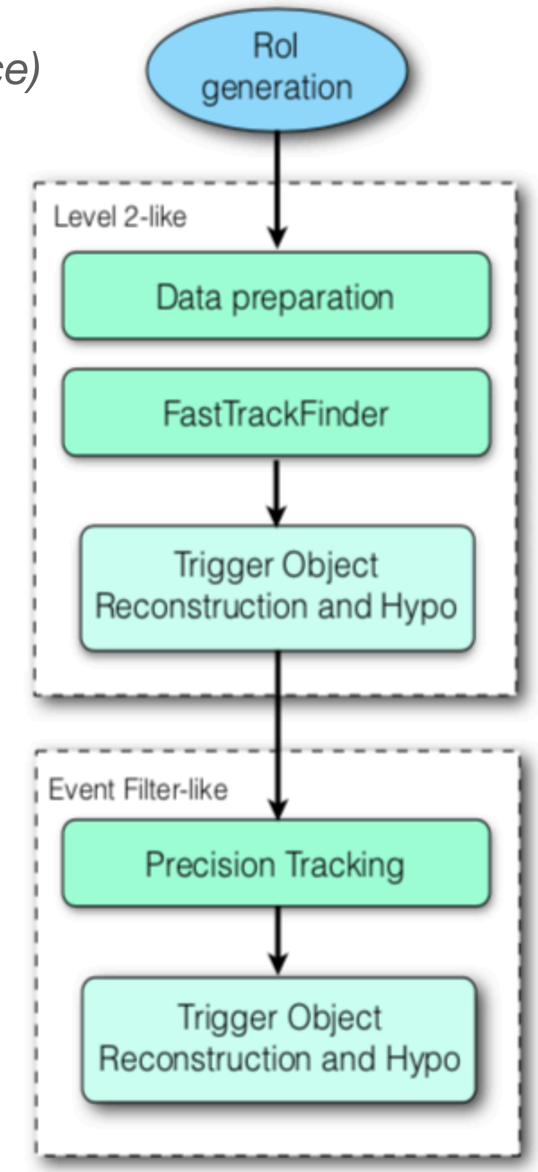

Fast tracking sufficient for e.g. object identification

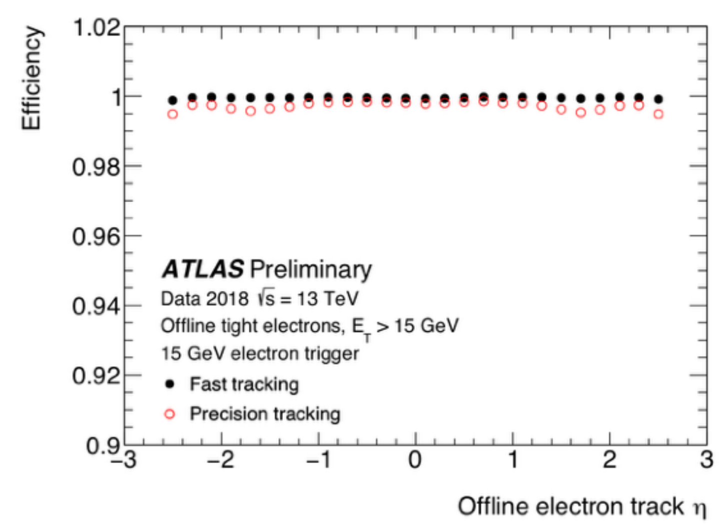

Precision tracking needed for e.g. *b*-jet tagging

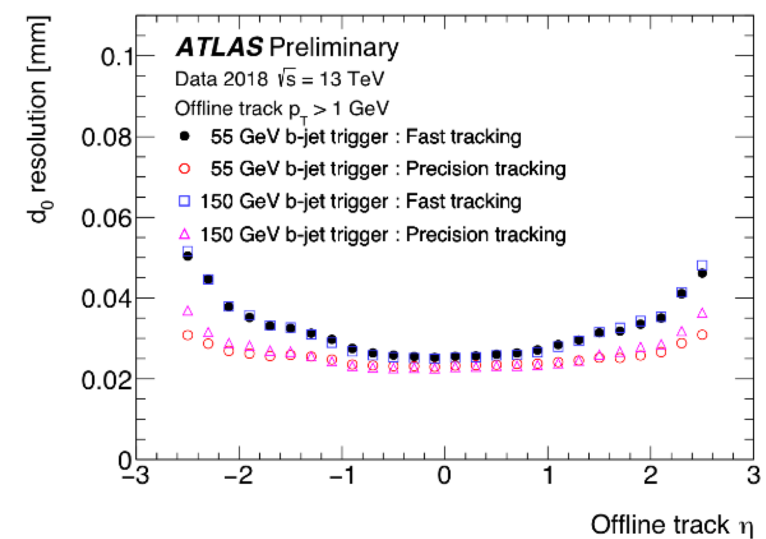

ATL-COM-DAQ-2020-059

- Rewriting of software tracking for Run-3 currently ongoing *[no public results yet]*
  - Improvements to offline tracking also included
  - Machine learning-based improvements: see next slide

On track to have **high-rate *full-scan tracking*** to be used in reconstruction of HLT objects (including long-lived signatures) & non-standard workflows

Further improvements: machine learning

ATL-COM-DAQ-2021-003

CPU time grows linearly with the number of tracking seeds (due to combinatorics)
→ reduce the number of fake seeds as soon as possible

ML algorithm
(Classifier based on Kernel Density Estimator)
can predict **probability of pair of hits belonging to the same track**

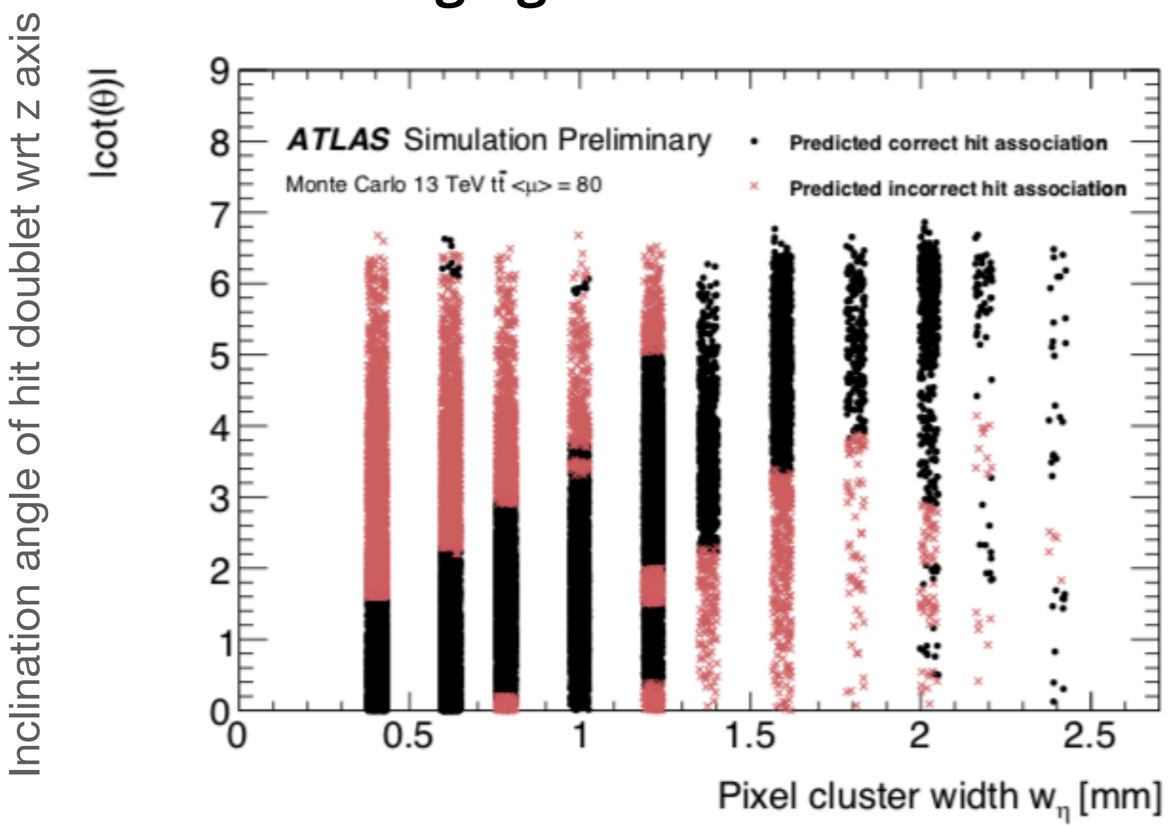

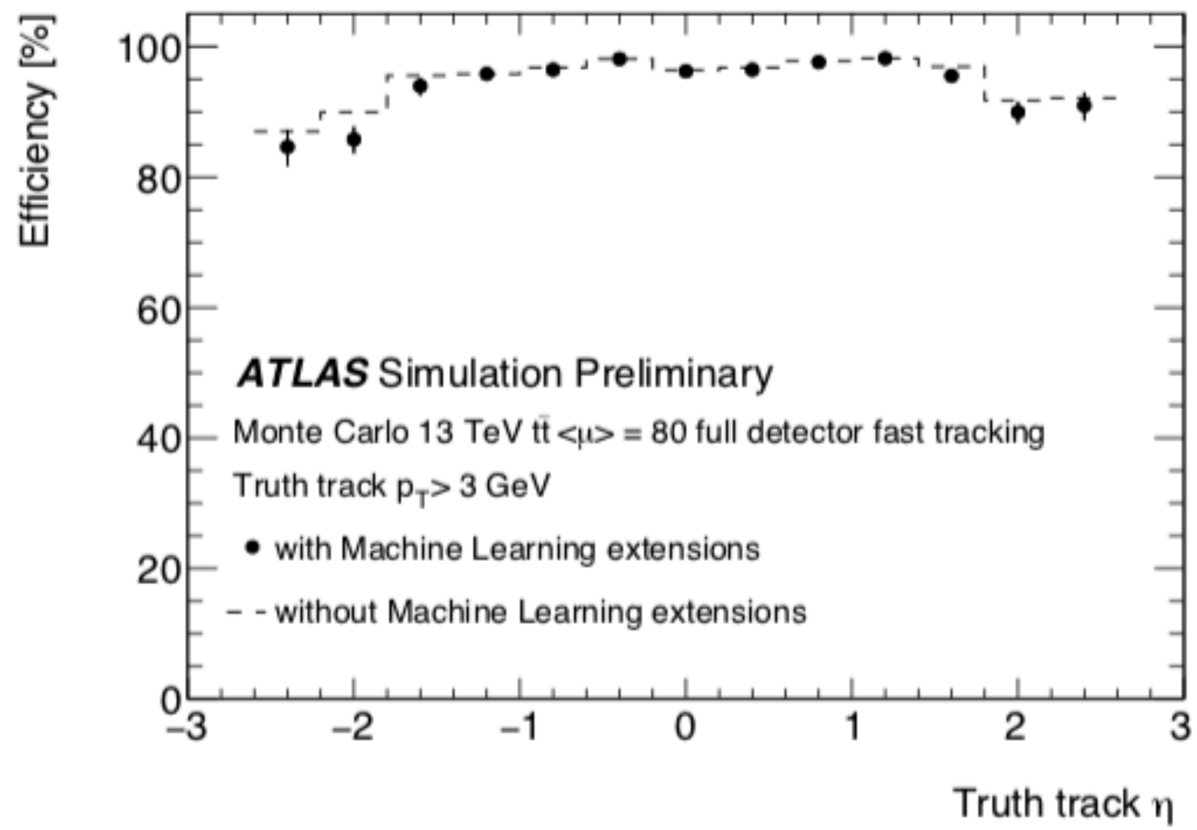

Total Speed-up Factor	Seed Generation	Seed Processing	Track Fitting
2.3×	1.3×	3.3×	1.5×

$<\mu>$	Efficiency Loss (%)	Total Speed-up Factor
40	0.7	1.6×
60	0.7	2.1×
80	1.1	2.3×

For use in the trigger: trained predictor implemented in Look Up Tables (LUT)

There will be new triggers…

All graphics from H. Russell's slides, HEP Software Foundation Trigger & Reco WG, 2019

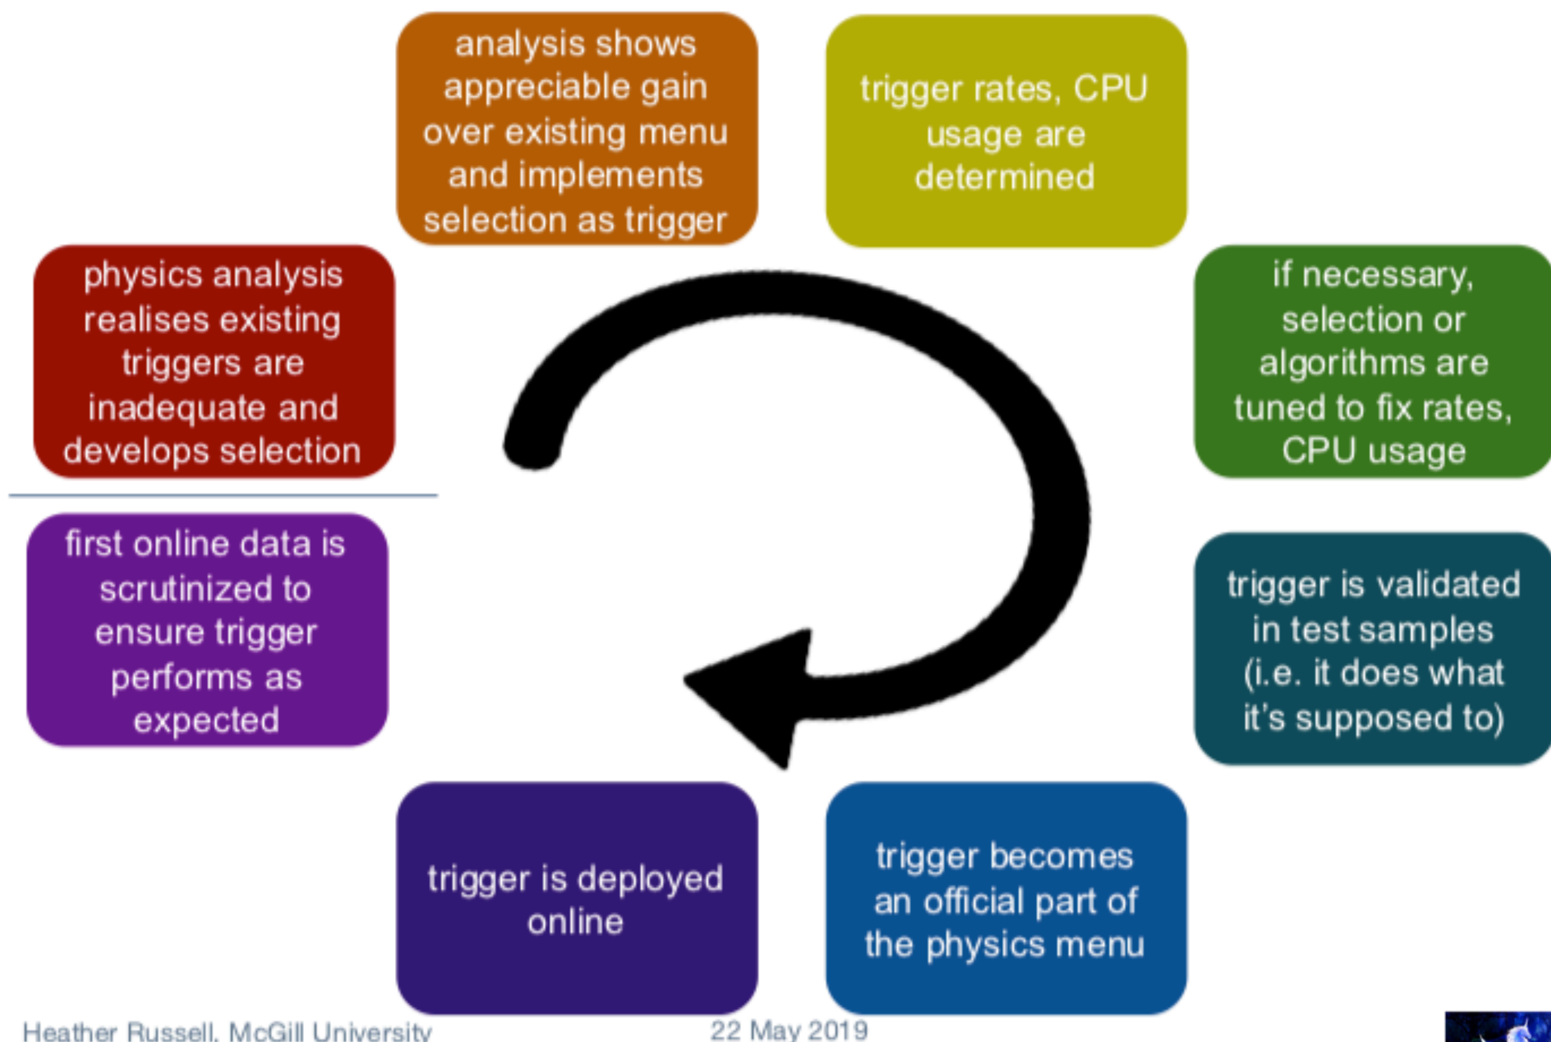

Now is the time to request trigger to record data in Run-3!
(good physics motivations & theory/experimental cross-talk always welcome)

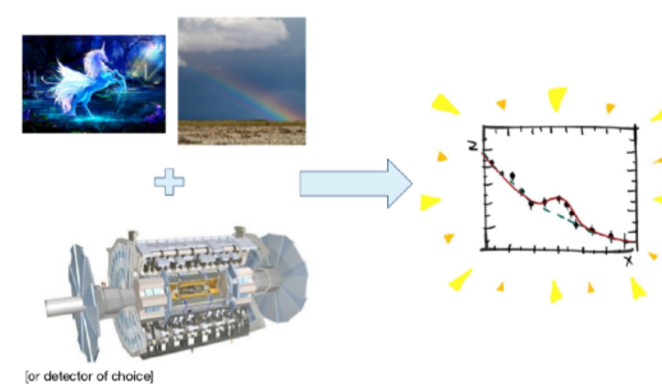

Non-standard workflows towards Run-3

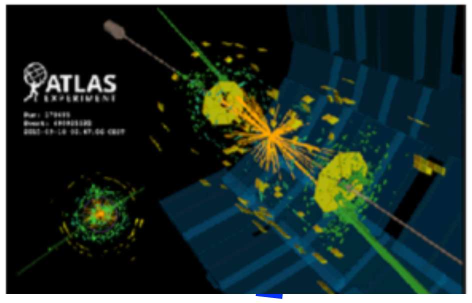

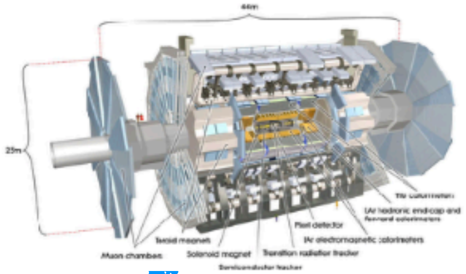

Online ◀ ▶ Offline

Event selection (trigger) ▶ Object reconstruction and calibration ▶ Data analysis

Data
I like

What if I really want this data?

Data I would have used to like if I had been a data hipster

Data
you like

What we can record in full after the trigger

Data ~nobody likes

Data produced by the LHC
(multiplied by large number)

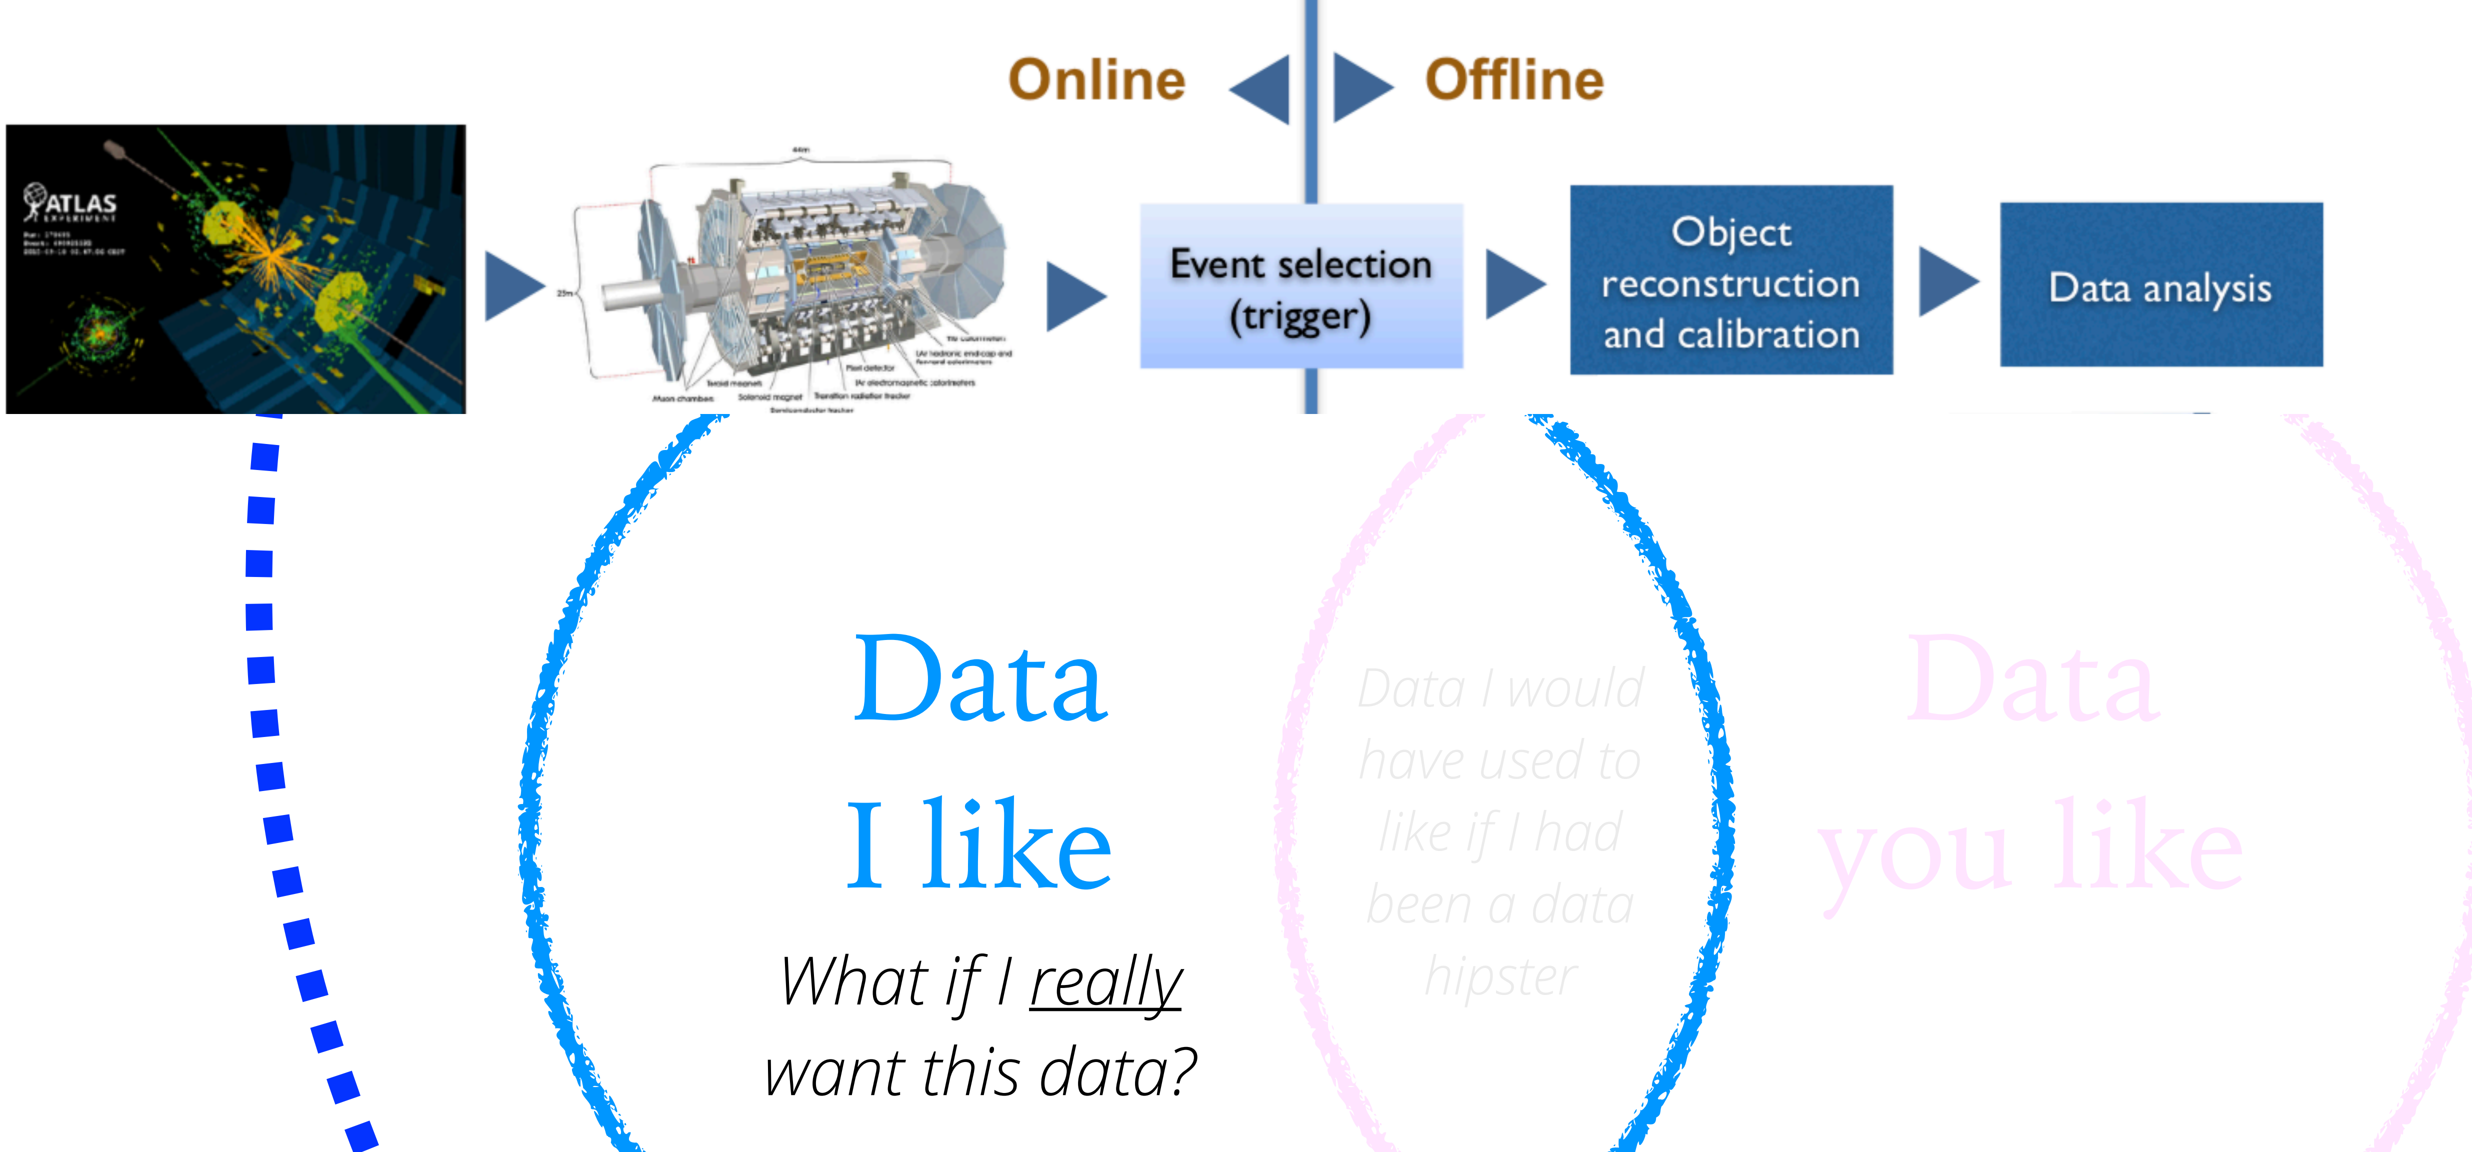

1. Analyze as much data as possible, **as fast as possible**
- **Real-time analysis** (*Trigger-Level Analysis / Data Scouting / Turbo Stream*)

2. Save a little **less data** per each event
- A prerequisite of real-time analysis, but can also save more (*Partial Event Building / Selective Persistency*)
- Requires making **hard choices** on what information to keep for further analysis

3. Save data for **further reconstruction, later**
- Periods between data taking can be long…reconstruct when offline CPU available again (*Delayed stream / Data Parking*)
- Byproduct: make data & workflows FAIR and sustainable!

A paradigm change, started in earlier runs

Turbo stream (LHCb)
Data Scouting (CMS)

Asynchronous data analysis

First record data, then reconstruct/analyze it

↓

Real-time data analysis

Reconstruct/analyze data as soon as it is read out so that only (**smaller**) final-state objects or histograms need to be stored

Regular trigger & data analysis path

Turbo/Data Scouting/TLA path

Turbo stream (LHCb),
Data Scouting (CMS),
Trigger-level Analysis (ATLAS),

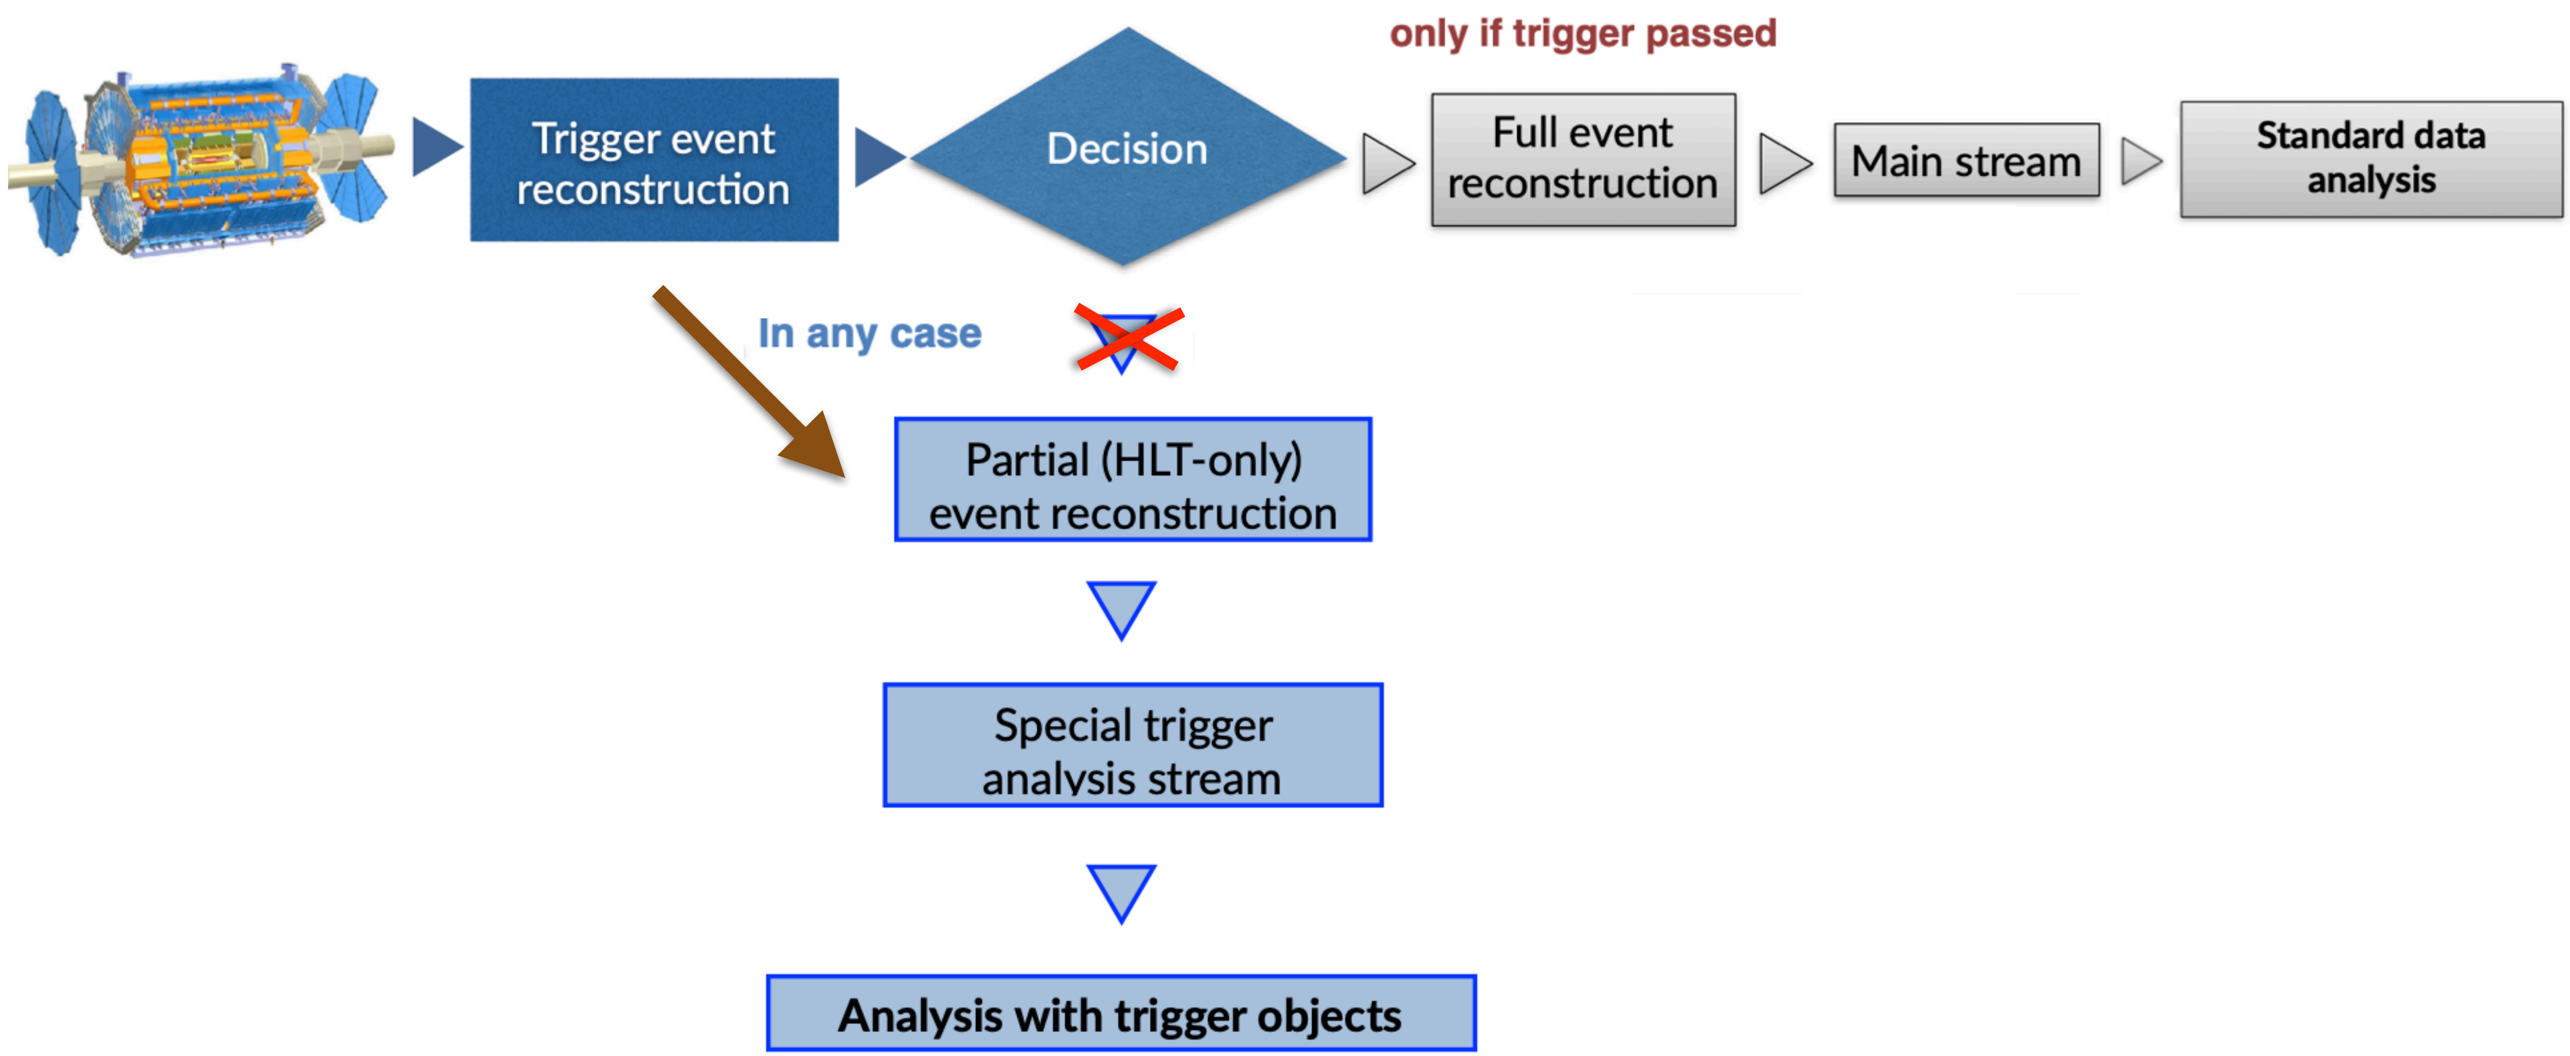

(Near-)real-time analysis of LHC data

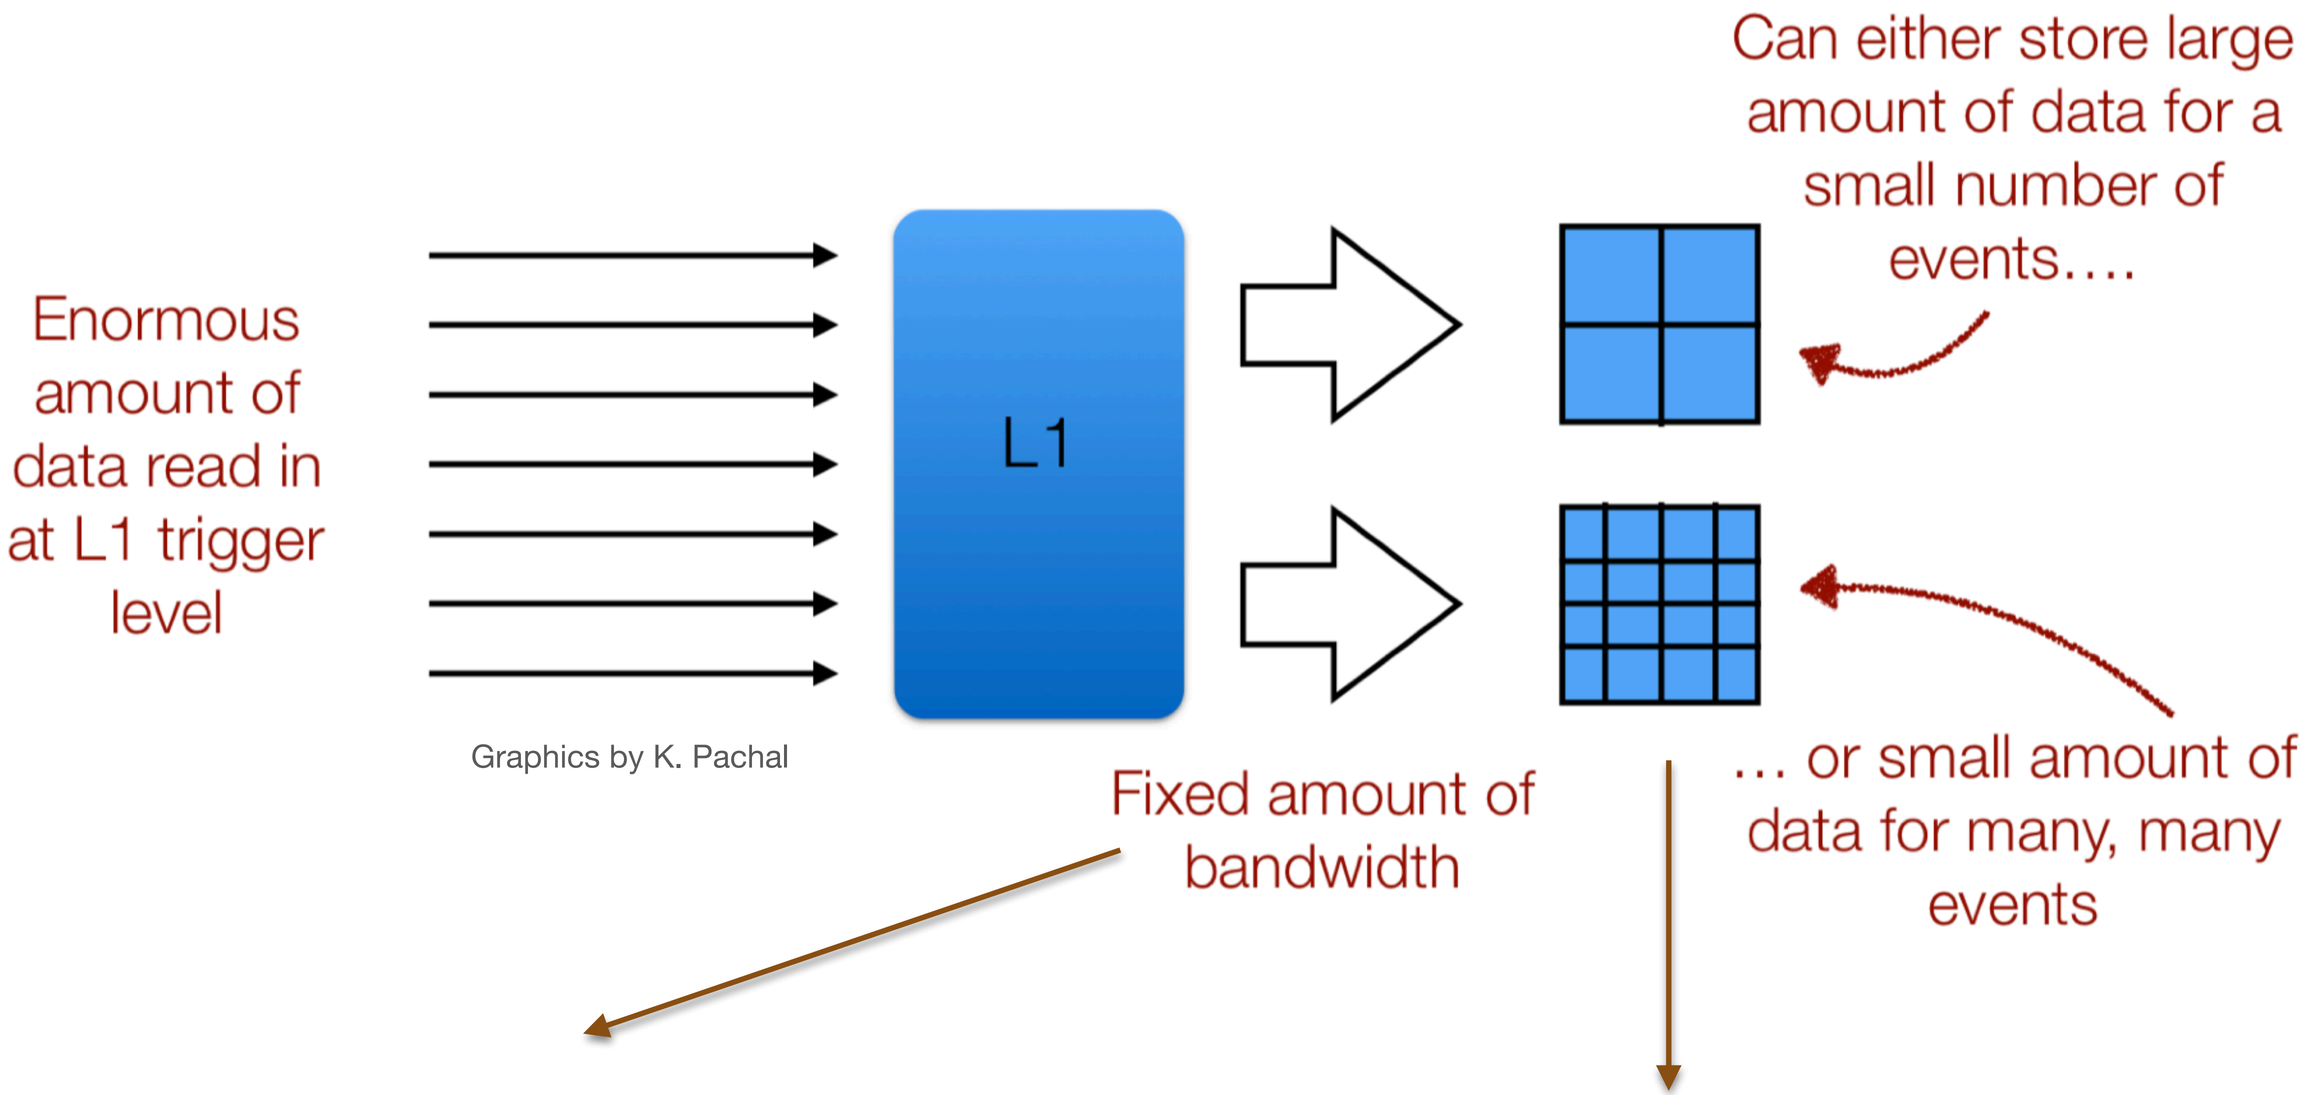

Perform as much "analysis" as possible @ HLT

- Reconstruction & calibration
- First preselection to skim "backgrounds"

Reduced data formats:

- Only keep final trigger objects (drop raw data)
- Save only "interesting" parts of the detector
- A combination of the two

Overcoming storage (and CPU) bottlenecks

Save many more smaller events

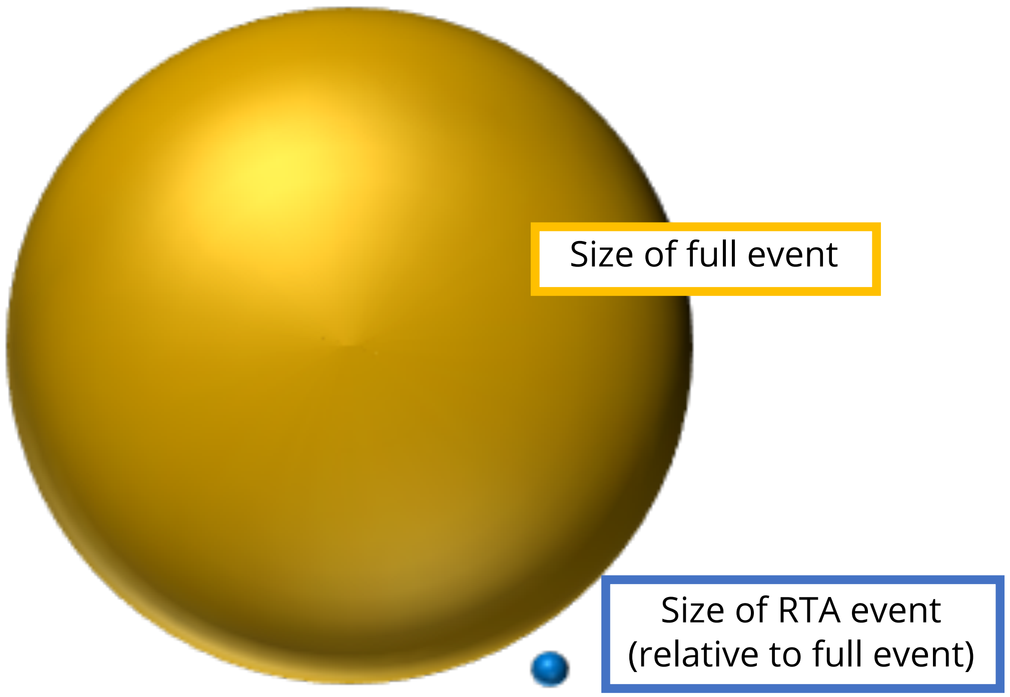

TRIG-2019-04

Stream	Average event size
Physics, express	1 MB
Trigger-level analysis	6.5 kB
Calibration	1.3 kB to 1 MB
B-physics and light states	1 MB

- Allows to record and store much higher event rates

Overcoming storage (and CPU) bottlenecks

Save many more smaller events

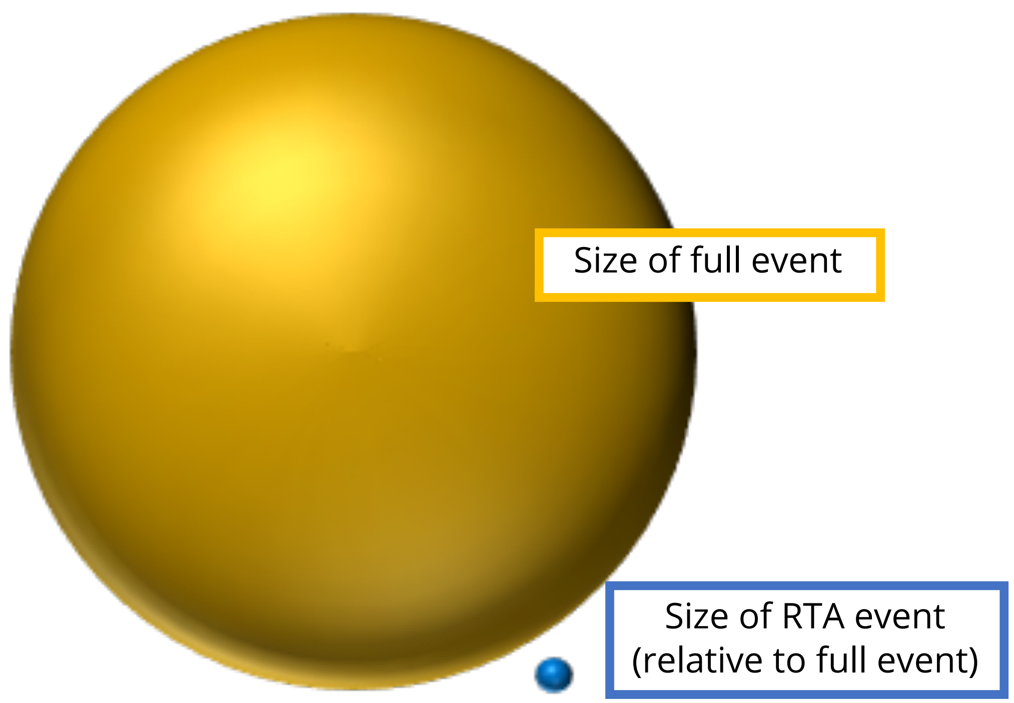

TRIG-2019-04

Stream	Average event size
Physics, express	1 MB
Trigger-level analysis	6.5 kB
Calibration	1.3 kB to 1 MB
B-physics and light states	1 MB

- Allows to record and store much higher event rates

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults

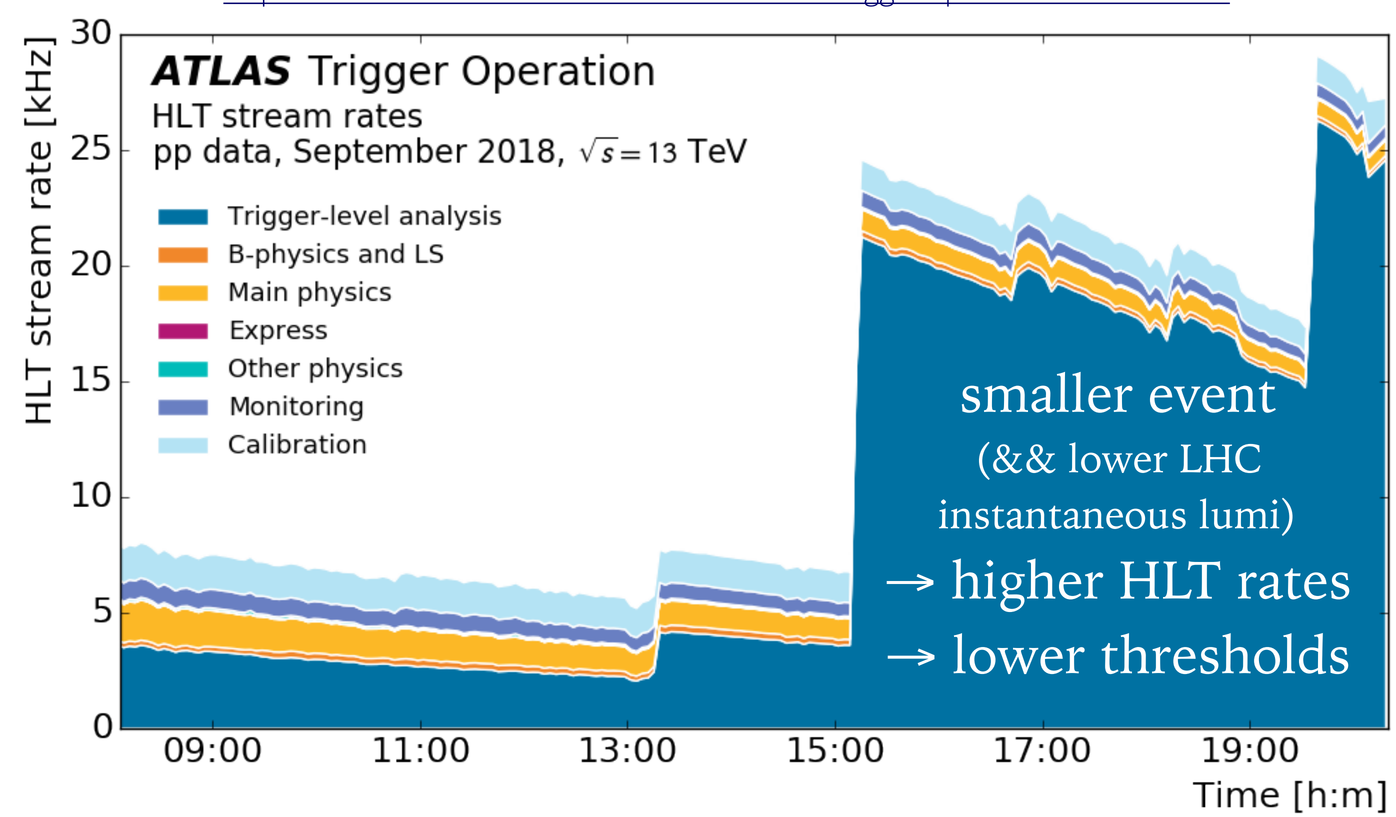

Use all the CPU, all the time

- LHC end-of-fill ➞ unused HLT farm nodes
- Can lower the HLT thresholds to record more trigger-level events
  - Note: this does not work with lumi-leveling, so not clear this will happen in Run-3

Run-3 plans: extend to physics objects beyond jets

More with less: Partial Event Building (=Selective

Real-time analysis is necessary for searches that would otherwise have been impossible due to trigger constraints

Traditional offline analysis still required for a number of searches/final states where all raw information is needed (but we could do better)

Partial Event Building / Selective Persistency as a middle way: save raw data && trigger objects only in the regions of interest, re-reconstruct later

Example of customizable outputs @ LHCb:

- keep trigger objects only (7 kB)
- **keep trigger objects + "on-demand" raw and/or reco in selected regions (< 200 kB)**
- keep everything (200 kB)

HSF Trigger & Reco / Institut Pascal discussion, July 2016: https://indico.cern.ch/event/835074/

H. Russell, EPS-HEP 2019,

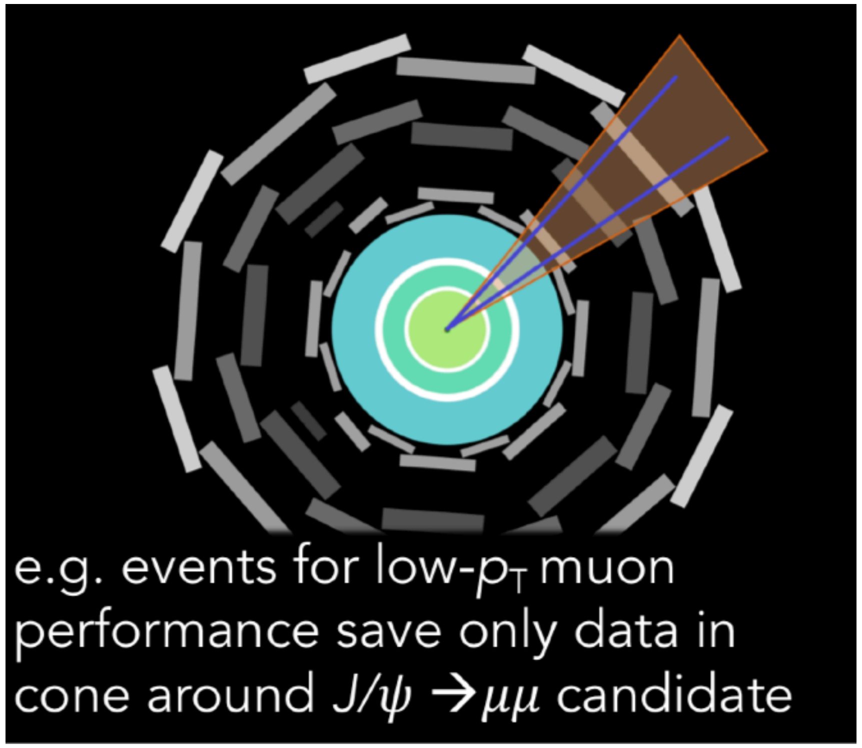

More, later: delayed streams (= data parking)

If **offline CPU availability** is the bottleneck to recording more data:
delay data reconstruction until LHC ends taking data and the (Tier-0) farm is free

Run-1: delayed stream (HT>500 GeV) brought large advantages for hadronic searches

- Dijet resonances (as a precursor to TLA): arXiv:1407.1376
- RPV stops (with b-tagging): arXiv:1601.07453
- Also used for 2012 jet energy scale derivation

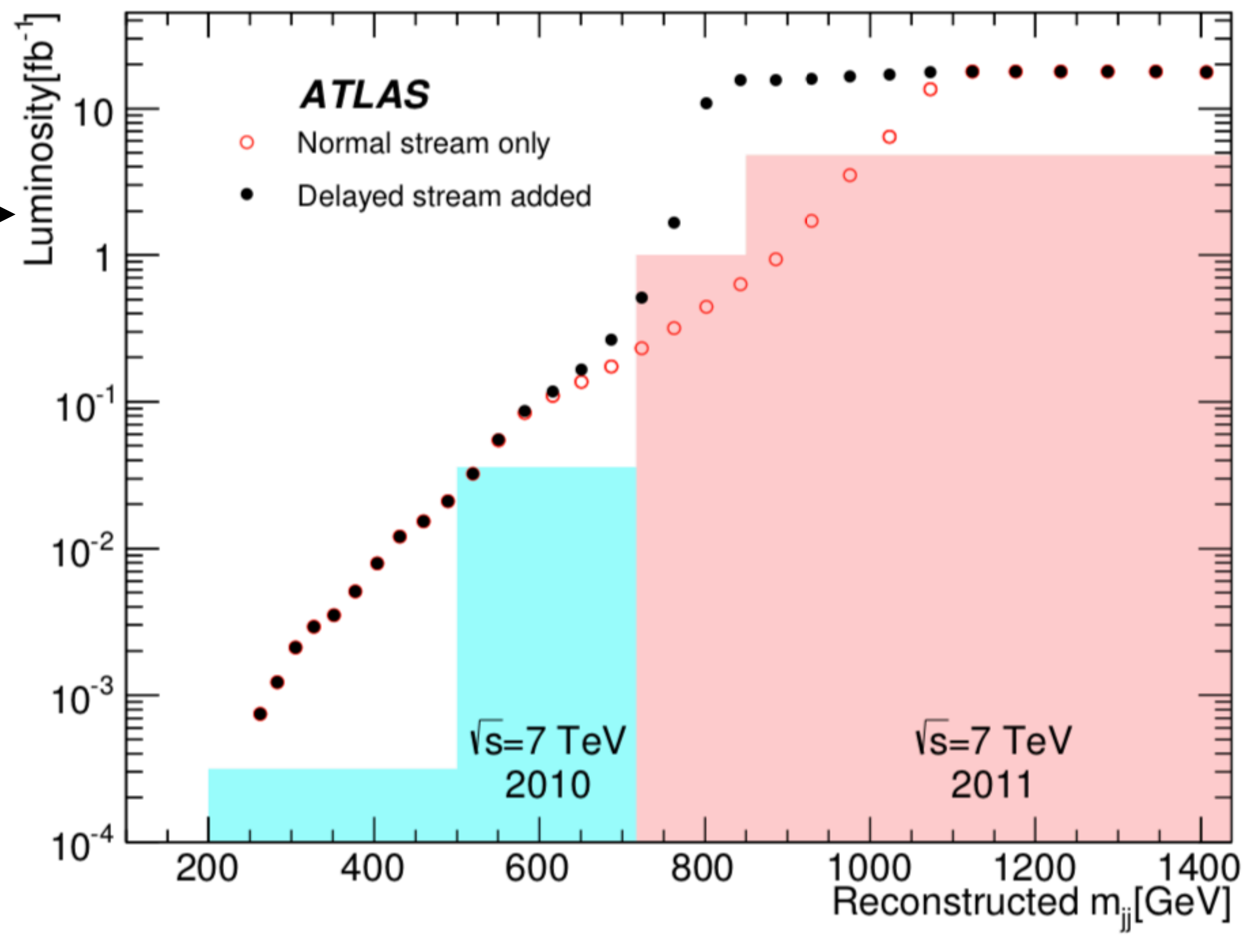

Run-2: also available as "safety net" in case Trigger Level Analysis saw events

- A public answer to yesterday's questions about RK is in the next slide

Run-3: plans to expand use of delayed stream

More, later: delayed streams (= data parking)

TRIG-2019-04, ATL-DAQ-PUB-2019-001

depending on their primary use case and their specific offline reconstruction needs. Figure 1 shows the average recording rate of the physics data stream of all ATLAS *pp* runs taken in 2018. Events for physics analyses are recorded at an average rate of ~ 1.2 kHz.[2] This comprises two streams, one dedicated to *B*-physics and light states (BLS) data, which averaged 200 Hz, and one for main physics data, which averaged 1 kHz. The BLS data are kept separate so the offline reconstruction can be delayed if available resources for offline processing are scarce due to high LHC uptime.

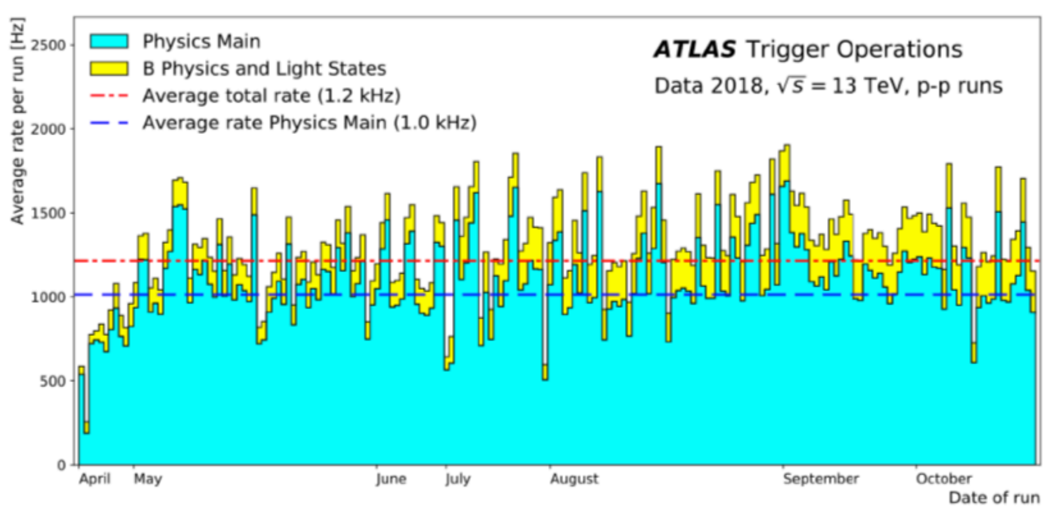

Figure 1: The average recording rate of the main physics data stream and the *B*-physics and light states data stream for each ATLAS *pp* physics run taken in 2018. The total average of all runs is indicated as a red dash-dotted line, and the total average of the main physics stream is indicated as a blue dashed line.

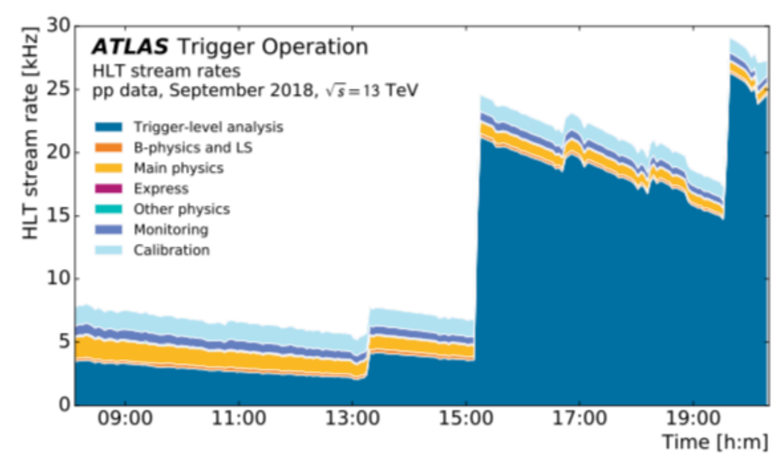

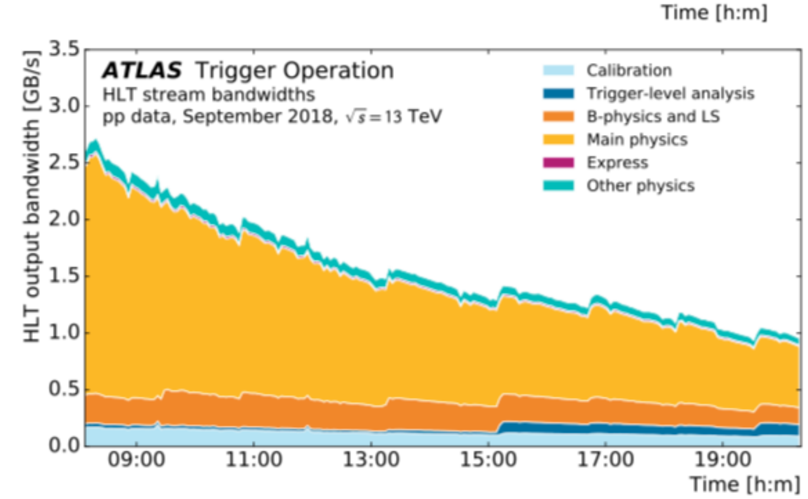

An example: some very exotic signatures

Mapping of exotic signatures to big picture of benchmark models not always easy → difficult to prioritize → may be difficult to decide what exactly to include in trigger menu

Signatures with a **common denominator**: unusual tracks/energy distributions, more or less localized in the detector

How do we make sure we don't miss these events?

increasing event rate

1. write dedicated trigger algorithms
2. save a mixture of raw data and trigger-level objects
3. save (custom-reconstructed) trigger-level objects only
4. save any of the above and reconstruct data later

increasing event size

5. [outlier detection]

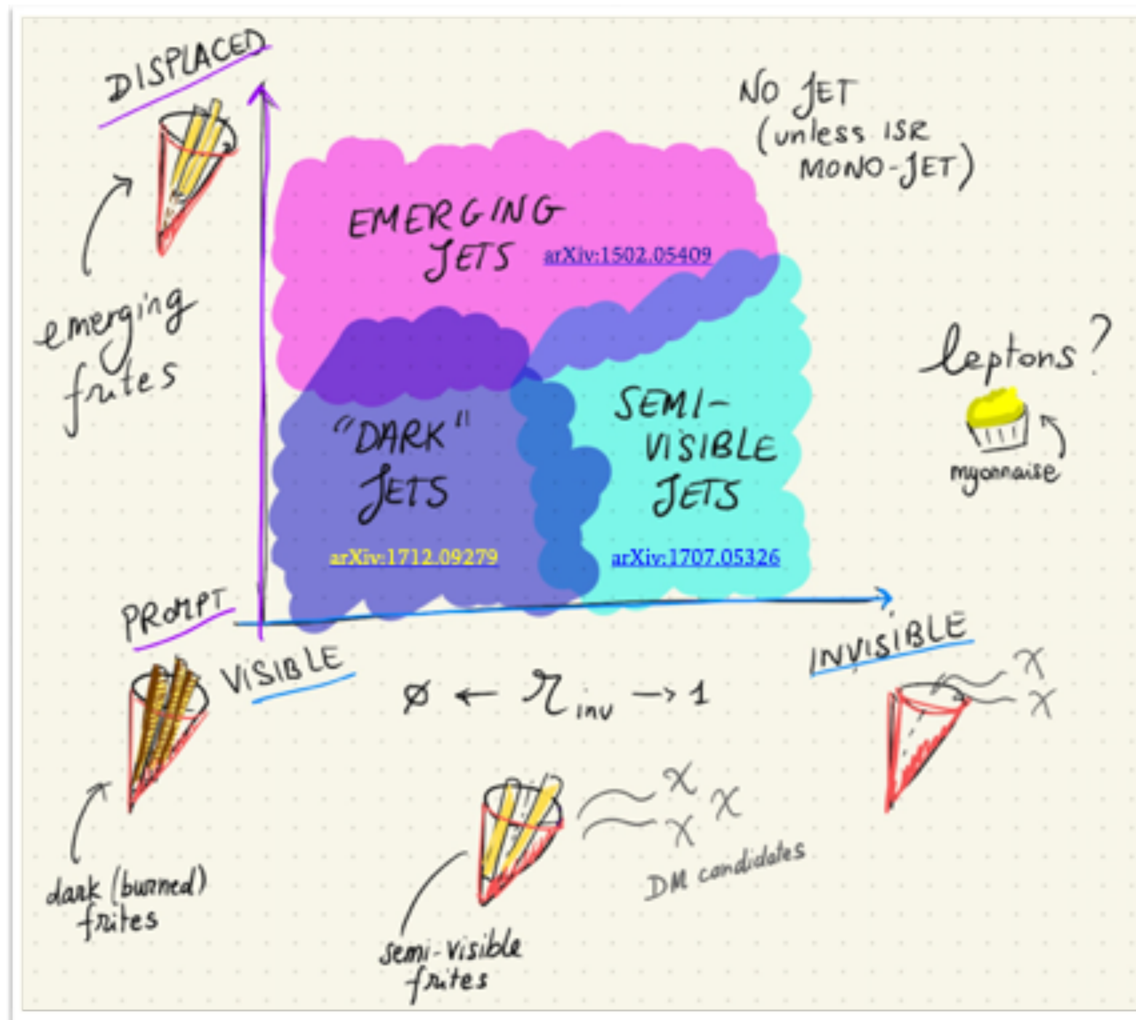

Inspired by K. Pedro & C. Fallon's talk @ DMLHC2019 and by this twitter thread

Conclusions

What is needed to operate the Run-3 trigger

Designing, implementing, operating…

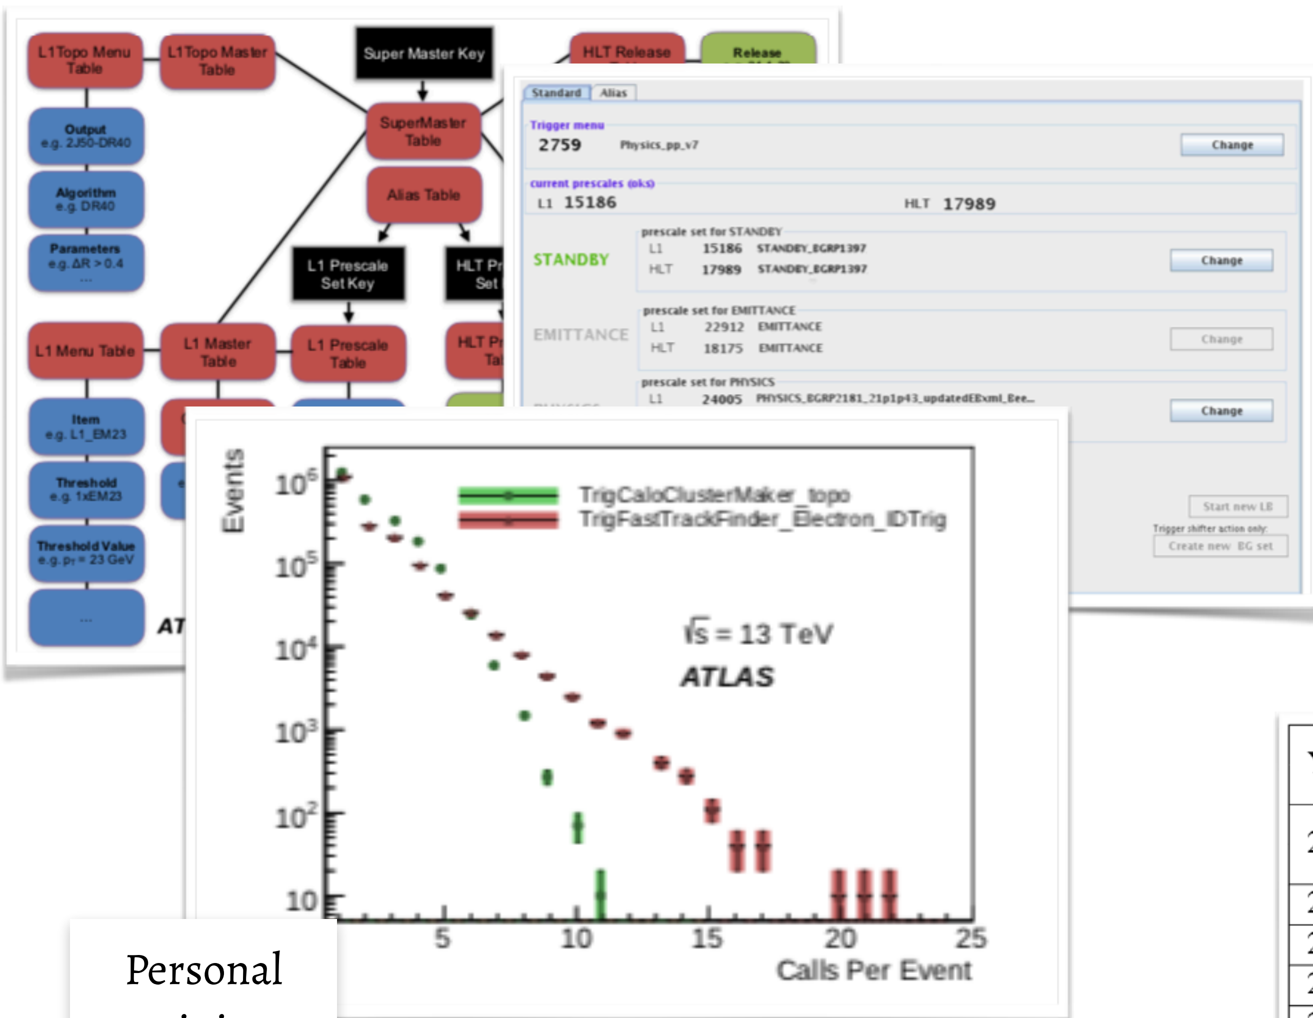

…and monitoring

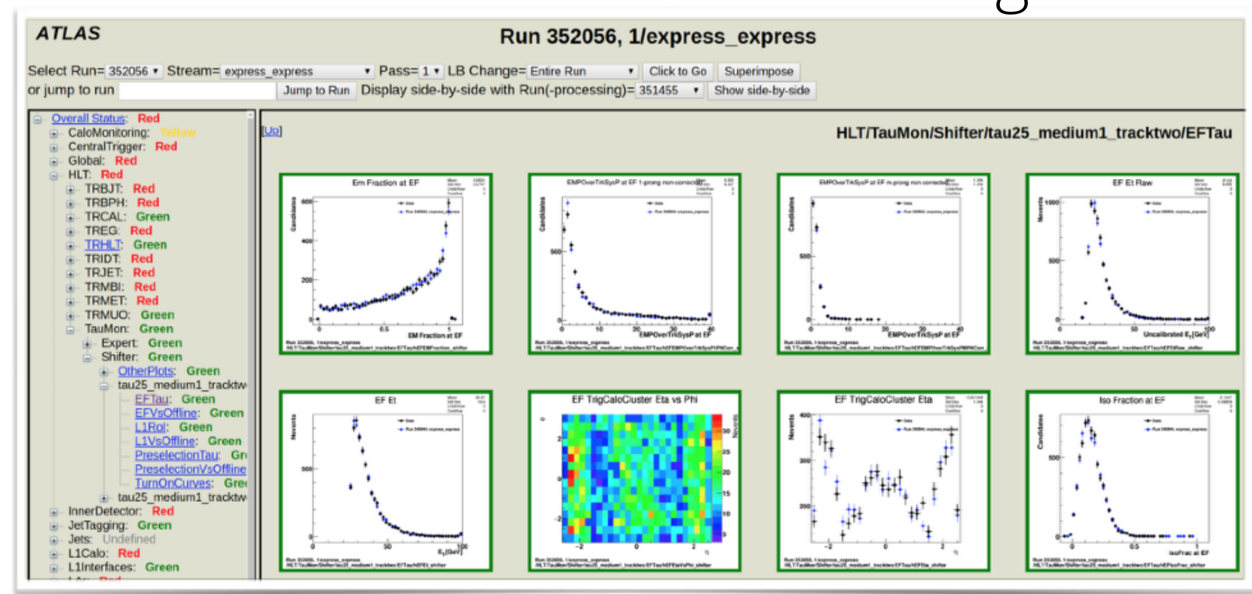

Year	Dataset	Trigger DQ Eff.		ATLAS DQ Eff.	Integrated Luminosity
		L1 [%]	HLT [%]	[%]	of good quality data
2015	*pp* @ 13 TeV (50 ns)	100.00	99.94	88.77	84 pb^{-1}
	pp @ 13 TeV	99.97	99.76	88.79	3.2 fb^{-1}
2016	*pp* @ 13 TeV	98.33	100.00	93.07	33 fb^{-1}
2017	*pp* @ 13 TeV	99.95	99.96	95.67	44 fb^{-1}
2018	*pp* @ 13 TeV	99.99	99.99	97.46	59 fb^{-1}
2015–2018	*pp* @ 13 TeV	99.57	99.94	95.60	139 fb^{-1}

arXiv:2007.12539, submitted to JINST

Personal opinion

Run-3 temptation: *"I'll get a hobby until we collect the entire dataset"*

In order to make the most of the Run-3 data, we need to make sure we dedicate enough experimentalist's time & funding & career prospects to technical / performance work

Conclusions

Interesting time for high energy collider physics:
we don't know what to expect from new physics
(but we have a prior: it should be *somewhere*)
we have the **LHC running now**,
and the data we discard is **gone forever**

Let's make the most of Run-3 LHC data
by recording as much *useful & (re)usable* data as possible!

***Closing note*: Real-time decision making (=trigger) challenges exist across all LHC collaborations and beyond HEP**
→ HEP Software Foundation trigger & reconstruction WG [Website] [Indico category] [Most recent whitepaper]

Backup slides

CATERINA DOGLIONI - LUND UNIVERSITY
@CATDOGLUND, SHE/HER
http://www.hep.lu.se/staff/doglioni/

LUNDS UNIVERSITET

European Research Council
Established by the European Commission

ATLAS EXPERIMENT

The ATLAS Run-2 trigger system