Resource usage and rate predictions in the ATLAS High Level Trigger

CHEP 2016 – San Francisco

Tim Martin

On behalf of the ATLAS Collaboration

University of Warwick

10 October 2016

Introduction: ATL-DAQ-PUB-2016-002

ATLAS operates a two level trigger system: **L1** & **HLT** (High Level Trigger) with up to 40,000 CPU cores running a menu of $\sim$ 1000 **chains** (algorithmic selections).

- This presentation will describe the ATLAS Trigger Cost-Monitoring Framework
  - Sampling of event execution data during HLT operation.
  - **Offline auditing** of resource usage/**'cost'** in the HLT (CPU & data requests).
- And the 'Enhanced Bias' mechanism.
  - Special data samples, **enriched in high-p_{T} events passing L1 trigger** which are likely to be selected by the HLT.
  - Taken such that the prescales* are **invertible**.
  - Used to **predict rates**

*Prescale factor p, ($p \geq 1$): Accept at random only $1/p$ events passing selection.

ATLAS Trigger Infrastructure

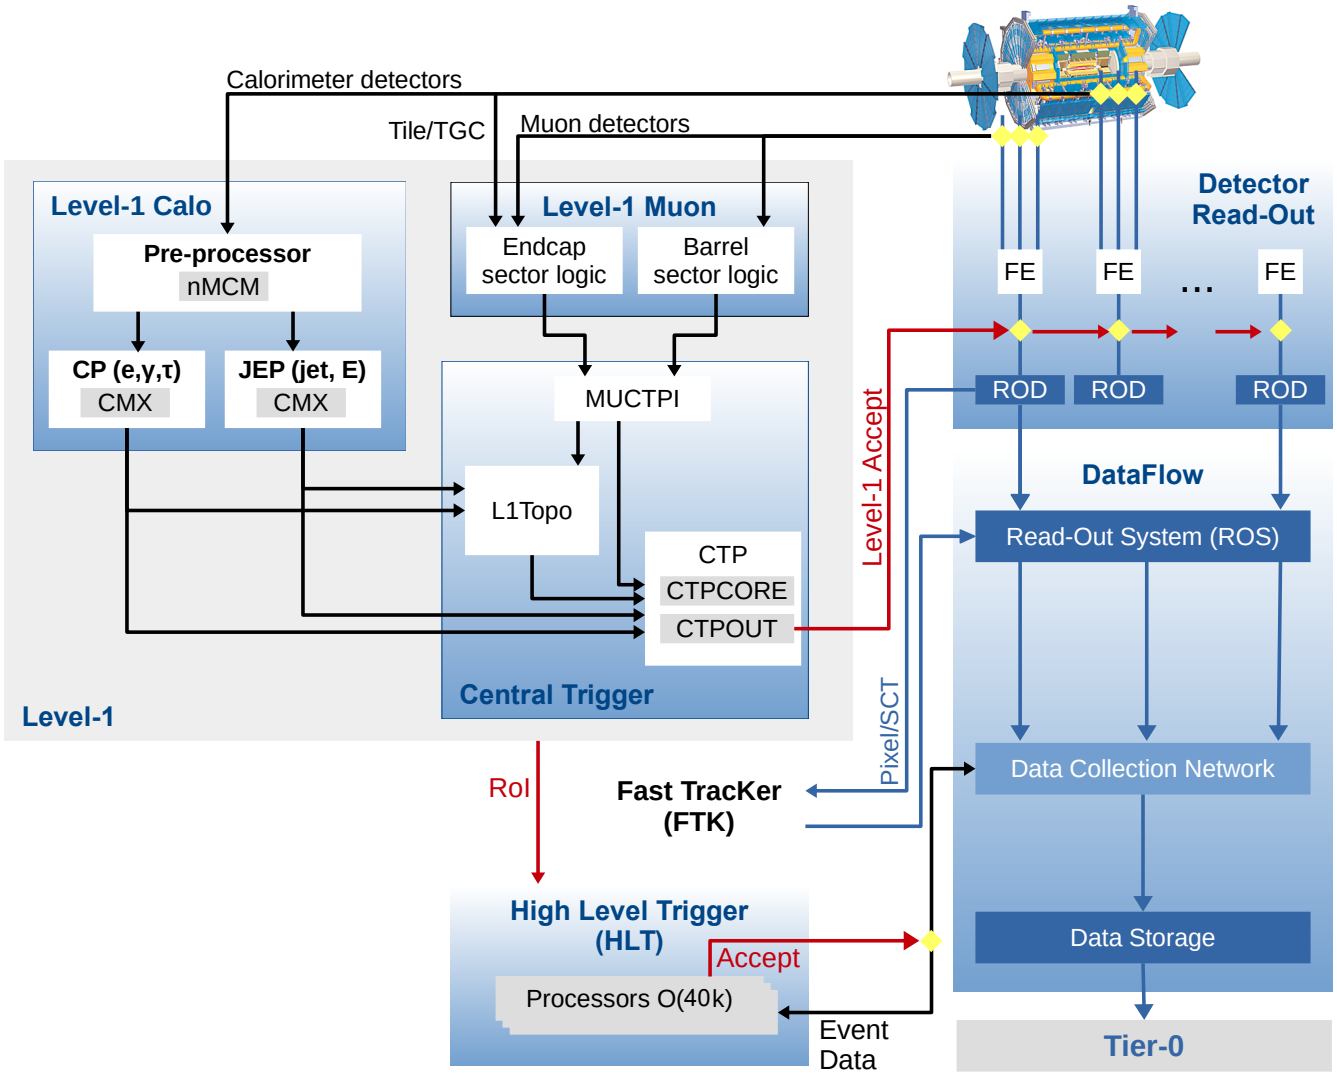

ATLAS Trigger Infrastructure

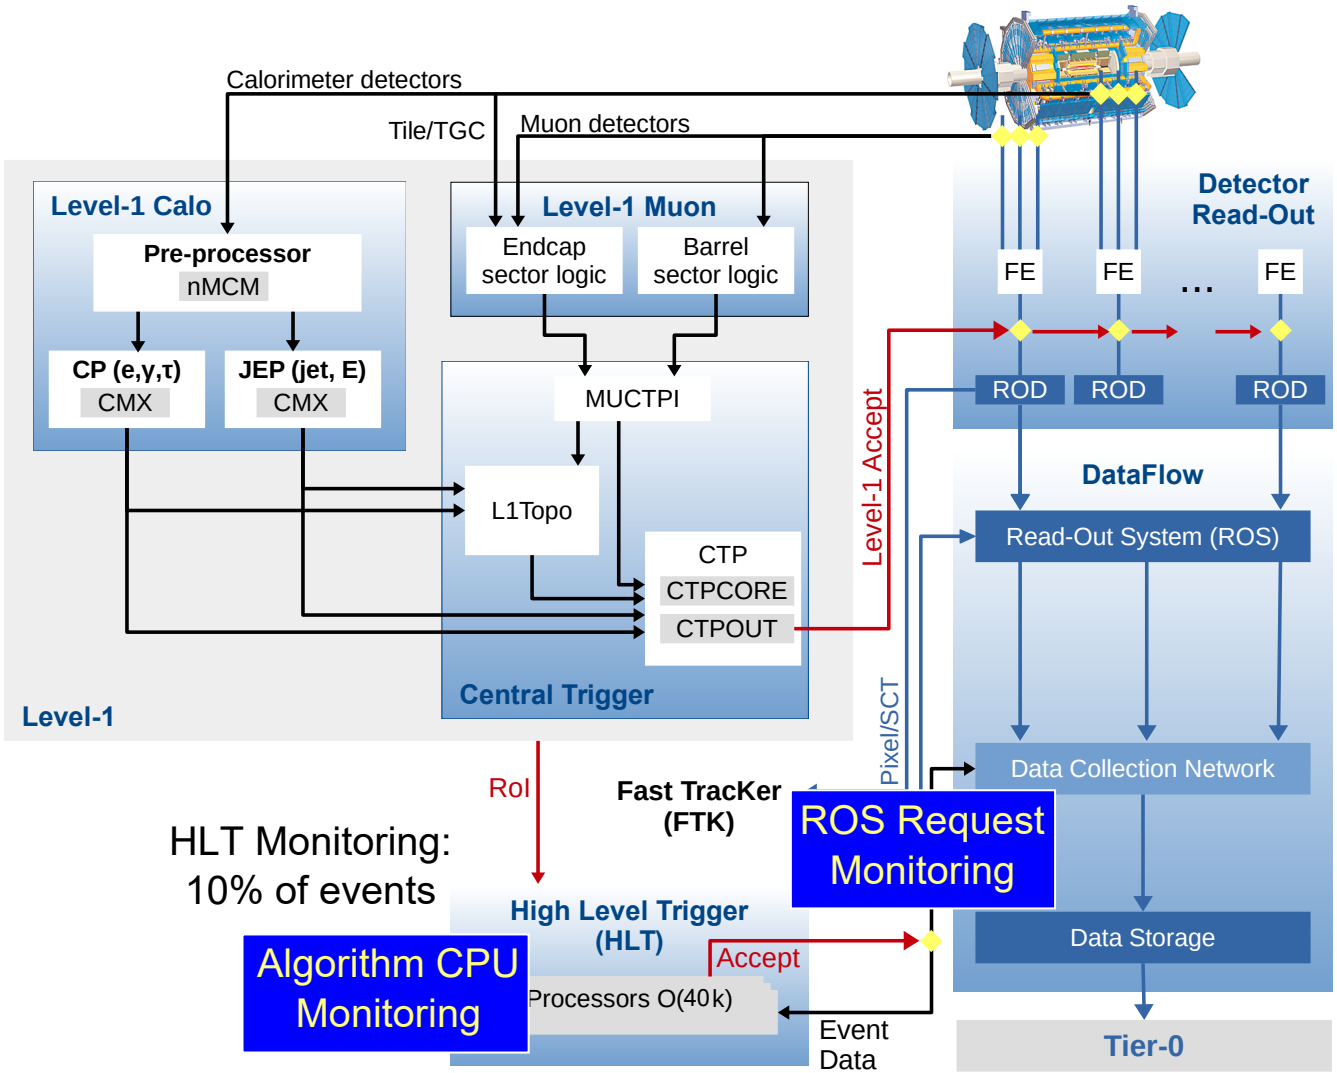

Processing HLT Cost-Monitoring Data

- 10% of HLT executions are monitored, at random (up to 10 kHz).
- Monitoring data buffered locally and bulk-exported (up to 800 Hz).
- Typical long run will generate **2.75 TB** of cost-monitoring data.
- HLT cost-monitoring data are deserialised and converted into ROOT ntuples in CERN's Tier-0 computing facility.
- Cost-monitoring framework processes the data **within 24h**
  - Set of '`Monitors`' of **high**-level and **low**-level information.
  - Outputs around 500,000 histograms.
  - Outputs tables of data in CSV, plus meta-data.

Cost-monitoring data are automatically made available to the collaboration via a web portal.

Processing HLT Cost Monitoring Data

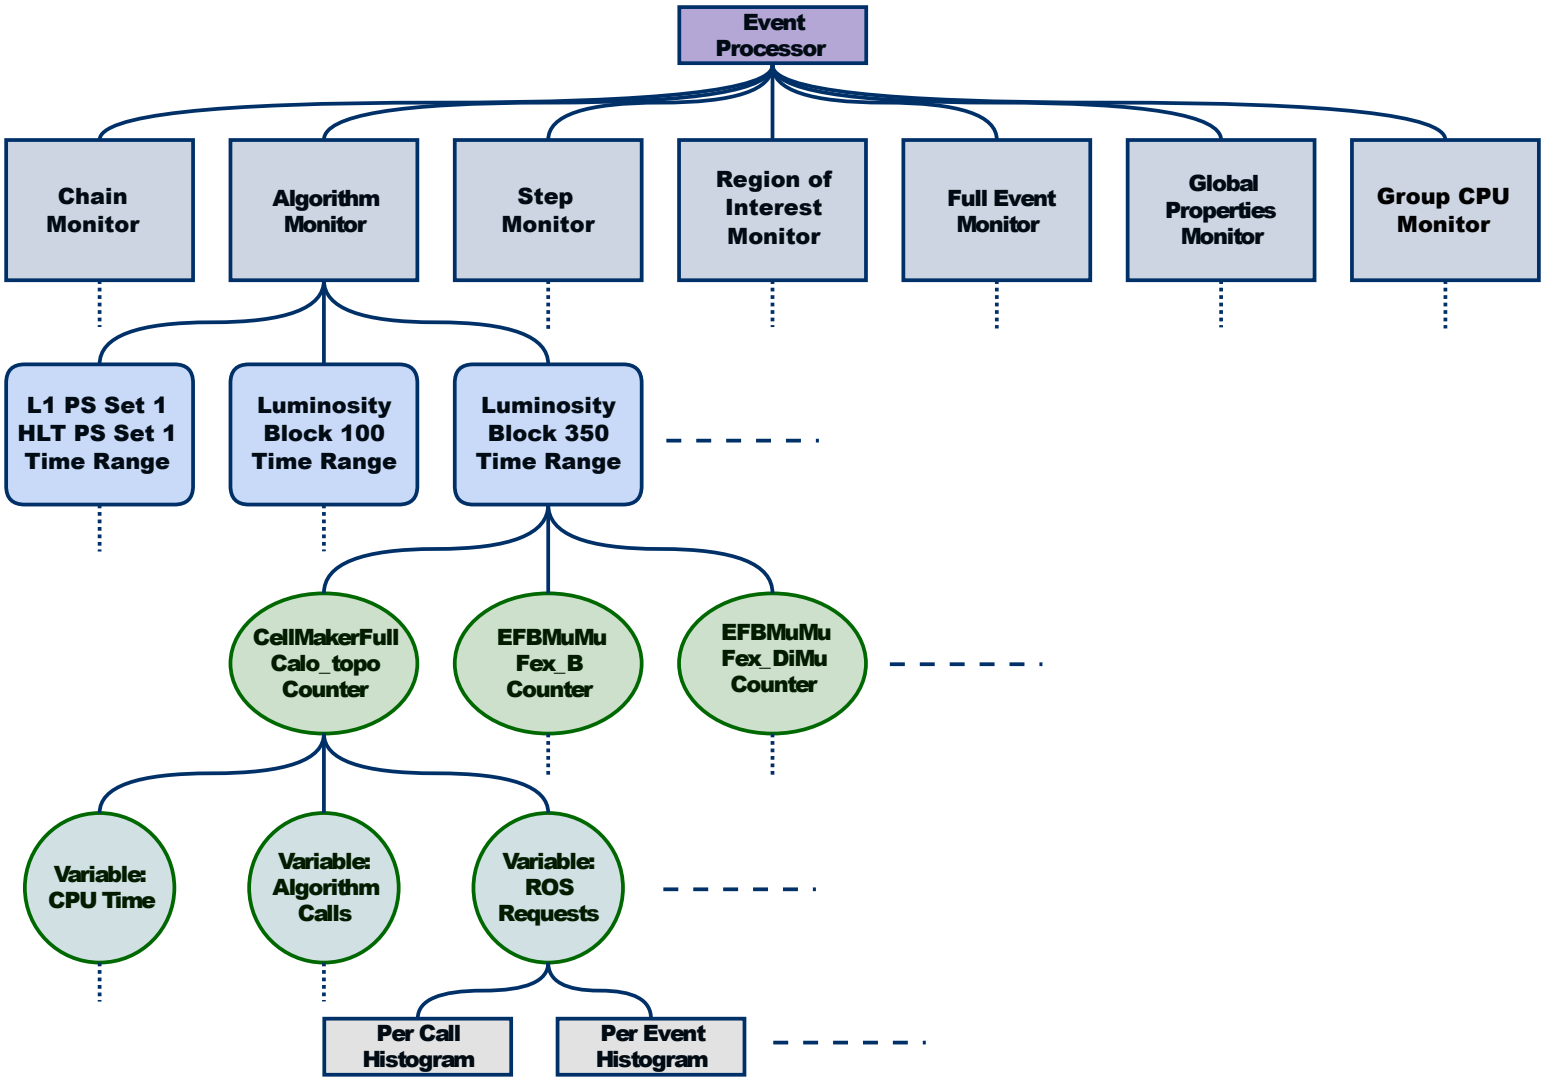

Trigger Chain Group CPU Usage

Fraction of total CPU use per physics group.

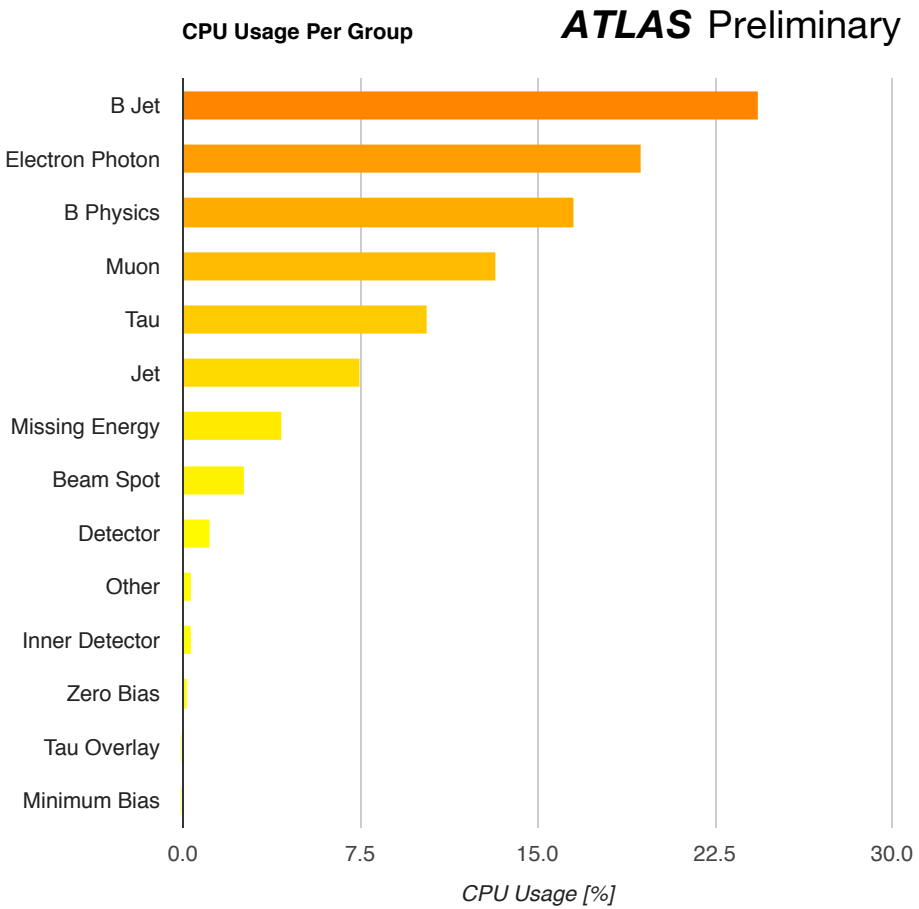

In total, 62% of 40,000 CPU cores are utilised in processing 77 kHz input from L1.

Data collected over 180 seconds at $\mathcal{L} = 1 \times 10^{34}$ cm^{-2}s^{-1}.

Algorithm CPU Usage [PUB-TRIG-2016-02]

- Low-level monitoring of individual feature extraction algorithms:
  - Calorimeter clustering within Region-of-Interest or full calorimeter $\rightarrow$ **double peak structure**.
  - Region-of-interest electron tracking.

CPU time per algorithm-call, and algorithm CPU time per-event.

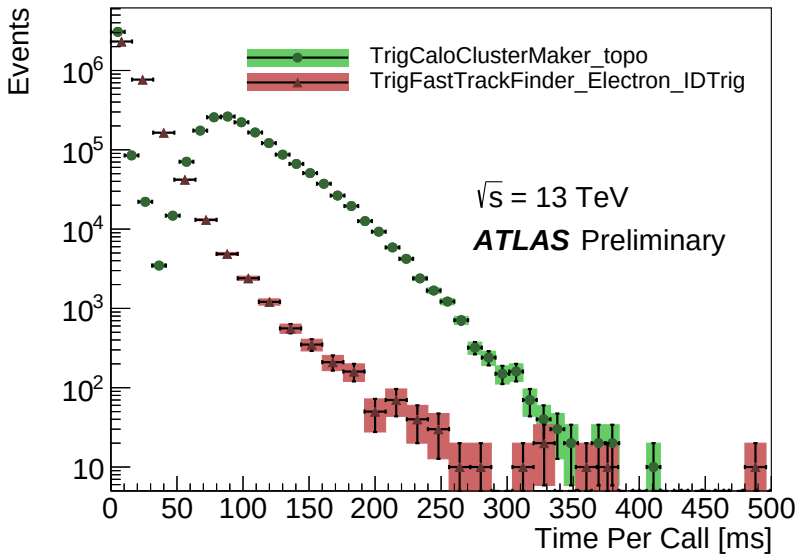

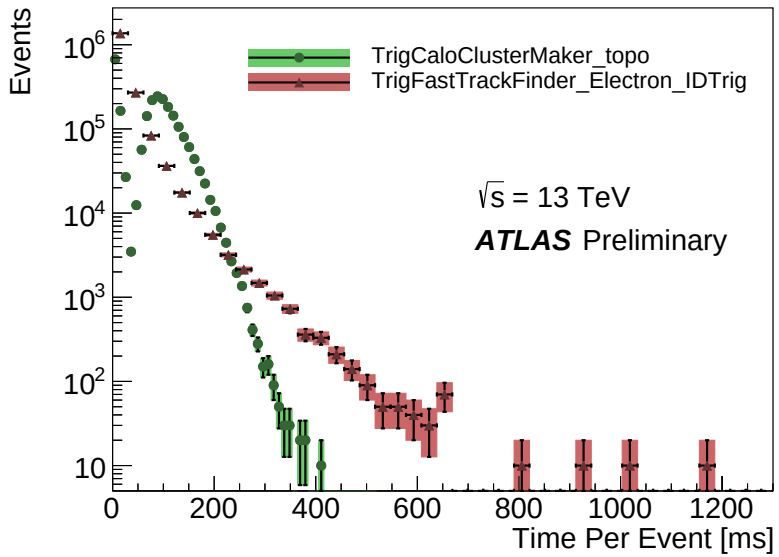

Algorithm CPU Usage [PUB-TRIG-2016-02]

- Low-level monitoring of individual feature extraction algorithms:
  - Calorimeter clustering within Region-of-Interest or full calorimeter $\rightarrow$ **double peak structure**.
  - Region-of-interest electron tracking.

Calls per event (Regions-of-Interest), and fractional CPU usage of all algorithms in the event.

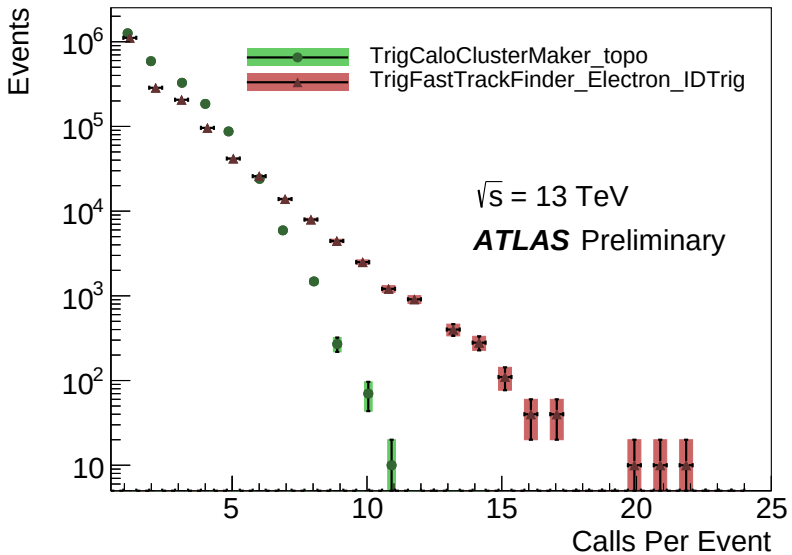

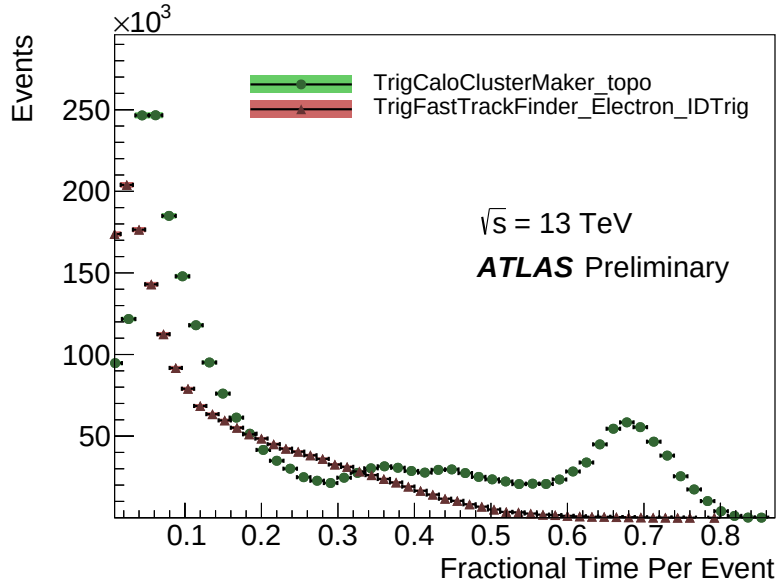

Predicting Trigger Rates

With Enhanced Bias Data

Taking Enhanced Bias Data

- A special set of chains are enabled recording 300 Hz extra events (on top of regular physics) for 1h.
- 1×10^6 Enhanced Bias (EB) events are collected (only L1 bias).
- Each chain targets physics at a different *rate* from low to high.
- EB chains have singular prescale values at L1 and the HLT
  - For low-p_T, prescaled L1 items, a ~ 5 kHz random L1 trigger is used and the L1 decision is inspected by the HLT.
  - 'L1 Seed Rates' below assume $\mathcal{L} = 1 \times 10^{34}$ cm^{-2}s^{1}

Name	**Seeding**	**Output [Hz]**	**L1 Seed Rates [kHz]**
Random	Random	60	> 500
Low	Random	60	50–500
Medium	Random	60	20–50
Primary	Direct	110	0.1–20
High	Direct	10	< 0.1

Taking Enhanced Bias Data

- A special set of chains are enabled recording 300 Hz extra events (on top of regular physics) for 1h.
- 1×10^6 Enhanced Bias (EB) events are collected (only L1 bias).
- Each chain targets physics at a different *rate* from low to high.
- EB chains have singular prescale values at L1 and the HLT
  - For low-p_T, prescaled L1 items, a ~ 5 kHz random L1 trigger is used and the L1 decision is inspected by the HLT.

Random: Unbiased low p_T jets, minimum bias

Low: 6 GeV μ, 15 GeV jets, 2×7 GeV e, 12 GeV τ

Medium: 15 GeV μ, 50 GeV jet, 35 GeV E_T^{miss}, 30 GeV τ

Primary: 20 GeV μ, 4×15 GeV jets, 22 GeV e, 60 GeV τ

High: 400 GeV jet, 6×15 GeV jets, 80 GeV E_T^{miss}

Predicting Trigger Rates with EB Data

- Rate predictions on **arbitrary selections** are possible with EB data.
- The new chain is defined and its selection is applied on all $N = 1 \times 10^6$ events yielding a raw result $r_e = 0, 1$ in event e.
- The rate of the new chain, c, is:

$$\mathrm{Rate}(c) = \frac{\sum_{e=1}^{N} w_{\mathrm{EB}}(e) w_c(e)}{\Delta t}$$

- Here $w_{\mathrm{EB}}(e)$ is the **enhanced bias weight**, it undoes the prescales used to collect the EB data.
- $w_c(e)$ Depends on the chain's **raw result** r_e and any **prescales** to be simulated.
- Δt is the **time period** of the EB dataset, typically 1 hour.

Calculating the Weights

- The **Enhanced Bias** weight $w_{\mathrm{EB}}(e)$, is a property of the **event**.
- It is calculated for event e with the **raw decision** r_{je} and **prescale** p_j of the $j = 1, 2, 3, 4, 5$ Enhanced Bias chains.

$$\frac{1}{w_{\mathrm{EB}}(e)} = 1 - \prod_{j=1}^{\text{EB Chains}} \left(1 - \frac{r_{je}}{p_j}\right)$$

- Different formulations of the chain weight $w_c(e)$ yields:
  - The **rate of a single chain** c
  - The **total rate of parallel chains (OR)** $c \cup d$
  - The **combined rate of parallel chains (AND)** $c \cap d$.

$$w_c(e) = \frac{r_{ce}}{p_c} \quad \Big| \quad w_{c\cup d}(e) = 1 - \left(1 - \frac{r_{ce}}{p_c}\right)\left(1 - \frac{r_{de}}{p_d}\right) \quad \Big| \quad w_{c\cap d}(e) = \frac{r_{ce}\, r_{de}}{p_c\, p_d}$$

Predictions vs. Reality

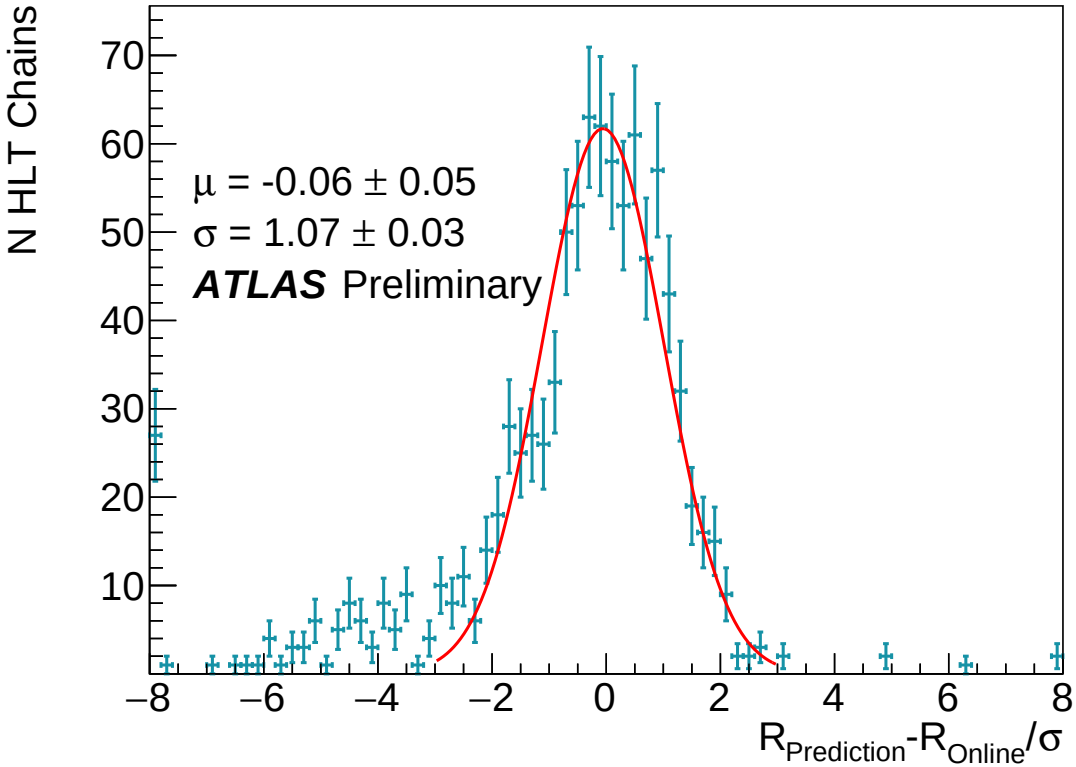

Comparison of 957 HLT chains rate prediction with their online rates. The predictions are normally distributed given their statistical error.

With a mean fractional error of 10% (predictions) and 2% (online)

Predicted Trigger Rates

Predicted L1 rates scan via L1 emulation, including statistical error.

Mechanism allows for quick rates evaluation of arbitrary selections.

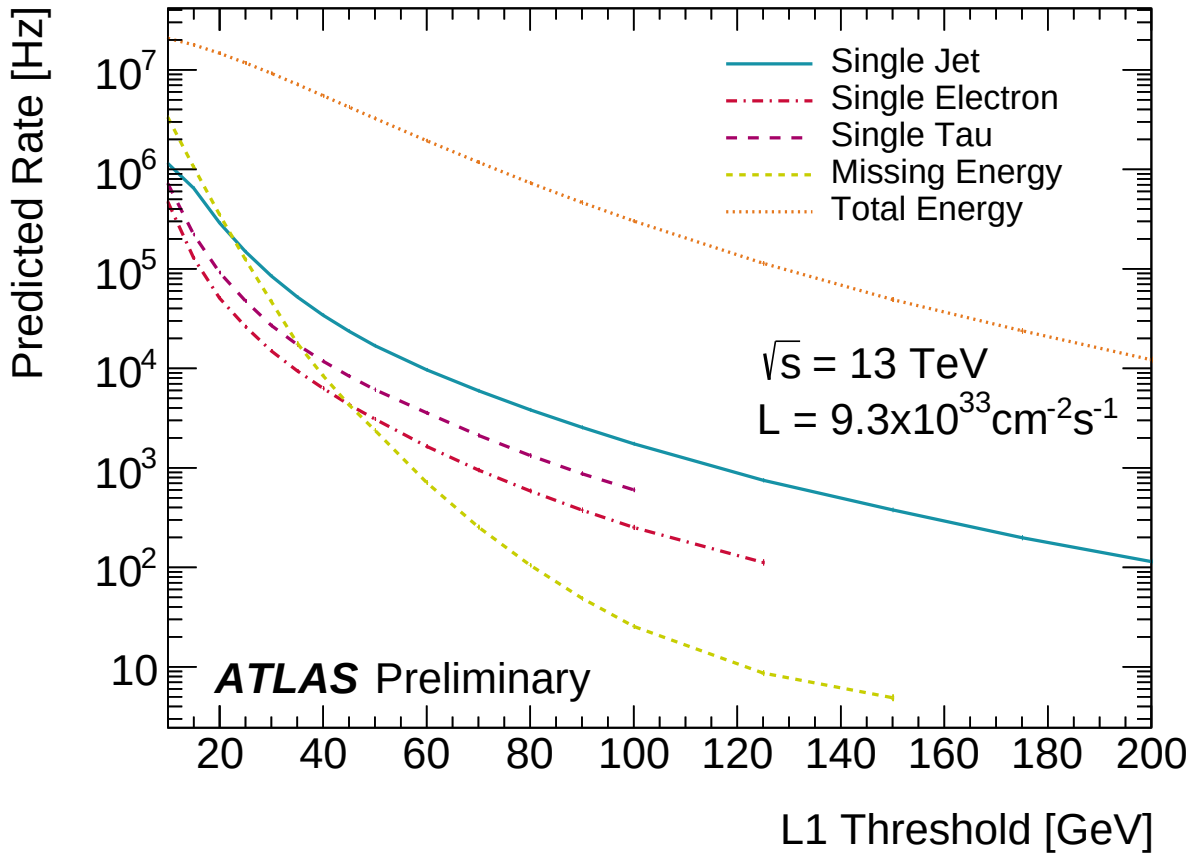

Conclusions: ATL-DAQ-PUB-2016-002

Detailed monitoring data from the ATLAS High Level Trigger are processed automatically within 24h of run finish & made available to the collaboration via web portal.

Offline execution monitoring of the HLT at a **high-level**, including the *total execution time* and at a **low-level**, including *per-algorithm monitoring*. Allows for optimisations, monitoring & future planning.

The 'Enhanced Bias' mechanism allows for compact datasets ($\sim$ 1 million events) with the statistical power to evaluate rates for arbitrary HLT selections.

Rate predictions are calculable for individual trigger chains, groups of chains, the total rate, unique rates, overlaps between triggers.

Validate an entire trigger menu **before** it is deployed on the live system.