LHCC Poster Session - CERN, 13 March 2013

ATLAS TDAQ System Performance in 2012 Data Taking

The challenge: running conditions

In 2012 LHC delivered four times the integrated luminosity it provided in 2011, at the cost of much harsher operating conditions.

	LHC design	2011 (peak)	2012 (peak)
Bunch spacing	25 ns	50 ns	50 ns
Colliding bunches	2808	1368	1368
Average primary vertexes per bunch crossing (pile-up)	25	17	**37**

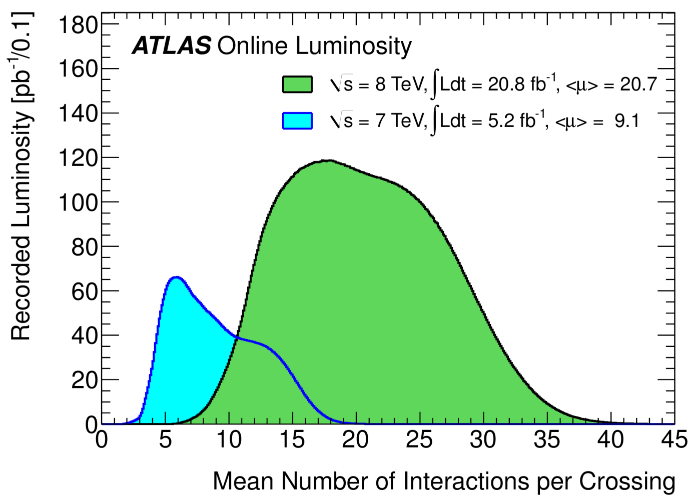

Higher pile-up causes higher CPU usage

An higher number of primary vertexes per bunch crossing increases the computational complexity of reconstruction algorithms.

Level-2 event processing time vs. pile-up

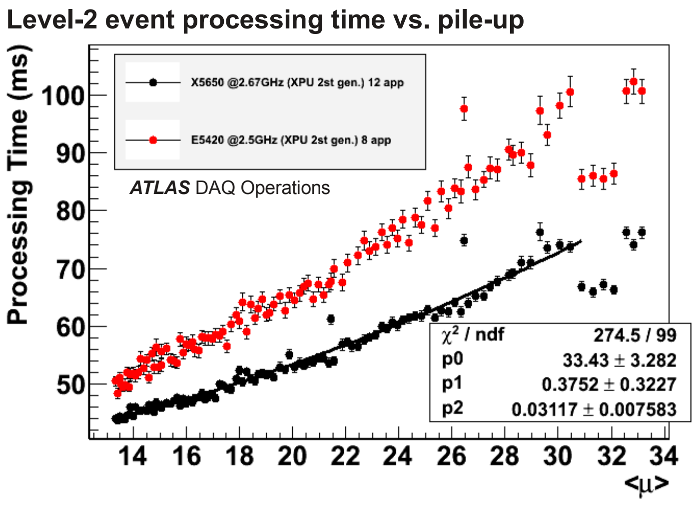

Event Filter event processing time vs. pile-up

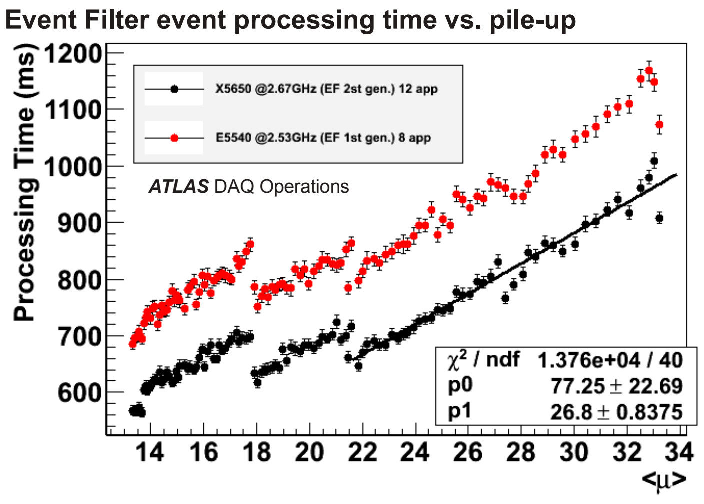

Higher pile-up means bigger events

As the number of primary vertexes increases, so does the number of physical objects and therefore the average event size.

Event size vs. pile-up for the main physics streams

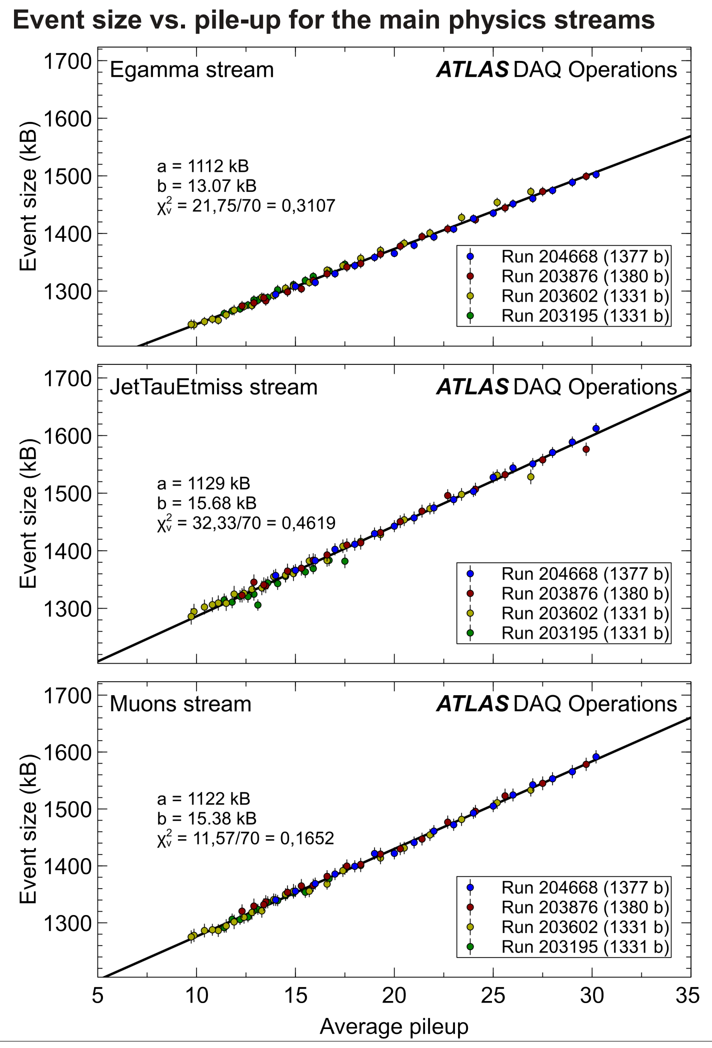

The ATLAS Trigger and Data Acquisition System

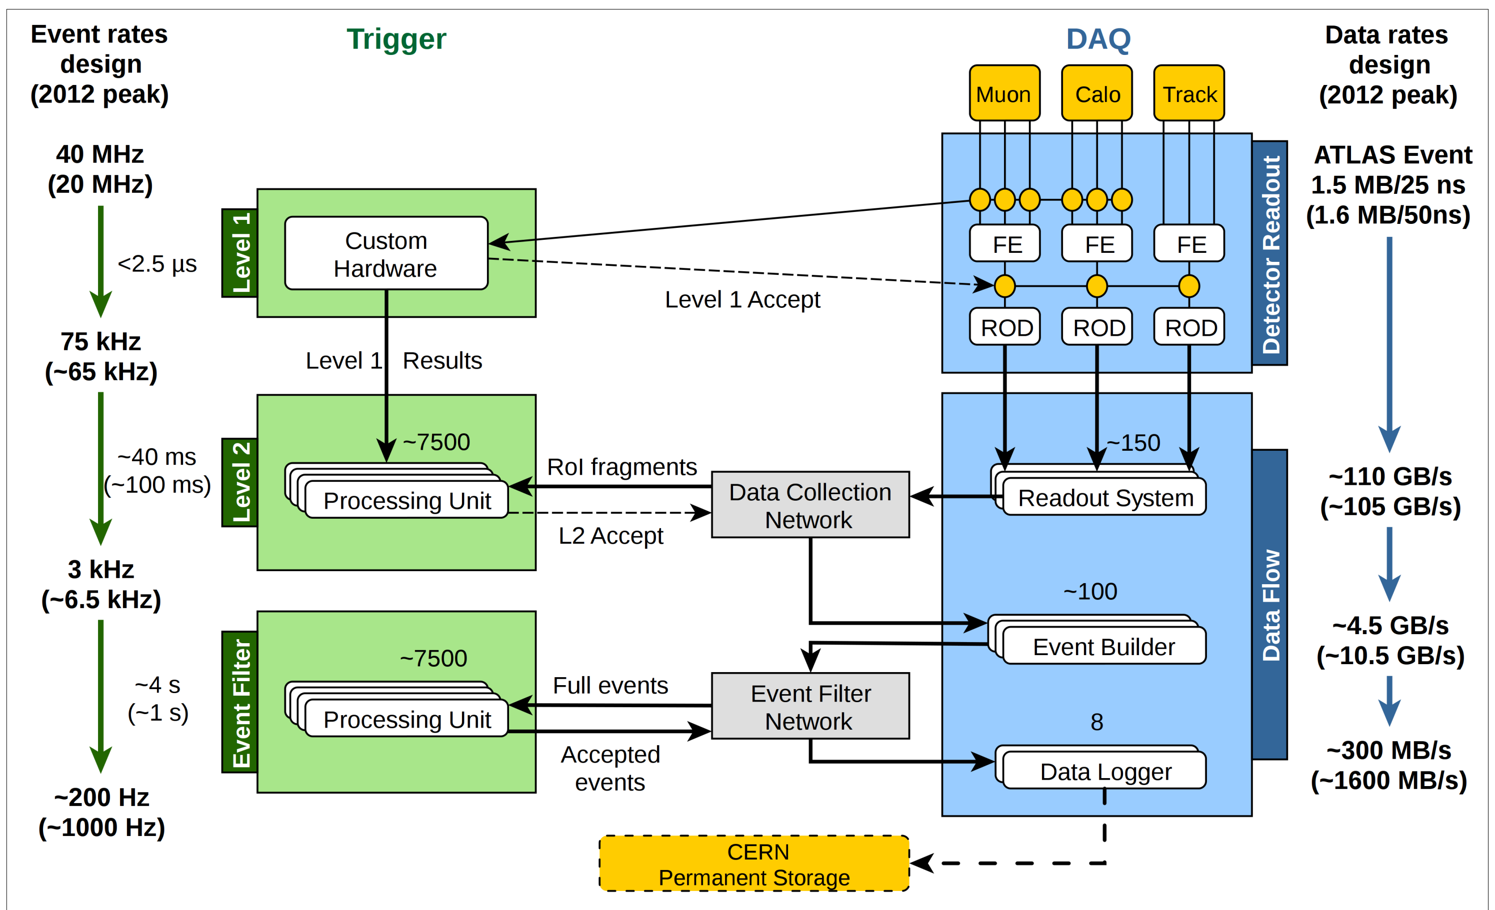

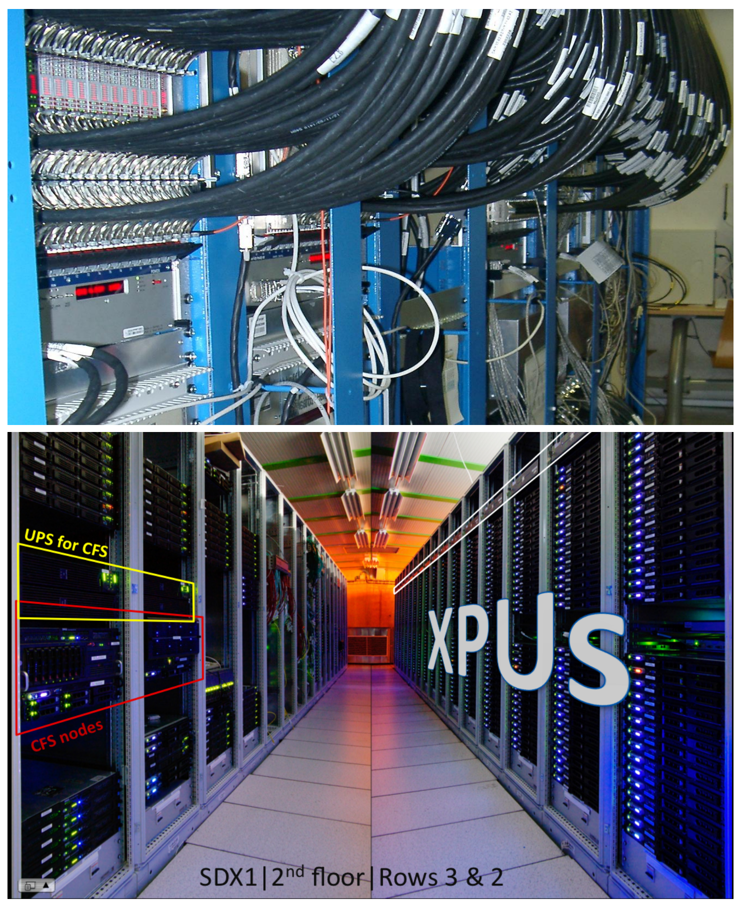

Level-1 operates on coarse-grained data from the muon and calorimeter systems.
Level-2 operates on "regions of interest" corresponding to a small fraction of the full event.
Event Filter operates on the fully built events.

The result: acquired more than twice the data than 2011 with the same efficiency!

The ATLAS TDAQ system was operated well beyond its design parameters, supporting the pursuit of the experiment's physics goals. This was made possible by a design that supports excellent "vertical" and "horizontal" scalability. Thus the system could be made to cope with the new conditions simply by adding additional hardware resources.

In particular, the Data Logging subsystem could be easily expanded to operate at more than five times its design input bandwidth.
The Event Builder, instead, could run at more than twice its designed input bandwidth thanks to the acquired operational experience enabling a less conservative usage of the available resources.

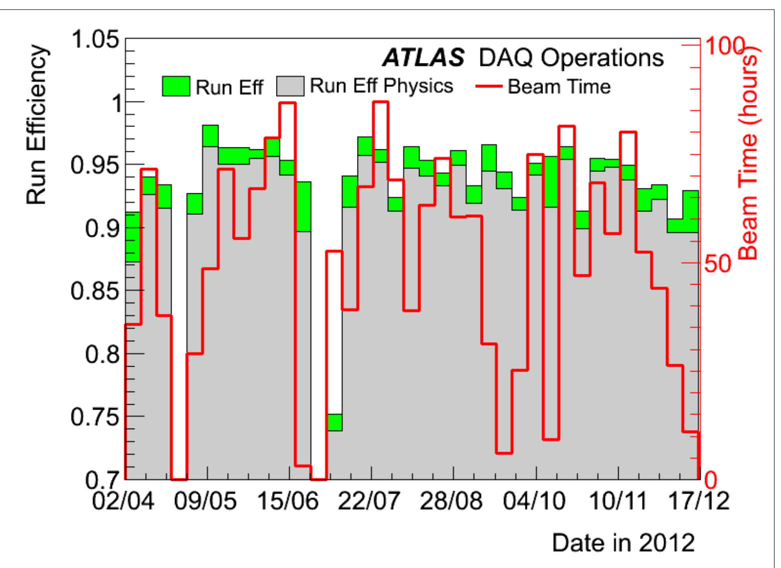

Despite more challenging conditions, ATLAS operated with the same level of efficiency of 2011.

	2011 efficiency	2012 efficiency
Stable beams	1550 hours	1800 hours
Running and not busy	95.0 %	94.4 %
Running and ready for physics	94.0 %	93.7 %

The difference between "not busy" and "ready for physics" accounts for the time it takes to ramp up the high voltage in the inner detector, an operation that can be only performed after stable beams have been declared.

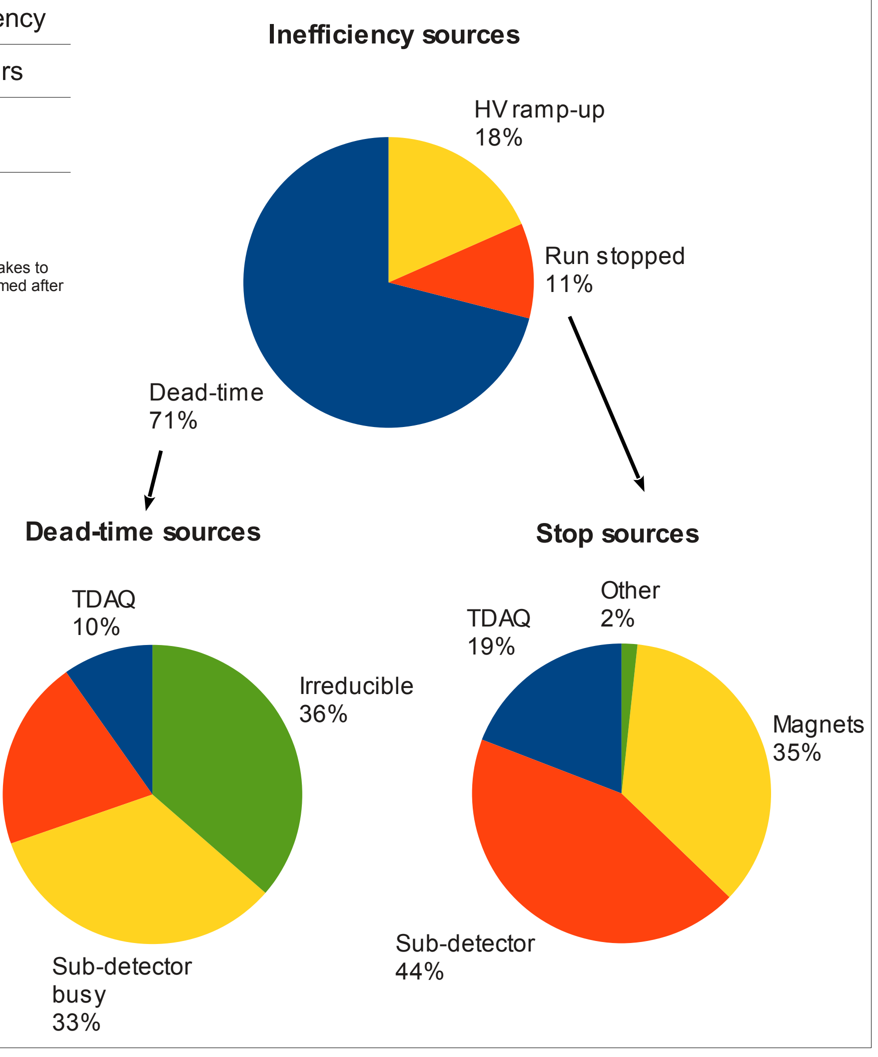

In 2012, the TDAQ community spearheaded an automation campaign for ATLAS:

- New software procedures enabled to deal with most error conditions without stopping the run
- The knowledge base of existing "expert" systems were vastly expanded in order to streamline the recovery from common errors

This resulted in:

- An ATLAS TDAQ system that "just works" and was responsible for small fractions of dead-time and stop-time
- A reduction of the number of instances in which the run needed to be stopped in order for an expert to manually fix problems, thanks to the adoption of the automated procedures by sub-detectors,

Towards 2015: an evolved system with additional operational flexibility and scalability

Currently Level-2 and Event Filter are strictly separate systems:

- Two computing farms with O(1000) heterogeneous nodes each
- Two different Ethernet networks
- O(100) Event Builder nodes interfacing the two

The evolution project:

- Merge Level-2, Event Builder and Event Filter into one single farm
- Each node performs all the steps required by the triggering process
- Enables greater flexibility in responding to varying operating conditions
  - no more static division of CPU power and network bandwidth between Level-2 and Event Filter
  - distributes the Event Builder over the whole farm

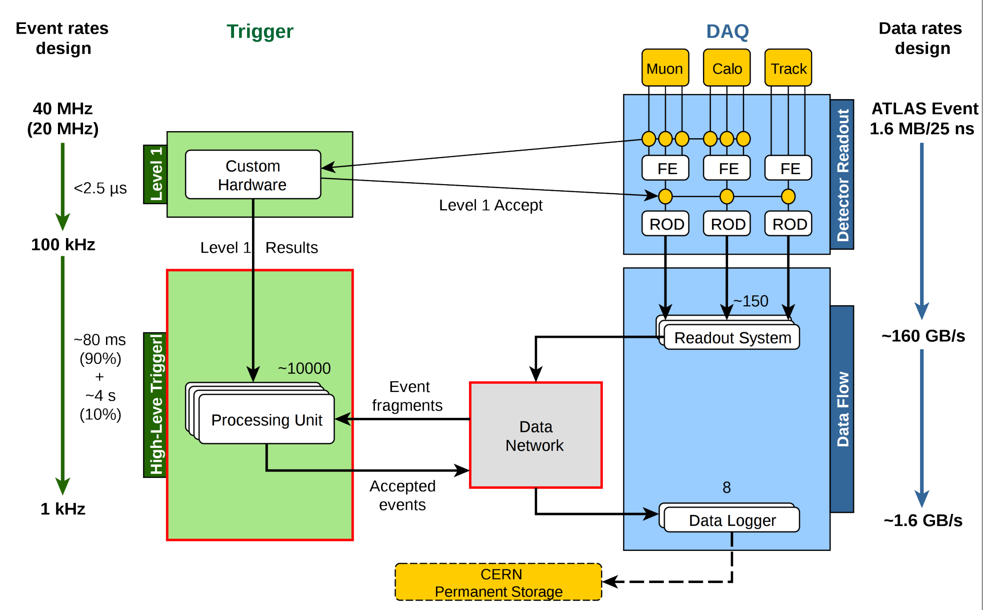

A typical run

Run 214494, started on Tue Nov 13, 11:54 UTC.
Peak luminosity: 7.08x10^{33} cm^{-2} s^{-1}
Peak average pile-up: 33.58
LHC delivered luminosity: 152.3 pb^{-1}
ATLAS recorded: 146.8 pb^{-1} [96.3%]

Evolution of the Level-1 total trigger rate

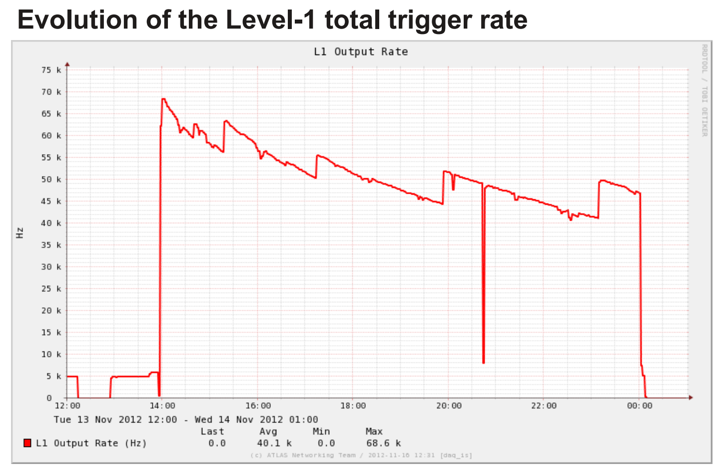

Evolution of the Event Builder input bandwidth (i.e. after the Level-2 selection)

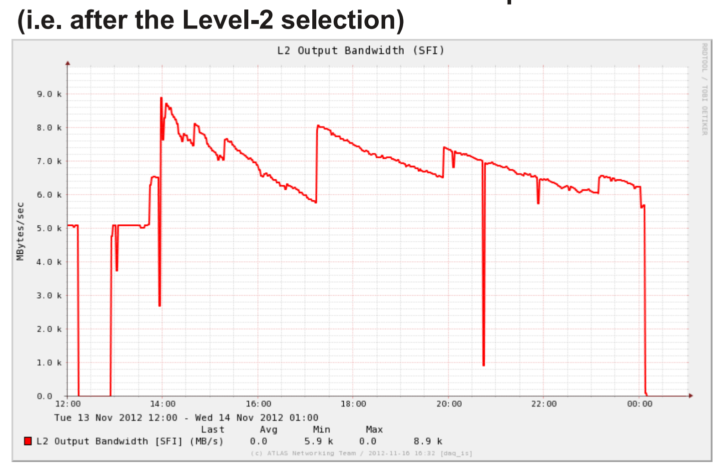

Evolution of the Data Logger input bandwidth (i.e. after the Event Filter selection)

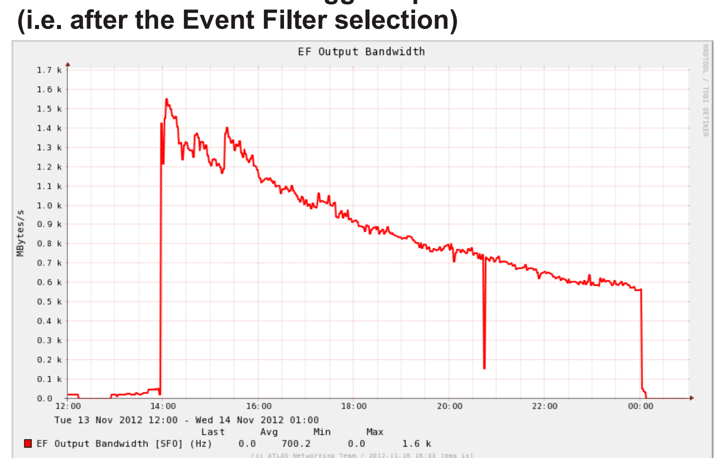

The trigger rates and data-flow bandwidths clearly reflect the decreasing instantaneous luminosity as the LHC fill progresses. After the trigger rates drop to a configured threshold, a "levelling" mechanism enables additional triggers as the fill progresses, to further exploit the resources freed by the drop in luminosity.

Evolution of the event size (measured at the Event Builder)

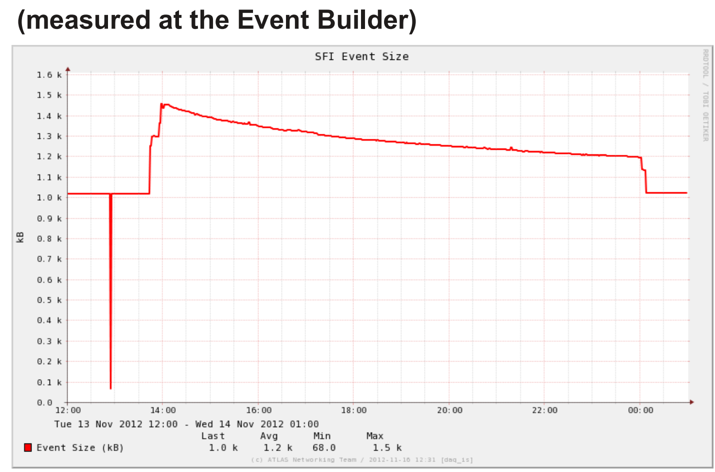

The effect of pile-up on the average event size is clearly visible in the evolution over a LHC fill. Due to this, even if trigger rates remained constant throughout a fill, higher bandwidth usage is to be expected at its beginning.

Tommaso Colombo (CERN and Universität Heidelberg) for the ATLAS Collaboration

ATLAS EXPERIMENT