The ATLAS Level-1 Trigger System

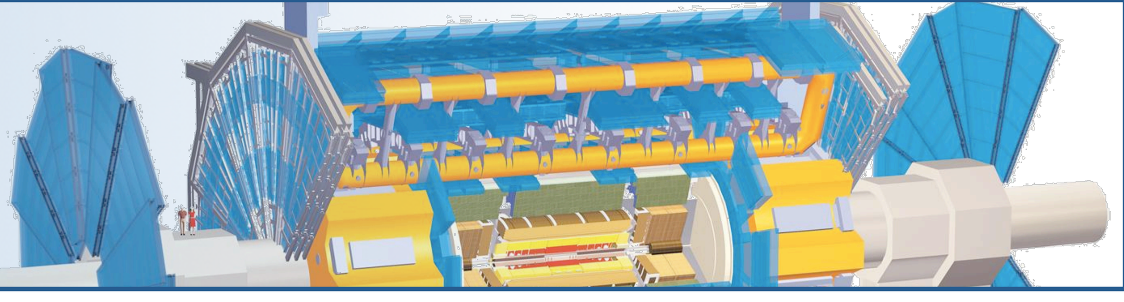

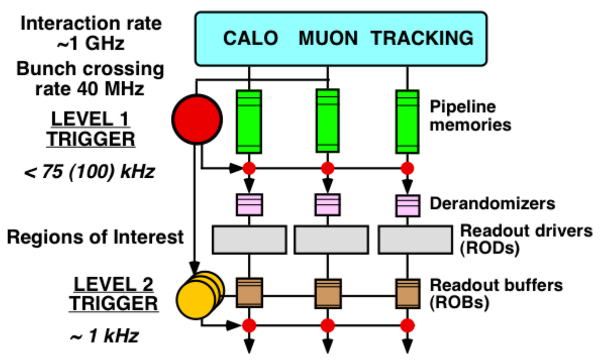

Figure 1: Schematic diagram of the Level-1 Trigger's position in the ATLAS Trigger system. The High Level Trigger (Level 3) is not shown. The final data recording rate is approximately 300Hz

The Level 1 Trigger of the ATLAS Detector is a hardware-based (extensively ASICs), synchronous (with the 40MHz LHC Clock), pipelined system consisting of a Central Trigger Processor (CTP) fed by signals coming primarily from dedicated trigger hardware in the Calorimeter (L1Calo) and the Muon (L1Muon) detector systems. The output of this trigger system is a single bit "Level 1 Accept" (L1A) signaling to the detector front-end readout systems, via the Trigger Timing and Control system (TTC), whether to readout the pipeline to buffers or discard the data for an LHC bunch-crossing.

Pipeline memory sizes limit the level 1 trigger decision latency to 2.5μs

The L1A is formed from the logical OR of 256 trigger items, with the trigger rate of each item constantly monitored.

The ATLAS detector trigger and data acquisition (TDAQ) system is designed to operate with a **maximum L1A rate of 75khz** (upgradeable to 100kHz)

As the LHC instantaneous luminosity continues to increase, the ATLAS LVL1 Trigger must continue to adapt to keep the output L1Accept rate within the prescribed limits, while maintaining signal trigger efficiency

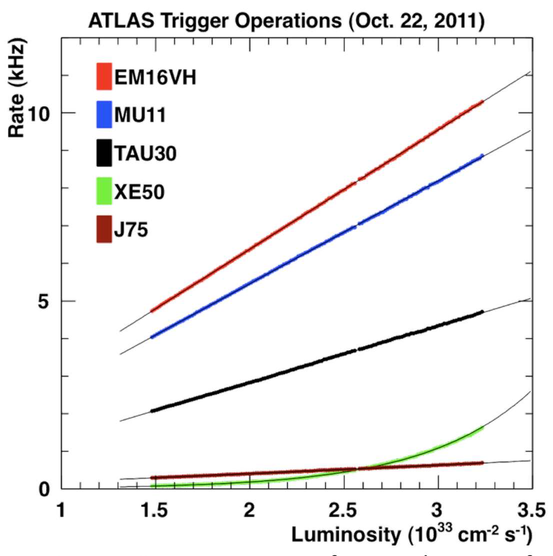

Figure 2: Trigger rates for a selection of trigger items (EM16VH is an electron/photon trigger, MU16 is a muon trigger, TAU30 is hadronic tau, XE is missing E_T, J75 is a jet trigger). Not all triggers scale linearly with luminosity

The LHC accelerator is continually aiming to deliver higher instantaneous luminosities to the detectors. In 2011, the accelerator was able to deliver upwards of 20 inelastic collections per colliding bunch (assuming a 71.5mb inelastic cross-section), with typically 1300 colliding bunches per LHC orbit (period 88μs).
In 2012, we can expect this to continue to increase.

ATLAS LVL1 Trigger Key Stats

Input signal rate: **40MHz**

Target Latency: **2μs (with 500ns contingency)**

- 0.75μs for longest time-of-flight (TGC part of L1Muon)
- 0.4μs worst case transmission time, 80m cables to trigger hardware (in USA15 Cavern)
- 0.4μs worst case transmission of L1A back to Front End Electronics
  - **Only 0.5μs actual processing time available**

L1A rate limit: **75kHz** (1 in 534 bunch crossings)

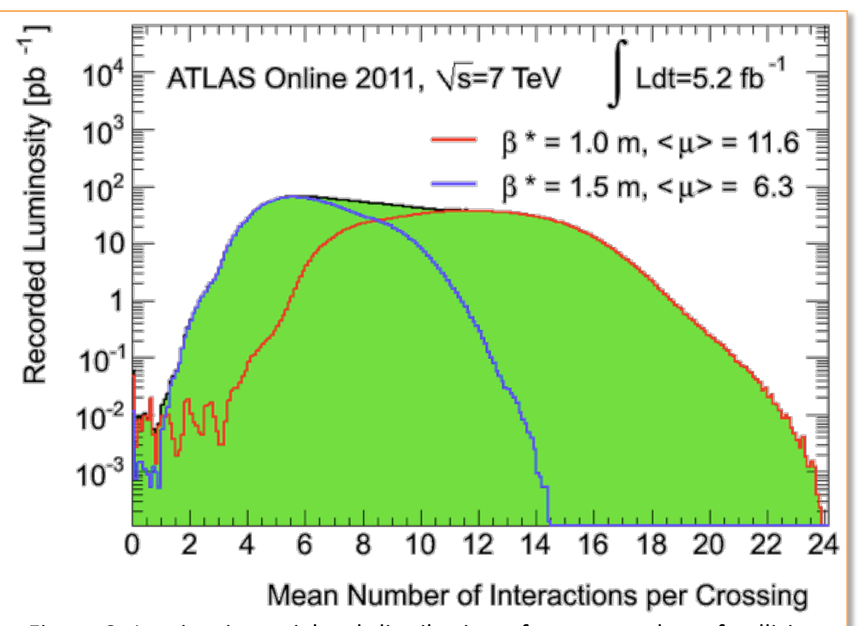

Figure 3: Luminosity-weighted distribution of mean number of collisions per bunch crossing, for 2011 proton-proton data. In Sep. 2011, the beam's β* was reduced from 1.5m to 1.0m, resulting in an increased collision rate per bunch

L1Calo – Calorimeter Trigger

Inputs (from detector)

- Analogue LAr and Tile Calorimeter signals, summed in granularities of $\Delta\eta\times\Delta\phi = 0.1\times0.1$ (central region) up to 0.4×0.4 (forward regions) = **7168 Trigger Towers**, split over EM and Hadronic layers

Processing (in dedicated electronics located in USA15 Cavern)

- Digitize analogue signals (40MHz sampling), **measure deposited transverse energy** and perform **bunch crossing identification** in each tower (using a FIR Filter and calibrated Look Up Table)
- Sliding window algorithm to find local energy maxima across towers and compared to threshold (**EM, TAU** triggers – up to 16) and perform summations to apply isolation vetoes, if required
- Summation in depth to form 0.2×0.2 **jet elements**, used for jet finding (J, FJ) and whole-calorimeter energy summations (**TE and XE, XS**)

Outputs (sent to Central Trigger Processor)

- **EM, TAU, J** (jet), **FJ** (forward jet - |η|>3.2) multiplicities for each threshold in the trigger menu
- **XE** (missing-E_T), **XS** (missing-E_T significance), **TE** (total energy) bits indicating pass or fail for each of the pre-specified thresholds in the trigger menu

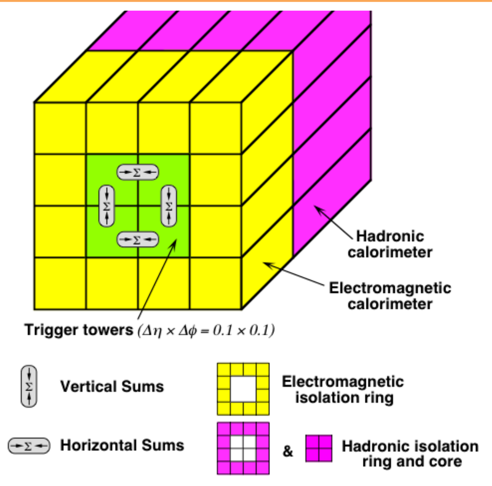

Figure 4: Trigger towers used in the EM and TAU trigger algorithms. The vertical or horizontal sums must be a local energy maximum

The trigger cross-sections, shown right, demonstrate the challenges faced by the Level 1 Trigger as luminosity is increased. Flat trigger cross-sections indicate a linear dependence on luminosity, and are directly caused by the underlying physics events; the rise in cross-section in the 8TeV region for all such triggers is consistent with the increased centre-of-mass energy.

Multi-object triggers, such as 4J10 (require 4 Jet ROIs satisfying the 10GeV threshold) have non linear rates, as expected for luminosities with multiple inelastic collisions per bunch (see "Multi object trigger rates" box, far right).

The missing E_T (XE) trigger rate is found to be particularly non-linear with luminosity. Forward Jet triggers also show a non-linear dependence. This behaviour has been identified as a consequence of the larger sizes of the trigger towers in the FCAL regions (|η|>3.2) and calorimeter end-caps (2.7<|η|<3.2). Studies showed increasing the tower noise cuts (see "Forward Region Noise Cuts") could calm these trigger rates without significantly affecting trigger performance. These new cuts were implemented for 2012 running.

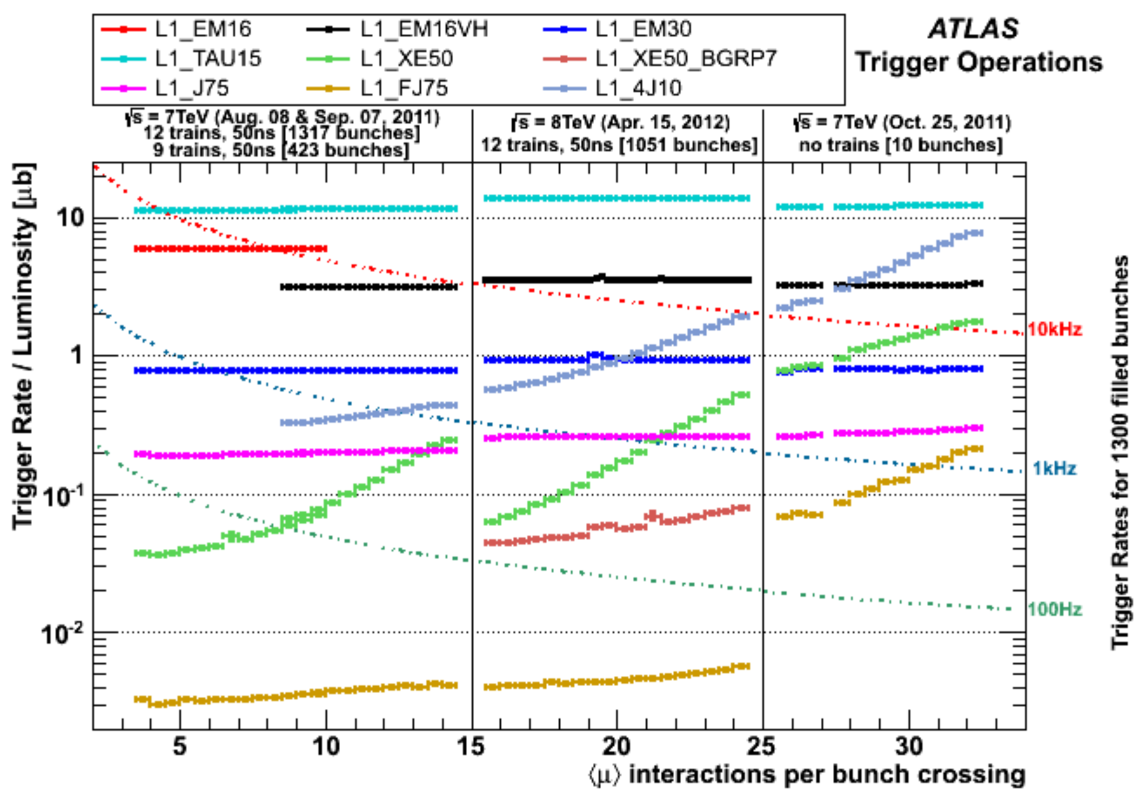

Figure 5: Trigger cross-section (Rate/Luminosity) as a function of interactions per colliding bunch. For 7TeV, <μ> is determined with an inelastic cross-section of 71.5mb. For 8TeV, 73mb is used. See text for further explanation of diagram

VH Thresholds

During 2011, a modification was made to a few of the EM thresholds, so that the threshold energy required for a trigger-tower pair was dependent upon η. This was motivated by measurements that L1_EM trigger efficiency (with respect to offline reconstructed electrons) was not uniform across the detector. In some regions it was possible to raise the threshold without significantly altering the efficiency turn-on. Simultaneously, a veto on the hadronic core value (allow no more than 1GeV) was also introduced.

For these two actions, a significant reduction in the trigger rates (~50%) for this item was the reward (see figure 5)

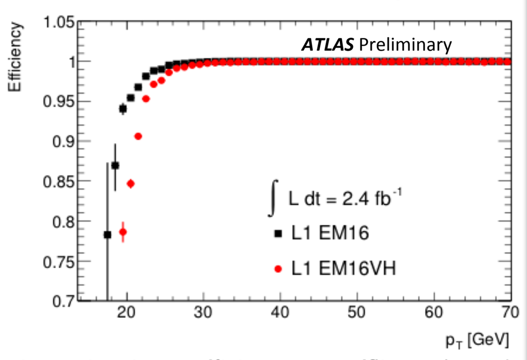

Figure 6: Trigger Efficiency wrt offline selected electrons from a Z *Tag and Probe* study

Forward Region Noise Cuts

The high occupancy of these towers (due to their relatively large size), coupled with unbalanced overlaying of the bipolar analogue pulse shapes they produce (see figure 15), leads to significant fluctuations in the analogue signal measurement (the ADC count, after digital sampling).

These "pileup" effects are quantified by measuring the standard deviation of ADC distributions obtained at different values of bunch-specific μ. A linear dependence is observed, and noise cuts suppressing the pileup fluctuations are derived. Preliminary studies suggest minimal impact on efficiency, for a significant reduction in trigger rate!

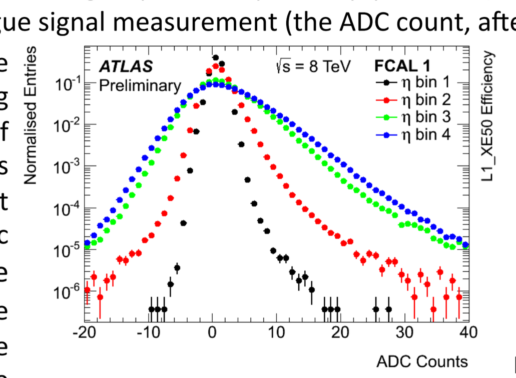

Figure 7: ADC distribution in different regions of FCAL trigger towers, in zero-bias (random events) data. The standard deviation of this distribution is measured for different values of μ

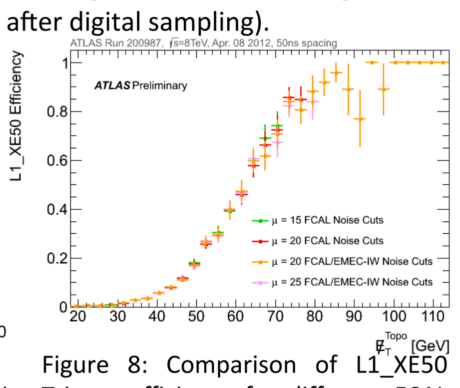

Figure 8: Comparison of L1_XE50 Trigger efficiency for different FCAL noise cuts, using a sample of candidate W events, wrt MET calculated using topological clusters.

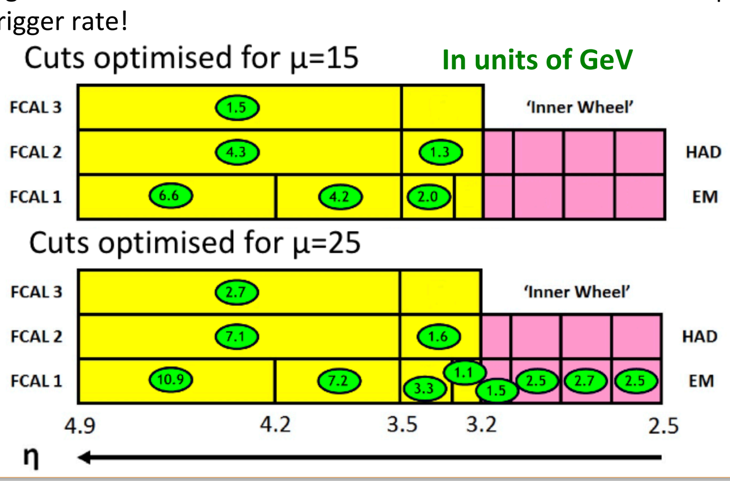

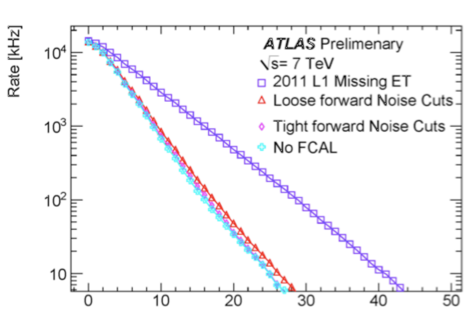

Figure 9: Reduction in XE rate with and without new noise cuts, and with the FCAL trigger towers excluded from trigger algorithms entirely

L1Muon – Muon Trigger

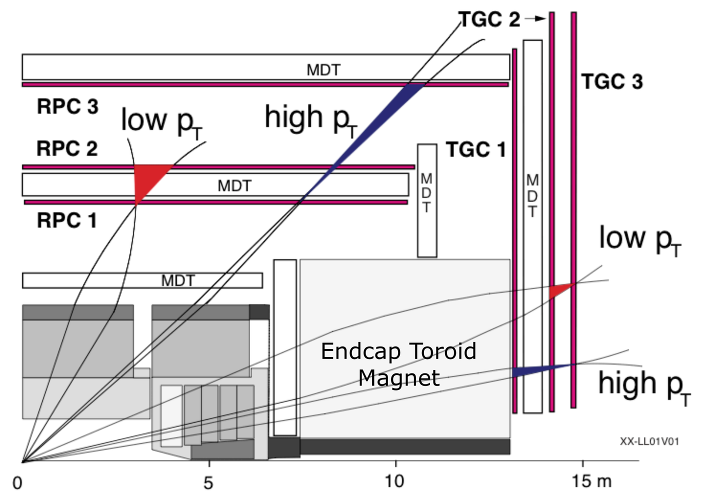

Figure 10: Longitudinal slice through ATLAS detector, showing location of RPC and TGC muon trigger chamber stations. Low-pT "Two RPC station" thresholds are for muons with approx. 6-10GeV transverse momentum. Muons are bent by the toroid magnet

Inputs (from detector)

- Over 800k hit strips from Resistive-Plate Chambers (**RPC**) in the barrel region (|η|<1.05) and Thin-Gap Chambers (**TGC**) in the endcap regions (1.05<|η|<2.4), spread over three planes ("stations") of doublet chambers

Processing (on-detector ASICs and in electronics located in USA15 Cavern)

- **Coincident matrices ("roads")** search for track candidates consistent with six different programmable thresholds, in granularities of 0.1×0.1. This logic is performed in 104 independent sectors
- Multiplicity counting for the different thresholds, accounting for chamber overlaps between sectors. This logic forms the **MUCTPI**

Outputs (sent to Central Trigger Processor)

- Multiplicities for **six MU (muon) thresholds:** three low-p_T "Two station" thresholds (require coincident hits in two planes in RPC, three in TGC), three high-p_T "Three station" thresholds (require three hits in all regions).

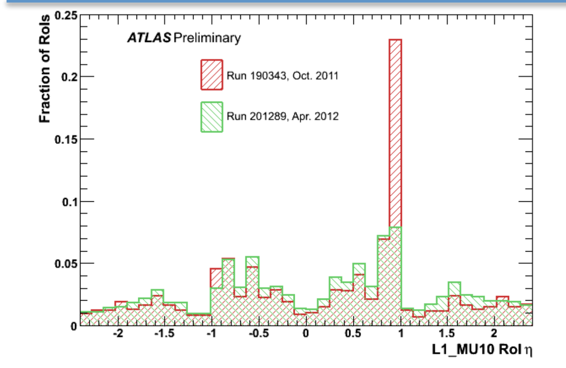

Figure 11: Distribution of L1_MU10 RoI, a trigger requiring 2 coincident station hits in the RPC region

Additional Shielding for 2012

A narrow region between the barrel and endcap regions of the detector was known to have inadequate shielding of the beampipe for the muon chambers. This led to significantly more triggers (primarily affecting the low-momentum 2-station thresholds) in this region (η≈1.0) due to gamma rays radiating from the beam. With additional shielding installed between 2011 and 2012, this effect was reduced along with significant reductions in trigger rates

Simulated single muon events are used to calculate coincidence windows for each point (RoI) of the pivot-plane: the window effectively indicates where hits in the other stations/planes are permitted (specified in radial and azimuthal directions as deltas from the line joining the pivot point to the origin) in order to pass the trigger threshold. In 2010 and early 2011, the windows were conservatively determined with 90% of hits associated to truth muons over the threshold being required in the window. Also a lack of cross-talk simulation in the TGC region meant windows here were artificially widened to recover efficiency in the data (at the cost of higher trigger rates).

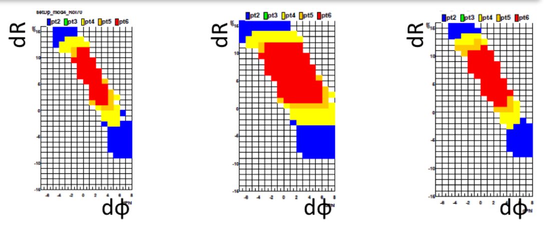

Figure 12: Example coincidence window from TGC. Left is original calculation (no cross-talk in simulation). Middle is the artificially widened windows to recover efficiency. Right is the newly optimized windows (with cross-talk simulated)

In mid-2011, the RPC windows for MU4, 6, 10, and 11 were tightened to require only 90% of clusters (instead of hits) from truth muons lie within the window. The smaller windows reduced trigger rates (see first section of the figure 10) and were validated to not impact on trigger efficiency. New windows for all thresholds were again optimized for 2012 running (with added cross-talk simulation for TGC regions, so the windows could be made smaller)

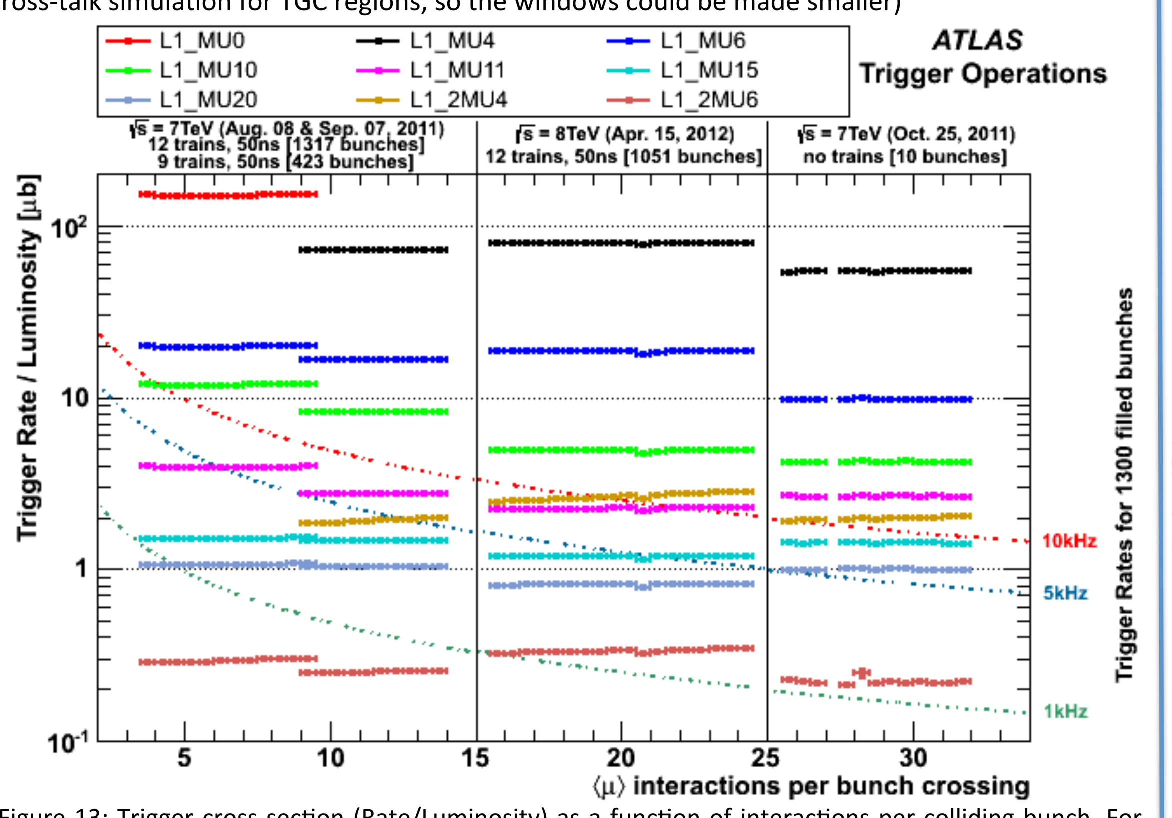

Figure 13: Trigger cross-section (Rate/Luminosity) as a function of interactions per colliding bunch. For 7TeV, <μ> is determined with an inelastic cross-section of 71.5mb. For 8TeV, 73mb is used. Muon trigger rate has a significant contribution from late hits, hence low rates for the well-spaced bunches of the third section of the plot (in first two sections, late hits contribute to neighbouring bunch crossings). See text for further explanation of figure

Multi-Object Trigger Rates

The non-linearity of trigger rates with respect to luminosity (figure 5 & 13) can be understood in terms of the number of interactions per bunch crossing. If the probability for a single collision to cause a single object trigger to fire (e.g. L1_MU6) is given by f, then in a bunch crossing with n collisions, the probability for the single object trigger to fire is 1 – prob(no trigger).
Assuming p is small, we have the following:

$$L \propto n \quad , \quad Rate_{Trig} \propto P_{Trigger}$$

$$P_{Trigger} = 1-(1-f)^n \approx fn$$

$$\Rightarrow Rate_{Trig} \propto L$$

For, e.g,. a two-object trigger, the probability is given by 1- prob(1 or no triggers):

$$P_{Trigger} = 1-(1-f)^n - nf(1-f)^{n-1}$$

$$\approx fn - nf(1-(n-1)f) = (n^2-n)f^2$$

Therefore multi-object triggers are non-linear

TGC Fake Muon Triggers

Figure 14 shows that a large proportion of the fake muon triggers (unmatched to an offline reconstruction muon) come from the endcap region of L1Muon. It is believed that this is due to low-momentum protons from secondary interactions near the endcap toroid. The protons are able to traverse a path through the TGC that corresponds to a muon originating from the centre of ATLAS.
Simulation has shown that by incorporating information from TGC located just inside of the endcap toroid (to-date, such information has not been used by L1Muon), the fakes can be distinguished from real muons coming from the origin, and so significant trigger rate reductions should be possible. This will be implemented as part of the 2014 "phase0" upgrade

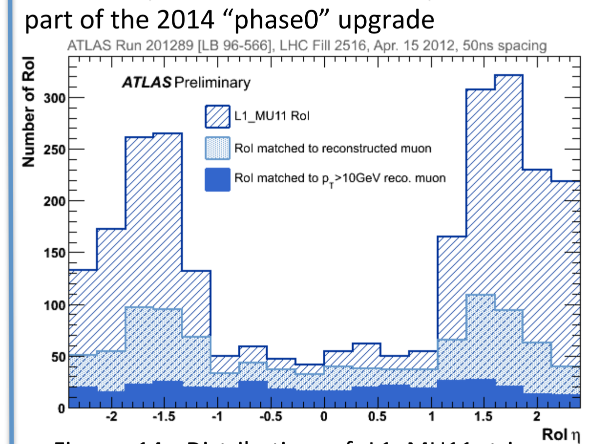

Figure 14: Distribution of L1_MU11 triggers, shown with the distribution of those matched to an offline reconstruction muon. Most "fake" triggers occur in the TGC endcap trigger system

L1 CTP – Central Trigger Processor

Inputs

- Trigger inputs from L1Calo, L1Muon and specialized detectors (MBTS, Lucid, BPTX, ALFA)
- LHC Timing Signals (Bunch crossings, beam orbits)
- Busy signals from TTC partitions (about 40)

Outputs

- Broadcast L1A and LHC timing information to front-end detector Read-Out Drivers (ROD) via the TTC
- Send Region of Interest (RoI) information to LVL2 Trigger

Processing

- Align trigger inputs, in-time, with respect to one another – **CTPIN**
- Synchronize LHC timing signals with beam pickup **- CTPMI**
- Form trigger items, specified in the trigger menu, from the trigger inputs
- Prescaling of trigger items and bunch group masking
- Generate two forms of preventative deadtime: simple (fixed number of bunches, 5 in 2011, 4 in 2012) and complex (leaky-bucket algorithm)
- Logical OR of trigger items after prescaling, masking, deadtime and busy veto – this forms the Level 1 Accept (L1A)
- Online monitoring of trigger, deadtime and busy information (accumulative and bunch-by-bunch counters) – **CTPMON**

BGRP7 Mask

In 2012, a new bunch group mask (BGRP7) was provided that masks the first three bunches of any bunch train. This was motivated by the strong bunch-position dependence of the missing energy (XE) trigger, caused by the bipolar analogue pulse shapes of the Calorimeter trigger detector, which span several bunches. Figure 5 shows the L1_XE50_BGRP7 trigger that uses this bunch group mask. It can be seen the rate is lowered and the luminosity dependence is significantly calmed.

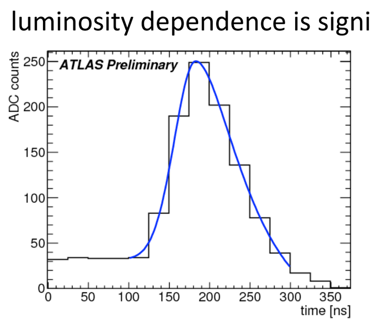

Figure 15: Bipolar pulse shape for a typical L1Calo trigger tower. The pulse spans several bunch crossings (the full pulse undershoot is not shown)

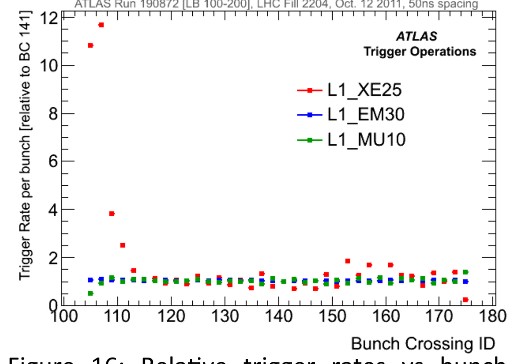

Figure 16: Relative trigger rates vs bunch crossing for a single bunch train of a typical 2011 LHC Fill. The XE trigger exhibits significantly higher rates at the start of the train. The slightly decreased trigger rate in the first bunch for L1_MU11 is due to there being a non-negligible contribution to muon trigger rate from late hits, particularly in the region that was unshielded in 2011 (see figure 11)

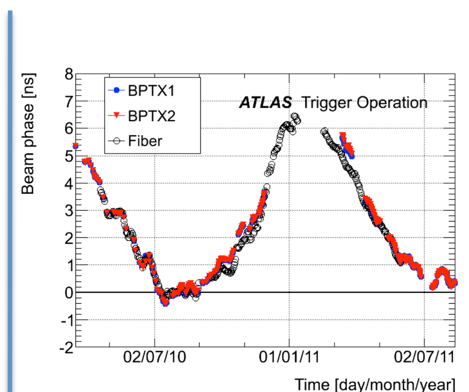

Figure 17: Phase shift of LHC beam pickup to LHC timing signal. The phase shift is constantly monitored by the CTP and corrected for before broadcasting the clock to the the detectors via TTC

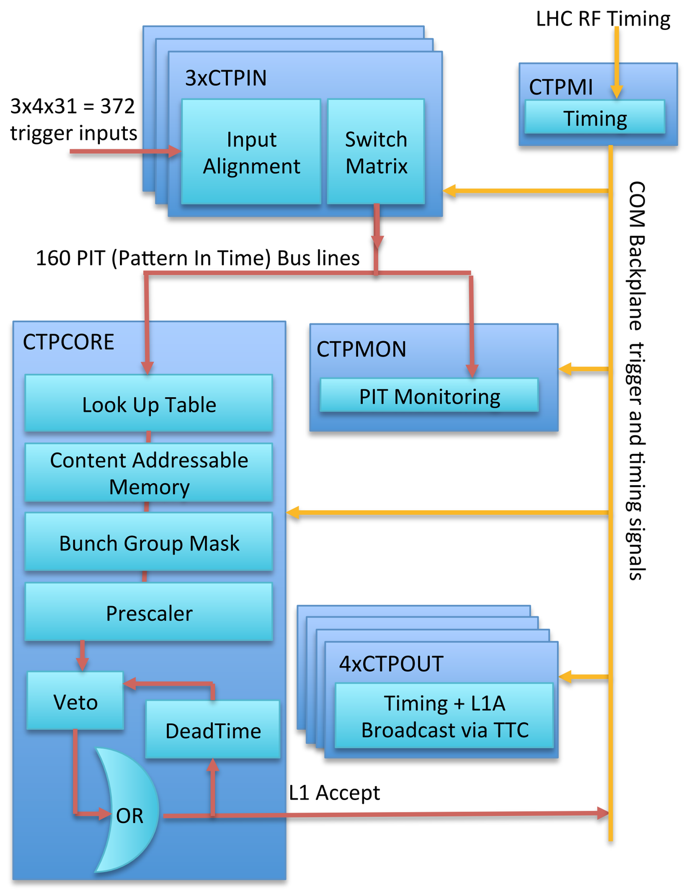

Figure 18: Outline of the CTP hardware. The Content Addressable Memory supports up to 256 trigger items, formed from up to 160 trigger inputs (current PIT Bus limit). The permitted CTP latency is about 100ns (4 bunch crossings). CTPCORE sends RoI summary information to LVL2 and DAQ systems (not shown in figure)

Preliminary Upgrade Plans

Since 2011, many features of the CTP have been operating at maximum capacity. Upgrades are planned for the 2013 long shutdown of LHC, which will almost double the capacity of the CTP hardware (CTPCORE and CTPOUT):

- The 160 PIT Bus line signals will be multiplexed to deliver double data rate(80MHz) signaling on the existing backplane – effective 320 PIT Bus lines
- Additional optical inputs (using SNAP12 ribbon-fiber receivers) for new trigger inputs (e.g. a new L1Calo Topological processor). Will migrate CTPIN connections from electrical DDR PIT Bus to optical inputs
- Three additional L1A partitions – secondary partitions can be used for commissioning new detector systems, or calibration runs. They will have their own deadtime and selection of trigger item OR for L1A ("Veto", "Deadtime" and "OR" of figure 18). CTPOUT would need a redesign too.

	Used Now	Available Now	Planned Upgrade
CTPIN Input Signals	212*	372	
CTPIN Integrating monitoring counters	138	768	
PIT Bus Lines	160	160	320
CTPCORE Trigger items	241	256	512
CTPCORE Bunch group masks	8	8	16
CTPCORE Max number of AND terms	6	256	
CTPCORE Max number of OR terms	6	12	
CTPCORE per-bunch trigger item counters	12	12	256
CTPOUT Cables to TTC partitions	20	20	25
CTPMON per-bunch monitoring counters	88	160	

* Only 160 items can be forwarded to CTPCORE, so some items are for CTPIN monitoring only

Will Buttinger, on behalf of the ATLAS Collaboration