

Towards a Level-1 tracking trigger for the ATLAS experiment

Alessandro Cerri*, on behalf of the ATLAS Collaboration

University of Sussex

E-mail: a.cerri@sussex.ac.uk

Among the upgrades for the High-Luminosity LHC era, the ATLAS collaboration is studying and developing the availability of inner detector tracking information at the first level of its three-tiered event selection chain. This will provide additional flexibility and rejection power: essential ingredients in order to cope with the demanding conditions of the upgraded LHC, as well as with unforeseen bandwidth constraints. The current state of the feasibility and performances studies is discussed.

*Technology and Instrumentation in Particle Physics 2014,
2-6 June, 2014
Amsterdam, the Netherlands*

*Speaker.

1. Introduction

The future plans for the LHC accelerator allow, through a schedule of phased upgrades, an increase in the average instantaneous luminosity by a factor 5 with respect to the original design luminosity. The ATLAS experiment [1] at the LHC [2] will be able to maximise the physics potential from this higher luminosity only if the detector, trigger and DAQ infrastructure are adapted to handle the sustained increase in particle production rates. This paper describes the changes expected to be required to the ATLAS detectors and trigger system to fulfil the requirement of such high luminosity scenario. The increased number of interactions per bunch crossing will result in higher occupancy in the detectors and increased rates at each level of the trigger system. The trigger selection will improve the selectivity partly from increased granularity for the sub detectors and the consequent higher resolution. One of the largest challenges will be the provision of tracking information at the first trigger level, which should allow a large increase in the rejection power at this stage of the selection and yet still allow the full physics potential of the experiment to be fulfilled. In particular, the electroweak scale still requires the thresholds on the transverse momenta of particles to be kept as low as possible. Tracking can provide essential information toward this aim early in the trigger chain. Studies to understand the feasibility of such a system have begun, and are proceeding in two directions: a fast readout for high granularity silicon detectors, and a fast pattern recognition algorithm to be applied just after the Front-End readout for specific sub detectors. Both existing, and novel technologies can offer solutions. The aim of the ongoing studies we review on this paper is to determine the parameter space to which this system must be adapted.

This paper will first give an overview of the upgrade plans, including the physics motivation and goals, and then focus on the tracking system and its readout. Performance constraints and implications for the Level-1 tracking hardware will be discussed, and potential strategies for track finding and reconstruction will then be outlined.

2. The LHC Upgrade Path

The Large Hadron Collider accelerator and experiments upgrade program is staged in three separate phases. The currently ongoing *Phase 0* is aimed at an instantaneous luminosity of 1.6×10^{34} Hz/cm² (*Run 2*). Data taking is planned to begin in early 2015 and last until mid-2018, when the *Phase I* upgrades will require a second shutdown of the complex until early 2020. The following *Run 3* data taking period will occur at an instantaneous luminosity of $2 - 3 \times 10^{34}$ Hz/cm². At this point the nominal centre-of-mass energy of the LHC will have been attained, and the design instantaneous luminosity surpassed. The LHC upgrade path doesn't stop here, but foresees yet another upgrade (*Phase II*) aimed at an instantaneous luminosity roughly $5 \times$ larger than the initial machine design. This High Luminosity LHC (HL-LHC) scenario is quite challenging both from the experiments and accelerator point of view, and can in some respects be seen as a new project, with significant R&D programs currently ongoing. Among these is the upgrade of the ATLAS trigger system and the use of tracking information starting in the first level of the trigger.

3. Physics Reach for HL-LHC

The HL-LHC phase has the goal of expanding the LHC physics capabilities through the inte-

gration of 3 ab^{-1} of pp collisions, corresponding to roughly $300 \text{ fb}^{-1}/\text{year}$. This will give ATLAS access to high-statistics searches and precision measurements. In the SUSY sector, for instance, a larger parameter space will be excludible: direct neutralino and chargino searches for gauge-mediated decays with three leptons in the final state will give exclusion capabilities up to approximately $m_\chi \sim 1 \text{ TeV}$ [3]. Higgs physics will receive a similar boost, with the "rare" $H \rightarrow \mu\mu$ channel measurable with a 7σ significance and a signal-to-background ratio of approximately 10^{-3} [4].

4. The ATLAS Trigger System Upgrade Path

The LHC upgrade phases described in section 2 are to be seen as incremental. From the ATLAS TDAQ point of view, this means that the Phase I and II systems will build incrementally - through upgrades and modifications - on-top of the hardware that was used for data collection in 2010 and 2011.

The ATLAS Run I TDAQ architecture has been described thoroughly elsewhere [5]. It consists of a three-tier system with a first level (Level-1) realised in custom hardware, and the following two (Level-2 and Event Filter, collectively identified as HLT or High Level Trigger) implemented in a cluster of commercial CPUs.

In order to optimise the use of the available bandwidth, the flow of information to the High Level Trigger is initiated through detector Regions of Interest (ROI). ROI's correspond to areas identified by Level-1 primitives as "active" in the given event. At later stages the High Level Trigger can then process larger portions of the detector, up to the full system.

In Phase 0, the TDAQ system architecture changes, with the merging of the Level-2 and Event Filter trigger levels into one homogeneous HLT entity [5]. The upgraded HLT will receive global tracking information from Fast TracKer (FTK): a dedicated pre-processor implemented with custom hardware [6]. The incremental deployment of FTK will be completed by Run 3.

The Level-1 trigger system is being refurbished for phase 0, providing the ability to subdivide the ATLAS detectors into separate read-out partitions and implementing the possibility of performing certain topological selections based on incoming detector primitives [7]. The calorimetric information available at Level-1 is also being upgraded to achieve higher pile-up immunity [8].

The main improvements of the ATLAS detector during the Phase I upgrades will include a further refinement of the Level-1 calorimetric primitives [9], with improved granularity, and the installation of a new layer of muon detectors [10] in the forward ($1.3 < |\eta| < 2.4$) region. These upgrades will reduce fake rates and improve the overall detector resolution.

The HL-LHC ATLAS detector will face pile-up as high as $\langle \mu \rangle \sim 140$ interactions per beam crossing with bunch crossing intervals as short as 25 ns and an equivalent 1 MeV neutron fluency of $\sim 1.5 \times 10^{16}$ particles/cm². Extensive detector upgrades are foreseen for this phase [11], including a redesign of the software and computing infrastructure, a full replacement of the ATLAS inner tracker with a combination of pixel and strip detectors (see section 5 later), the upgrade of the experiment's liquid argon calorimeter electronics, as well as upgrades of the muon drift chambers, the forward detector and the experiment's shielding from accelerator radiation. The experiment TDAQ system upgrade highlights will be the introduction of tracking information at Level-1 (L1Track) and the upgrade of muon detector electronics to cope with the higher collision rate.

5. The ATLAS Inner Tracker Phase II Upgrade

Several requirements and constraints led ATLAS to plan a full replacement of its inner tracker [11]. Among the most important ones is the need to face a fluency $5\times$ larger than in Run I and the desire to achieve a better acceptance and performance for physics (leading to the extension of the tracker η coverage, the reduction of its mass and the ability to measure as many as 14 points per charged track).

As opposed to the existing ATLAS tracking detector which includes a Transition Radiation Tracker, the new tracker (ITK) will be composed of only two subdetectors: an innermost pixel detector, surrounded by a double-sided strip detector. The pixel detector is planned as an evolution of the current ATLAS pixel tracker, made of approximately 640×10^6 channels (vs $\sim 80\times 10^6$ in the Run I version), each corresponding to a pixel element area approximately 1/2 to 1/4 of that in the current ATLAS pixel detector. The detector will have the ability to be read-out at 1 MHz event rate.

The strip detector will be based on double-sided silicon wafers, with a 40 mrad stereo angle and a larger volume coverage than the Run I SemiConductor Tracker [12]. The 74×10^6 channels will be distributed in a large volume, with each front-end chip covering larger detector elements than the inner pixel layers. This would imply a slower read-out time which has substantial consequences on the experiment's TDAQ chain. In particular, the current front-end and read-out technologies would not allow the off-detector transfer of the full event payload in time for L1 decision.

Given these restrictions, there are two design philosophies being considered: the first one adapts to the available off-detector bandwidth, introducing a new trigger level and performing seeded readout of tracking detector regions to limit the amount of additional data to be transferred. The second strategy approaches the problem as the data flows off-detector, and foresees the implementation of preliminary pattern recognition stages within or as close as possible to the detector front-end chain. These options will be discussed in the remainder of this document.

6. A New ATLAS Trigger Level

The first approach mentioned in section 5 introduces a new trigger level (Level-0) with a nominal accept-rate of 500 kHz and a 6 μ s latency. All the ATLAS Level-1 components compatible with this performance (including barrel and end-cap muon trigger detectors, as well as early information from the calorimetric system) are “promoted” to Level-0 selection. These Level-0 primitives will provide RoIs, used to target ITK regions for early read-out on time to be used by the L1Track system. The Level-1 decision can then be taken using all the usual Level-1 primitives implemented in ATLAS, plus the tracks from the L1Track system.

This design relies on the current performance figures of the strip tracker front-end readout ASIC (ABC130 [11]), as implemented in fourth quarter of 2013. This device foresees two buffering levels, corresponding to the new Level-0, and the subsequent Level-1 triggers. The main ITK readout bandwidth limitations come from two serial links: a first serialised daisy-chain link is used to transfer data from the chain of 5-12 ABC130 chips – required to read-out a detector ladder – to the hybrid chip controller (HCC) integrated circuit. The HCC collects the ABC130's data and transmits it off-hybrid through a second serial *stave link* line. The ABC130 daisy chain link is

implemented – for redundancy – as a ring closed through the HCC chip, effectively allowing for the possibility of bandwidth duplication.

This design assumes that the Level-0 ABC130 buffer is loaded with hits at the full bunch crossing rate (40 MHz). Its content is transferred – on a Level-0 accept – to the Level-1 buffer, which provides both full detector read-out up to the maximum Level-1 accept rate of 200 kHz (*L1A* data), and data from part of the detector following *regional read-out requests* (R3). R3 regions are seeded by a specific Level-0 primitive and identified through an RoI map, driving the selection of Level-1 buffer regions to be sent to the Level-1 track finding system. Both the R3 and the L1A data are propagated through the ABC130 daisy-chain link.

6.1 Optimised Performance Studies

The main focus of the current studies has been the viability of the Level-0/Level-1 approach, given the goals of the L1Track system. Some resource optimisation has been implemented in the ABC130 usage in order to guarantee that the R3 data is ready for L1 processing within $6\mu\text{s}$ from the LHC collision (*R3 latency*).

In the remainder of this paper, it will be assumed that on average 10% of the tracker data is requested for regional read-out following a Level-0 trigger. Three main optimisations had to be implemented in order to satisfy the $6\mu\text{s}$ R3 latency constraint:

- the prioritisation of the R3 data over the L1 information in the off-detector transmission;
- the exploitation of the redundant daisy-chain link to duplicate the daisy-chain bandwidth to the HCC;
- the doubling of the HCC *stave link* bandwidth, from 160 to 320 Mbps;

Figure 1 shows the latency for 95% of the R3 data for ITK hybrids (one corresponding to each line) as a function of the Level-1 accept rate. The plot reports simulated results for the innermost barrel modules, where detector occupancy is highest. All data is simulated assuming a 500 kHz accept rate at Level-0, 10% R3 occupancy, 160 Mbps *stave link* bandwidth and the use of both daisy-chain links to the HCC. The two sets of curves exemplify how the implementation of R3 data prioritisation over Level-1 data leads to a significant milder dependency of the R3 latency on the Level-1 accept rate, yielding latencies well below $6\mu\text{s}$ for all the simulated hybrids.

As different ITK areas are simulated, it becomes clear that R3 prioritisation and the exploitation of the dual daisy-chain link are not sufficient: the HCC itself becomes the bottleneck in the off-detector propagation of the hits. This is illustrated by figure 2, where the Level-0 latency of end-cap detector elements (where the length of ABC130 daisy chain can get as large as 12 chips) is shown to go well above $6\mu\text{s}$ at the nominal Level-1 accept rate of 200 kHz. The comparison of the performance with infinite transmission buffer size (dashed) and that with realistic dimensioning of the same buffer (continuous) demonstrates that the stave-link bandwidth is a substantial limiting factor. In order to understand how to circumvent this bottleneck, the same system has been simulated with the doubling of the HCC stave link bandwidth: figure 3 shows how the performance of the system improves when the stave link bandwidth is raised to 320 Mbps, effectively allowing the full tracker to provide R3 data within the allotted $6\mu\text{s}$ latency.

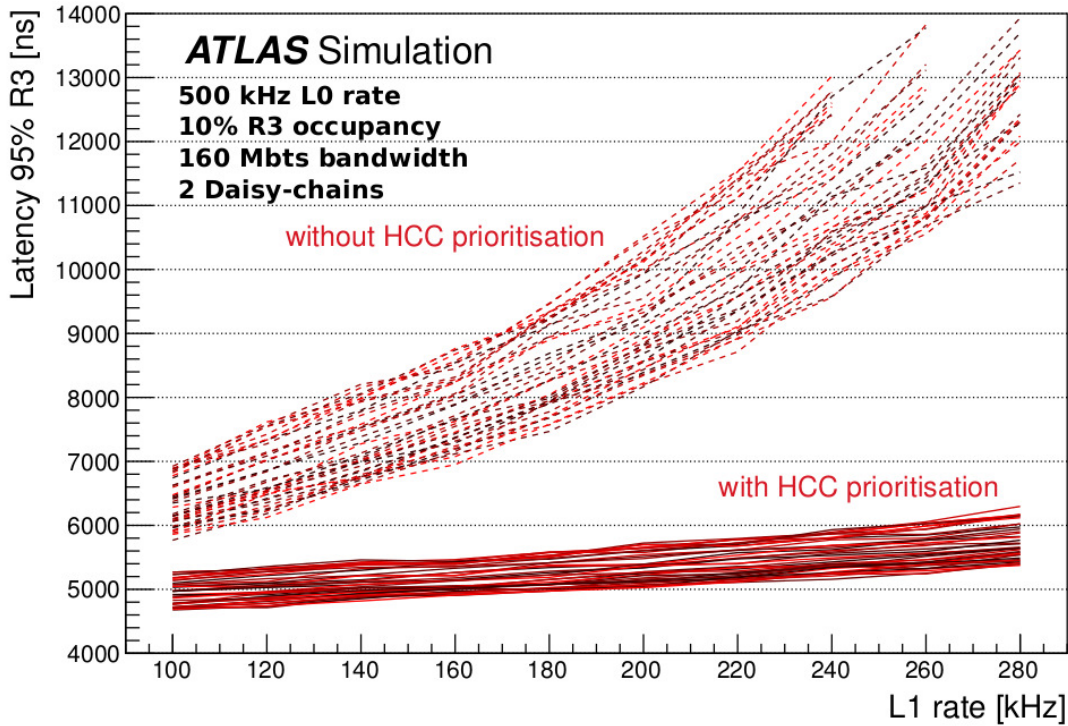


Figure 1: Latency of R3 requests in ITK's innermost barrel hybrids as a function of the L1 accept rate [13]. Each curve corresponds to a different hybrid, while the two family of curves differ in the implementation of R3 prioritisation, as discussed in the text.

7. Considerations on Track Finding and Reconstruction

The previous sections demonstrate that the relevant data can be transferred off-detector on-time for track reconstruction at Level-1. The next issue is the actual identification and reconstruction of tracks. Two approaches are being considered so far: the first one relies on studies performed for the FTK project [14]. Assuming that an architecture very similar to the FTK design [6] is used, the main point to address for the L1Track counterpart is how many associative memory patterns are required, as well as how many patterns are actually firing (as the latter will determine the latency of the track-fitting stage). The current FTK design (aimed at completion for the Phase I upgrade) is based on roughly 10^9 associative memory patterns. This is based on the ITK predecessor, and thus will need to be adapted to the use of 12 instead of 8 detector layers and a higher ($\sim 2 - 3\times$) pile-up. These factors will contribute to a higher number of patterns being activated per event. On the other hand, the number of patterns required and/or firing will be reduced by the fact that a higher transverse momentum threshold (2-4 GeV/c) can be used for Level-1 track finding, and by the use of Regions of Interest. Studies are in progress in order to determine if and how these mitigating factors are adequate.

An alternative approach abandons the RoI concept, and aims at on-detector or quasi-on-detector data reduction for L1 tracking: this could be accomplished for instance with the implementation of clustering algorithms in the front-end system, or the identification of "tracklets" with




Figure 2: Latency of R3 requests in ITK’s endcap hybrids as a function of the L1 accept rate [13]. Each numbered curve corresponds to a different hybrid. The continuous curves correspond to a realistic dimensioning of the HCC internal buffer, while the dashed lines are simulated assuming an infinitely large buffer to emphasise where the buffer is being saturated (separation of dashed and continuous lines).




Figure 3: Simulated latency of R3 requests in ITK’s endcap hybrids as a function of the L1 accept and R3 request rates [13]. The left histogram reports the latency when assuming the standard 160 Mbps stave link bandwidth, while the right plot shows the same latency when doubling the stave link bandwidth and doubling the number of daisy chain links. Both simulations assume HCC prioritisation, and simulate the "slowest" hybrid among those reported in figure 2.

the use of one or more pairs of ITK surfaces. Preliminary assessments [15] indicate that such approaches would effectively reduce the bandwidth requirements below 4 Gb/s per detector stave.

8. Conclusions and Outlook

ATLAS is planning the implementation of Level-1 track finding and reconstruction for the High-Luminosity-LHC era. This will provide additional flexibility to the whole TDAQ chain by reducing the bandwidth requirements for critical triggers. With L1Track, Level-1 trigger primitives for fixed thresholds will require a bandwidth reduced by as much as one order of magnitude, while trigger efficiencies for fixed bandwidth will improve by as much as a factor $\times 3$ [11]. The L1Track project is still under development, and the full specifications are expected by 2015, in order to support the full specification of the ATLAS inner tracker, which is going to be finalised with the publication of the ITK TDR in 2016.

References

- [1] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, 1008 JINST 3 S08003
- [2] L. Evans and P. Bryant (editors), *LHC Machine*, 2008 JINST 3 S08001
- [3] ATLAS Collaboration, *Prospects for benchmark Supersymmetry searches at the high luminosity LHC with the ATLAS Detector*, ATL-PHYS-PUB-2013-011, <https://cds.cern.ch/record/1604505>
- [4] ATLAS Collaboration, *Physics at a High-Luminosity LHC with ATLAS (Update)*, ATL-PHYS-PUB-2012-004, <http://cds.cern.ch/record/1484890>
- [5] A. Negri on behalf of the ATLAS TDAQ Collaboration, *Evolution of the Trigger and Data Acquisition System for the ATLAS experiment*, 2012 J. Phys.: Conf. Ser. 396 012033
- [6] ATLAS Collaboration, *Fast TracKer (FTK) Technical Design Report*, CERN-LHCC-2013-007 ; ATLAS-TDR-021, <https://cds.cern.ch/record/1552953>
- [7] R. Caputo et al., *Upgrade of the ATLAS Level-1 trigger with an FPGA based Topological Processor*, ATL-DAQ-PROC-2013-039, <https://cds.cern.ch/record/1628755>
- [8] M. Wessels on behalf of the ATLAS Collaboration, *Upgrade of the ATLAS Level-1 Calorimeter Trigger*, ATL-DAQ-PROC-2014-003, <https://cds.cern.ch/record/1703259>
- [9] ATLAS Collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, CERN-LHCC-2013-018 ; ATLAS-TDR-023, <https://cds.cern.ch/record/1602235>
- [10] ATLAS Collaboration, *New Small Wheel Technical Design Report*, CERN-LHCC-2013-006 ; ATLAS-TDR-020, <https://cds.cern.ch/record/1552862>
- [11] ATLAS Collaboration, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, CERN-LHCC-2012-022 ; LHCC-I-023, <https://cds.cern.ch/record/1502664>
- [12] ATLAS Collaboration, *The Expected Performance of the ATLAS Inner Detector*, ATL-PHYS-PUB-2009-002 ; <https://cds.cern.ch/record/1118445>
- [13] ATLAS Collaboration, *Approved plots for the L1Track Trigger project*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1TrackPublicResults>
- [14] C. Gentsos et al., *Future Evolution of the Fast TracKer (FTK) Processing Unit*, TIPP2014 proceedings, ATL-DAQ-PROC-2014-013, <https://cds.cern.ch/record/1710832>
- [15] ATLAS Collaboration, *A self seeded first level track trigger for ATLAS*, A Schöning 2012 JINST 7 C10010