Developments in the ATLAS Tracking Software ahead of LHC Run 2

Nicholas Styles[1], Massimiliano Bellomo[2], Andreas Salzburger[3], on behalf of the ATLAS collaboration

[1]Deutsches Elektronen-Synchtron (DESY), [2]University of Massachusetts, [3]CERN

E-mail: nicholas.styles@desy.de

Abstract. After a hugely successful first run, the Large Hadron Collider (LHC) is currently in a shut-down period, during which essential maintenance and upgrades are being performed on the accelerator. The ATLAS experiment, one of the four large LHC experiments has also used this period for consolidation and further developments of the detector and of its software framework, ahead of the new challenges that will be brought by the increased centre-of-mass energy and instantaneous luminosity in the next run period. This is of particular relevance for the ATLAS Tracking software, responsible for reconstructing the trajectory of charged particles through the detector, which faces a steep increase in CPU consumption due to the additional combinatorics of the high-multiplicity environment. The steps taken to mitigate this increase and stay within the available computing resources while maintaining the excellent performance of the tracking software in terms of the information provided to the physics analyses will be presented. Particular focus will be given to changes to the Event Data Model, replacement of the maths library, and adoption of a new persistent output format. The resulting CPU profiling results will be discussed, as well as the performance of the algorithms for physics processes under the expected conditions for the next LHC run.

1. Introduction

The first data-taking run of the Large Hadron Collider (LHC), which ended in 2013, was an enormous success. It provided almost 30 fb^{-1} of proton-proton collision data at centre-of-mass energies of $\sqrt{s}$ =7-8 TeV to the detectors. This data allowed a vast range of groundbreaking physics measurements to be made, including but not limited to the discovery of the Higgs Boson by the ATLAS and CMS experiments [1],[2]. It is planned that for the second run which will begin in 2015, the delivered luminosity will increase even further, with around 45 fb^{-1} per year at $\sqrt{s}$=13 TeV anticipated to be delivered to the experiments. This will be coupled with a move from a bunch spacing of 50 ns to 25 ns. During Run 2, the detectors will see an increase in the average number of interactions per bunch crossing, $\langle\mu\rangle$, starting at around 40 and potentially increasing up to 80 by the end of the run, compared to an average of around 25 for 2012 data, as shown in Figure 1. These multiple interactions per bunch crossing are referred to as 'pile-up'.

The ATLAS experiment [4] has used the long shutdown between Run 1 and Run 2 to perform various detector upgrades and consolidation. The ATLAS Inner Detector (ID) [5], as the subsystem closest to the interaction point, responsible for reconstructing the trajectory of charged particles, is particularly affected by increases in pile-up. The ID comprises three detector technologies; the high-granularity silicon Pixel detector forms the innermost part of the

Content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.
Published under licence by IOP Publishing Ltd

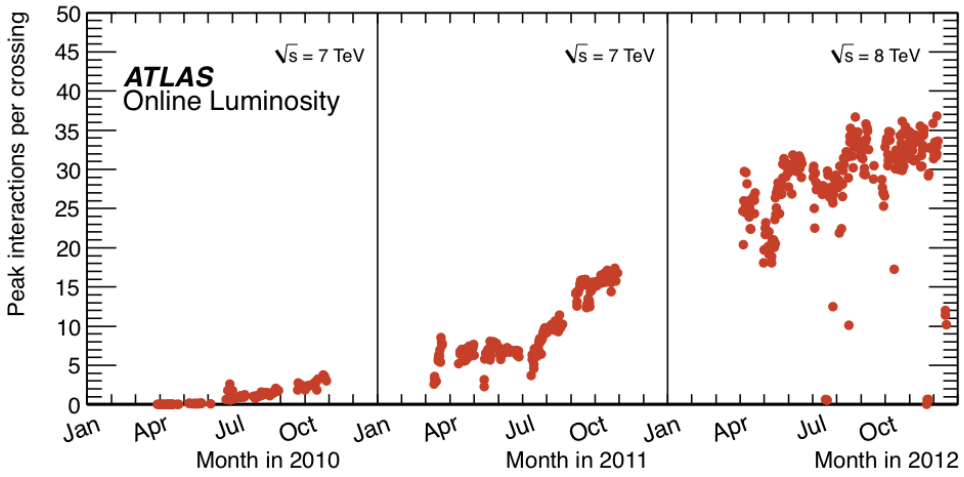

Figure 1. Evolution of the peak number interactions per bunch crossing observed by the ATLAS detector with time during the LHC Run 1 [3].

ID, followed by the silicon microstrip Semiconductor Tracker (SCT), and finally the Transition Radiation Tracker (TRT), based on drift tubes. During the long shutdown, a new innermost Pixel layer has been added to the detector, the Insertable B-Layer (IBL) [6], but this is far from the only activity going on during the shutdown period. Extensive work has also been undertaken to make significant improvements to the ATLAS event reconstruction software, part of the Athena [7] software framework. The ID track reconstruction algorithms [8] already take up a significant fraction of the total event reconstruction time, and the required CPU rises quickly as the multiplicity increases, exacerbating the situation further. Significant improvements to the ID tracking software were therefore essential to allow the goals for Run 2 data processing to be met; in particular, the requirement of 1 kHz first-pass event processing at the CERN Tier-0 computing centre [9].

2. Tracking Software Developments

In the following sections, the various developments made to the ATLAS ID Tracking software in order to help meet the 1 kHz Tier-0 processing target will be discussed.

2.1. Event Data Model Developments

The Event Data Model (EDM) describes the various elements of the track data in the ATLAS software framework. It is principally based upon the concept of `TrackParameters` objects, which describe the helix parameters of a track in a given reference frame. The previous EDM implementation [10] was analysed, and among the observations was that a large number of `dynamic_cast` operations were being used, necessitated by large number of virtual functions used in the EDM, which taken together constituted a non-negligible contribution to the overall CPU usage in track reconstruction. The EDM was redesigned using templated classes, and the concept of `MeasuredParameters` which inherited from TrackParameters, was removed, with measurements being distinguished from other types of TrackParameters by the presence of an associated covariance matrix. Through these changes, in addition to reducing the number of `dynamic_cast` operations, the EDM code base was also massively reduced in size, making it much easier to maintain in the future.

These changes were all made within the context of the development of a new persistent output

format for ATLAS reconstruction, which necessitated a re-design of the persistent output object of the track reconstruction, the `TrackParticle`. In redesigning this output object, priority was given to simplifying its inheritance structure and user interface. This was necessary in view of its future role as an object for use in data analyses, where previously there would have been further processing steps before reaching the objects used in analyses.

2.2. Maths Library Changes

Track reconstruction makes extensive use of linear algebra, and so there are potentially large savings available if the CPU spent performing these calculations can be reduced. The CLHEP maths library [11] was previously used throughout the ATLAS software, but after investigating potential alternatives (including SMatrix [12]), the decision was made to move to Eigen [13], a vectorised C++ maths library, for linear algebra within the ATLAS reconstruction software. Eigen was chosen since it offered the largest performance improvements for ATLAS use cases of the options investigated.

Macros, with optimisations for fixed size and symmetric matrices (which make up a large fraction of the matrices used in track reconstruction) were also developed to be used within the ATLAS track reconstruction, further improving the speed and efficiency of linear algebra calculations.

2.3. Track Seeding Optimisations

The first stage in reconstructing a track is the formation of track seeds from the various ID measurements. This was the focus of optimisation ahead of Run 2, leading to a number of changes in the strategy applied. Table 1 shows the percentage of seeds that result in a 'good' (i.e. satisfying various track quality criteria) track for various 3-measurement seed types (Pixel only, Pixel and SCT, and SCT only) at pile-up values of $\langle\mu\rangle = 0$ and 40. The presence of the IBL for run 2 allows more stringent requirements for seeds to be applied, which helps to mitigate the effect of increasing pile-up, and in the case of the best-performing seed type (3 SCT plus 1 extra measurement) the performance is actually improved with respect to the previous configuration at $\langle\mu\rangle = 0$.

Table 1. Percentage of seeds resulting in a good track being reconstructed, for the Run 1 (top row) and Run 2 (bottom row) configurations under different pile-up conditions. P represents a Pixel measurement being used in the seed, S an SCT measurement, and I an additional 4th measurement of either type.

$\langle\mu\rangle$	PPP	PPS	PSS	SSS
0	57%	26%	29%	66%
40	17%	6%	5%	35%
	PPP+I	PPS+I	PSS+I	SSS+I
0	79%	53%	52%	86%
40	39%	8%	16%	70%

During Run 1, the search region for seeds was defined by the maximum/minimum longitudinal impact parameter, while for Run 2 the region will be defined by the maximum/minimum vertex position in the longitudinal direction. This change was accompanied by tightening of track quality requirements.

A further development was made for the so-called 'TRT seeded tracks', which are reconstructed using measurement points not used during the initial seeding from Pixel and SCT measurements. TRT seeded track finding will only be processed for Run 2 in the case that a compatible seed cluster matching the track is found in the Electromagnetic calorimeter.

Together, these changes have resulted in a significant speed-up in the software, with no apparent loss of tracks useful for physics analyses.

2.4. Results of Developments

In implementing the changes described above (in addition to others not mentioned here,such as improved access to the magnetic field information), a huge number of software packages needed to be carefully modified, and in some cases substantially re-written. The result of this large effort by many people can be seen in Figure 2, which shows the average CPU time per event for reconstructing simulated $\sqrt{s} = 14$ TeV $t\bar{t}$ events with $\langle\mu\rangle = 40$, for different releases of the ATLAS reconstruction software.

Between releases 17.2.7.9 and 19.0.3.3, the changes to the EDM and Maths libraries were applied, in addition to other changes including moving from 32-bit to 64-bit, and applying compiler optimisations. Between 19.0.3.3 and 19.1.1.1, the optimisations of the track seeding strategy were applied (along with other changes). A significant, more than two-fold reduction in the overall event reconstruction time, and a corresponding reduction in the Inner Detector reconstruction time can be observed when going from 17.2.7.9 to 19.1.1.1, and the ~ 20 s average event reconstruction time is compatible with a 1 kHz Tier 0 event processing rate.

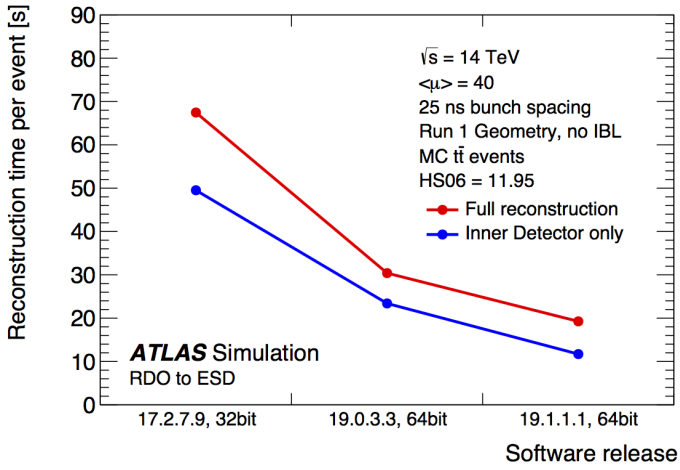

Figure 2. Average CPU time per event for reconstructing simulated $\sqrt{s} = 14$ TeV $t\bar{t}$ events with $\langle\mu\rangle = 40$ in various ATLAS software releases [14].

2.5. Developments required beyond Run 2

The High Luminosity LHC (HL-LHC) planned for the next decade, is foreseen to provide even higher integrated luminosities (a projected total dataset of 3000 fb^{-1}), which implies further increases in pile-up, with anticipated $\langle\mu\rangle$ of 140 or beyond.

A completely new tracking detector will replace the ID during the so-called 'Phase 2' upgrade [15]. In order to reconstruct events within whatever computing budget is available at the time, it is likely that dramatic changes will be required to the track reconstruction, to mitigate the 'combinatorial explosion' under such a high-multiplicity environment.

The current approach to track reconstruction is highly sequential, and so not ideally suited to make use of the increased parallelism offered by Many-Core/GPU technologies, and so this may imply that a fundamentally different approach is necessary in the future to make best use of available resources. Increased use of Vectorisation/SIMD approaches (i.e. within certain parts of the ATLAS code base itself) may also be a possible avenue of exploration.

3. Conclusions

An extensive program of software development has been undertaken within the ATLAS reconstruction framework during the LHC's long shutdown between Run 1 and Run 2, necessitated by the anticipated increase in pile-up for Run 2 LHC collisions, and the corresponding increases in CPU time for event reconstruction, in particular for the Inner Detector track reconstruction.

A variety of developments were made to the track reconstruction software, including moving to a redesigned Event Data Model, replacement of the Maths library for linear algebra with a vectorised library (Eigen), and optimisation of the track seeding strategy.

Together with changes elsewhere in the reconstruction software, this has resulted in a reduction by a factor > 2 of the CPU time required for event reconstruction. This will allow first-pass event reconstruction to run at the CERN Tier 0 site at a rate of 1 kHz.

References

[1] ATLAS Collaboration, G. Aad et al., *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys.Lett. **B716** (2012) 1–29, `arXiv:1207.7214 [hep-ex]`.

[2] CMS Collaboration, S. Chatrchyan et al., *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys.Lett. **B716** (2012) 30–61, `arXiv:1207.7235 [hep-ex]`.

[3] "ATLAS Experiment Luminosity - Public Results." https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults.

[4] ATLAS Collaboration, G. Aad et al., *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[5] ATLAS Collaboration, G. Aad et al., *The ATLAS Inner Detector commissioning and calibration*, Eur.Phys.J. **C70** (2010) 787–821, `arXiv:1004.5293 [physics.ins-det]`.

[6] ATLAS Collaboration, G. Aad et al., *ATLAS Insertable B-Layer Technical Design Report*, Tech. Rep. CERN-LHCC-2010-013. ATLAS-TDR-19, CERN, Geneva, Sep, 2010. https://cds.cern.ch/record/1291633.

[7] ATLAS Collaboration, G. Duckeck et al., *ATLAS Computing: technical design report*, Tech. Rep. CERN-LHCC-2005-022. ATLAS TDR-017, CERN, Geneva, 2005. https://cds.cern.ch/record/837738.

[8] T. Cornelissen, M. Elsing, I. Gavrilenko, W. Liebig, E. Moyse, and others on behalf of the ATLAS Inner Detector Software Group, *The new ATLAS track reconstruction (NEWT)*, J.Phys.Conf.Ser. **119** (2008) 032014.

[9] M. Elsing, L. Goossens, A. Nairz, and G. Negri, *The ATLAS Tier-0: Overview and operational experience*, J.Phys.Conf.Ser. **219** (2010) 072011.

[10] T. Cornelissen, N. Van Eldik, M. Elsing, W. Liebig, E. Moyse, et al., *Updates of the ATLAS Tracking Event Data Model (Release 13)*, ATL-SOFT-PUB-2007-003 (2007).

[11] "CLHEP - a class library for high energy physics." http://proj-clhep.web.cern.ch/proj-clhep/.

[12] "ROOT Mathematical Libraries - SMatrix Package." http://project-mathlibs.web.cern.ch/project-mathlibs/sw/html/SMatrix.html.

[13] "Eigen - a C++ template library for linear algebra: matrices, vectors, numerical solvers, and related algorithms." http://eigen.tuxfamily.org.

[14] "ATLAS Experiment Computing and Software - Public Results." https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults.

[15] ATLAS Collaboration, G. Aad et al., *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, Tech. Rep. CERN-LHCC-2012-022. LHCC-I-023, CERN, Geneva, Dec, 2012. https://cds.cern.ch/record/1502664.