ATLAS analysis workflows using the EventIndex and the Event Picking Server for massive event picking and enhanced processing

Dario Barberis[1,2,*], *Evgeny* Alexandrov[3], *Igor* Alexandrov[3], *Aleksandra* Dimitrievska[4], *Alexander* Iakovlev[3], *Simone* Pagan Griso[5], and *Semen* Turchikhin[1]

[1]University of Genoa and INFN, Genoa, Italy
[2]University of California, Berkeley CA, USA
[3]Joint Institute for Nuclear Physics, Dubna, Russia
[4]University of Birmingham, Birmingham, UK
[5]Lawrence Berkeley National Laboratory, Berkeley CA, USA

Abstract. The ATLAS detector produces a wealth of information for each recorded event. Standard calibration and reconstruction procedures reduce this information to physics objects that can be used as input to most analyses; nevertheless, there are very specific analyses that need full information from some of the ATLAS subdetectors, or enhanced calibration and/or reconstruction algorithms. For these use cases, a novel workflow has been developed that involves the selection of events satisfying some basic criteria, their extraction in RAW data format using the EventIndex data catalogue and the Event Picking Server, and their specialised processing. This workflow allows us to commission and use new calibration and reconstruction techniques before launching the next full reprocessing (important given the longer and longer expected time between full reprocessing campaigns), to use algorithms and tools that are too CPU or disk intensive if run over all recorded events, and in the future to apply AI/ML methods that start from low-level information and could profit from rapid development/use cycles. This paper describes the tools involved, the procedures followed and the current operational performance.

ATL-SOFT-PROC-2025-035
22 February 2025

1 Introduction

The ATLAS experiment [1] collects several billion proton-proton, proton-ion and ion-ion interactions at the LHC accelerator at CERN every year. These "events" are then processed several times, resulting in many different formats and versions of the same original information. During LHC Run 2 (2015-2018), for the main proton-proton interaction stream ATLAS recorded 19 billion events, grouped in 9 million files, for a total of 20 PB of raw data. During LHC Run 3 so far (2022-2024) the size of this dataset has already more than doubled, and more will come in 2025 and 2026. In addition, simulated events are generated to compare the results of the analysis of real data with different physics models.

*Presenter and corresponding author: Dario.Barberis@cern.ch

The EventIndex [2] system catalogues all ATLAS events, both real and simulated, and provides a set of tools to search and retrieve information on single events or on event groups, following user selections. Its design started in 2013 [3] and the first implementation was fully functional for the start of LHC Run 2 in 2015 [4]. The original version of the EventIndex provided a stable and reliable service throughout LHC Run 2, but like all software projects it had to be upgraded in order to stand the higher data rates for Run 3 and beyond [5].

ATLAS event data are processed to reconstruct physics analysis objects from the raw data, applying at the same time detector calibration and alignment procedures [6]. The data formats to be used for physics analyses (Analysis Object Data (AOD) and Derived Analysis Object Data (DAOD)) are by necessity reduced to contain only the physics quantities needed by the vast majority of the analyses, with the suppression of detector-level information that has already been used in the event reconstruction procedures.

Nevertheless, there are analyses that need to access such detector-level information, either to apply more refined calibrations or to execute additional algorithms to reconstruct more physics quantities, particularly particles with low transverse momentum that are only needed by a handful of analyses. These needs can be addressed by the novel workflow described in this paper, which makes substantial use of the Event Picking Server.

2 Event Picking Server

The Event Picking Server (EPS [7, 8]) is a tool designed and developed since 2019 to automate the event picking workflow, especially when users need to extract several thousand events at once. Its entry point for the users is a graphical interface where the user supplies a text file with the list of run and event numbers of all events to be retrieved, the event format and data (trigger) stream. Behind the scenes, a process splits the event list by run number, uses the EventIndex Query Service [9] to get the locations (as GUIDs [10]) of the files with those events, submits the event picking jobs to the PanDA workflow management system [11] that will run the jobs on the WLCG Grid, and finally notifies the user upon completion.

As the input files containing the events to be extracted are usually stored on tape at some WLCG data centre, this procedure may last several days (depending on the tape queues at that site). The retrieved events are grouped into datasets and made available in the EOS storage space [12] at CERN. The graphical interface can also be used to monitor the progress of the requests, which is a very useful feature given that the completion time can be of the order of days, or even weeks, when retrieving many raw data from files on tape.

3 Massive Event Picking Workflow

This novel workflow consists of three main stages:

1. The analysis team selects the events on the basis of all properties, except those involving objects to be reprocessed. In this way the number of events passed to the following stages is reduced to the bare minimum, but still has to contain events in the signal and background control and validation regions. As a rule of thumb, to avoid the event picking operation turning into a massive reprocessing campaign one should not request to stage from tape more than 10% of the raw data files; given the average number of events in ATLAS files, this means limiting the request to $O(10^{-4})$ of the total number of events (O(1M) events for the Run 2 data). The result of this first stage is a list of events in a text file, with a run-event number pair in each line.

2. The event list is given to the Event Picking Server, together with additional useful information to speed up the operations: year of data-taking, trigger stream, data format to be retrieved (usually raw data). The extracted events will be available in EOS storage at CERN, where data size is not a concern as long as we handle reasonable numbers of events as described above; for example, one million events occupy between 1 and 2 TB of disk space (depending on the data-taking year) in raw data format.

3. The extracted events are finally reprocessed with enhanced software and/or calibrations, producing new versions of AODs and DAODs containing the additional information to be used for the final analysis. Both the processing time and the output file sizes are not of concern, as long as the number of events is limited. At this point the analysis procedure can continue, using all necessary information.

4 Application examples

The first ATLAS analysis using massive event picking was the study of the $\gamma\gamma \rightarrow WW$ scattering in proton-proton interactions. After the first analysis round [13], it was realised that the remaining background in the signal region could be substantially reduced if particles with transverse momentum below the standard reconstruction threshold (700 MeV) could also be reconstructed and made available at analysis level. The "low-p_T" reconstruction algorithm can reconstruct tracks with $p_T > 100$ MeV with an acceptable efficiency and low fake rates, but it is slow and produces many reconstructed tracks; therefore it cannot be run on the complete data sample, but it can be run on pre-selected data. The procedure described in this paper was then (2019) developed and run manually on a small initial data sample of 50k events. In the meantime, the possibility to automate this procedure with the creation of the Event Picking Server was discussed, and in 2020 the first version of the EPS was used to extract and then process an additional sample of 100k events.

The decays $B_c^{\pm *} \rightarrow B_c^{\pm}(\rightarrow J/\psi\mu^{\pm}\nu_{\mu})\gamma$ produce low-energy photons that cannot be efficiently detected and measured in the ATLAS calorimeters. If the photon instead converts to an e^+e^- pair, it can be measured if the low-momentum electron and positron can be reconstructed, again using the low-p_T reconstruction algorithm. That analysis team used the EPS in 2023 to extract 650k events from the raw data collected between 2015 and 2018 and is now processing and analysing them.

Another analysis team is searching for heavy long-lived massive particles that are predicted to exist by Supersymmetric and other theoretical models, exploiting in particular the measurements of specific ionisation ($\mathrm{d}E/\mathrm{d}x$) in the four layers of the ATLAS Pixel detectors. During LHC Run 2 (2015-2018) the Pixel detectors, that are very close to the beam line and the interaction region, suffered from radiation damage and the average $\mathrm{d}E/\mathrm{d}x$ measurement for each track was corrected according to the integrated luminosity accumulated by the detector when each event was recorded [14]. With the higher LHC luminosity during Run 3 operation (2022-2026), correcting the average $\mathrm{d}E/\mathrm{d}x$ measurement will no longer be sufficient, as different Pixel detector modules suffer very different radiation damage; the correction has to be applied to each module individually, before calculating the average $\mathrm{d}E/\mathrm{d}x$. This means that the cluster measurements need to be propagated to the analysis level, increasing the event sizes substantially. As disk space strongly constrains any increase in data sizes, the solution is to analyse all events up to the point where the $\mathrm{d}E/\mathrm{d}x$ information is needed, extract those events using the EPS, run a modified reconstruction algorithm that saves the pixel cluster information only for the selected tracks, apply the charge corrections and then continue the analysis. The EPS will be used in late Winter 2024-2025 to extract over one million events belonging to the data-taking years 2015-2024.

5 Conclusions

The workflow described in this paper makes it possible to run non-default reconstruction algorithms on subsets of events even in case the execution of these algorithms requires additional computing time or increases the size of the output data files. This can improve the reach of existing analyses, and also allows detailed study of the effects of changes in calibrations and reconstructions on the data samples that may be most sensitive to these changes.

The analysis examples provided here are just the start of a series of analyses that can benefit from such workflows in the future, when the global data samples will become so large that it will be difficult to reprocess all the data every time new improvements become available to implement.

References

[1] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003. https://doi.org/10.1088/1748-0221/3/08/S08003

[2] Barberis D. et al., *The ATLAS EventIndex: A BigData Catalogue for All ATLAS Experiment Events*, *Comput.Softw.Big Sci.* **7** (2023) 1,2. https://doi.org/10.1007/s41781-023-00096-8

[3] Barberis D. et al., *The ATLAS EventIndex: an event catalogue for experiments collecting large amounts of data*, *J.Phys.Conf.Ser.* **513** (2014) 042002. https://doi.org/10.1088/1742-6596/513/4/042002

[4] Barberis D. et al., *The ATLAS EventIndex: architecture, design choices, deployment and first operation experience*, *J.Phys.Conf.Ser.* **664** (2015) 4, 042003. https://doi.org/10.1088/1742-6596/664/4/042003

[5] Gallas E.J. et al., *Deployment and Operation of the ATLAS EventIndex for LHC Run 3*, *EPJ Web Conf.* **295** (2024) 01018. https://doi.org/10.1051/epjconf/202429501018

[6] ATLAS Collaboration, *Software and computing for Run 3 of the ATLAS experiment at the LHC*, to be published in *EPJC* (2025). https://arxiv.org/pdf/2404.06335

[7] Alexandrov E. et al., *Development of the ATLAS Event Picking Server*, *Proc. 9th Int. Conf. "Distributed Computing and Grid Technologies in Science and Education" (GRID'2021)* (2021) Dubna (Russia). https://doi.org/10.54546/MLIT.2021.35.43.001

[8] Alexandrov E. et al., *The ATLAS Event Picking Service and Its Evolution*, *Phys.Part.Nucl.* **3** (2024) 437-440. https://doi.org/10.1134/S1063779624030055

[9] Rybkin G., *Query Service for New ATLAS EventIndex System*, *EPJ Web Conf.* **295** (2024) 01045. https://doi.org/10.1051/epjconf/202429501045

[10] GUID: https://www.rfc-editor.org/rfc/pdfrfc/rfc4122.txt.pdf

[11] Barreiro Megino F.H. et al., *PanDA for ATLAS distributed computing in the next decade*, *J.Phys.Conf.Ser.* **898**:052002 (2017). https://doi.org/10.1088/1742-6596/898/5/052002

[12] EOS: https://eos-docs.web.cern.ch

[13] ATLAS Collaboration, *Observation of photon-induced W^+W^- production in pp collisions at $\sqrt{s}$ = 13 TeV using the ATLAS detector*, *Phys.Lett.B* **816** (2021) 136190. https://doi.org/10.1016/j.physletb.2021.136190

[14] ATLAS Collaboration, *Search for heavy, long-lived, charged particles with large ionisation energy loss in pp collisions at $\sqrt{s}$ = 13 TeV using the ATLAS experiment and the full Run 2 dataset*, *J. High Energ. Phys.* **2023** (2023) 158. https://doi.org/10.1007/JHEP06(2023)158