Investigating the performance of ROOT-based data formats. Exploring storage settings and cloud solutions.

Caterina Marcon

8.IV.2022

INFN MILANO
Istituto Nazionale di Fisica Nucleare
Sezione di Milano

ATLAS EXPERIMENT

Optimizing ATLAS simulation workflow

Optimizing ATLAS simulation workflow: impact of different build types

- Comparison between multi-library dynamic, single dynamic library and static builds;
- Static build shows an improvement of up to 7%;
- The single library approach exhibits an increase of ~ 10% in executuion time;
- Promising results have been obtained with a single static library: improvement up to 6-7%.

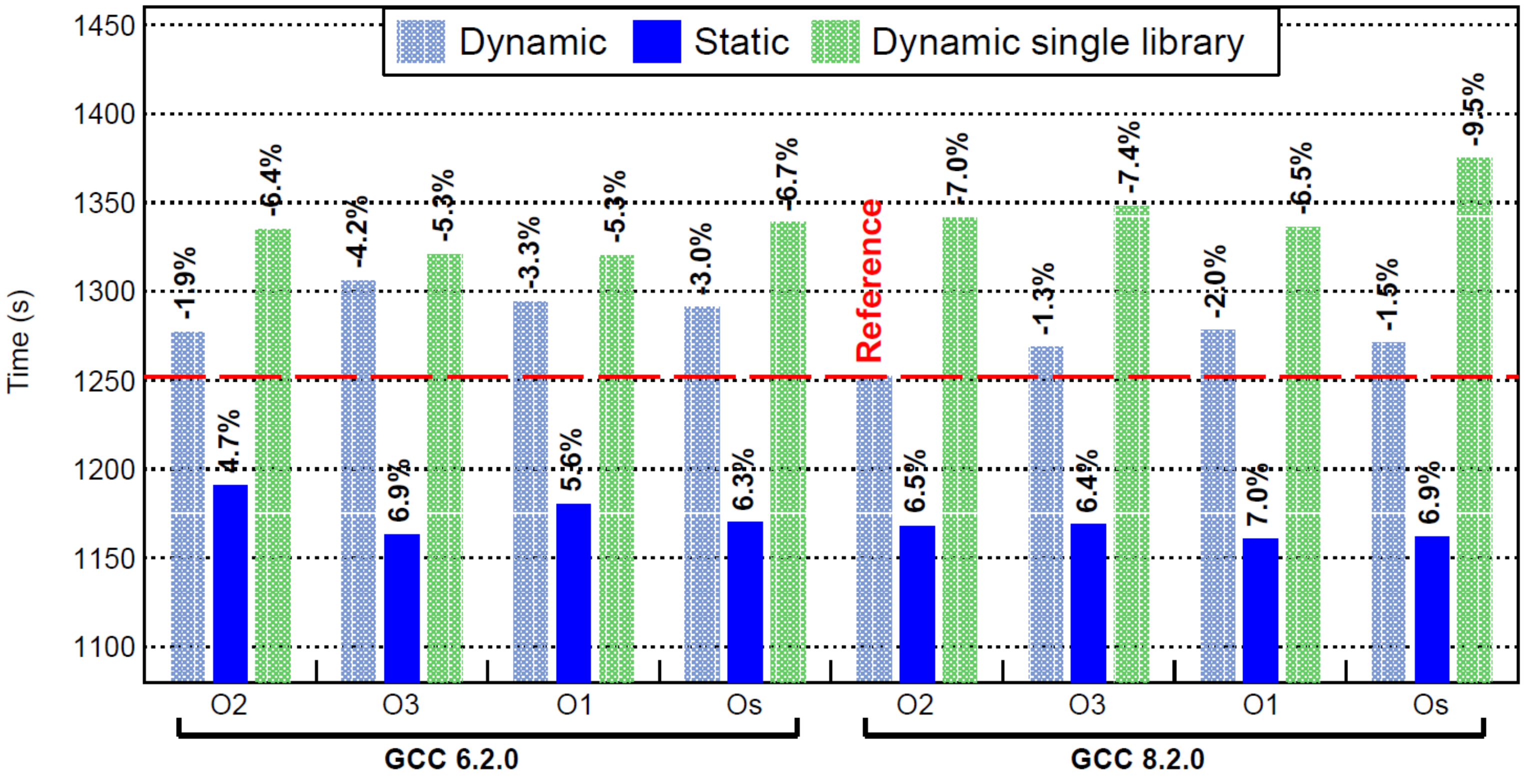

[1] https://indico.cern.ch/event/1106118/contributions/4693145/attachments/2376280/4059683/G4TF_Report_20Jan22_Final.pdf
[2] https://doi.org/10.1051/epjconf/202024505037
[3] https://doi.org/10.1051/epjconf/202125103005

Optimizing ATLAS simulation workflow: implementation of new shapes

- Ongoing effort to increase the efficiency of the ATLAS TRT geometry description;

- TRT is composed of 96 trapezoidal modules grouped in 3 types characterized by an increasingly larger cross-sectional area and the design comprises 1642 straws, which must be simulated individually;

- The current implementation takes advantage of boolean solids: two triangular prisms are merged together by their common face;

- In recent developments, two new shapes have been made available:
    - the arbitrary trapezoid Arb8;
    - the boundary representation Brep.

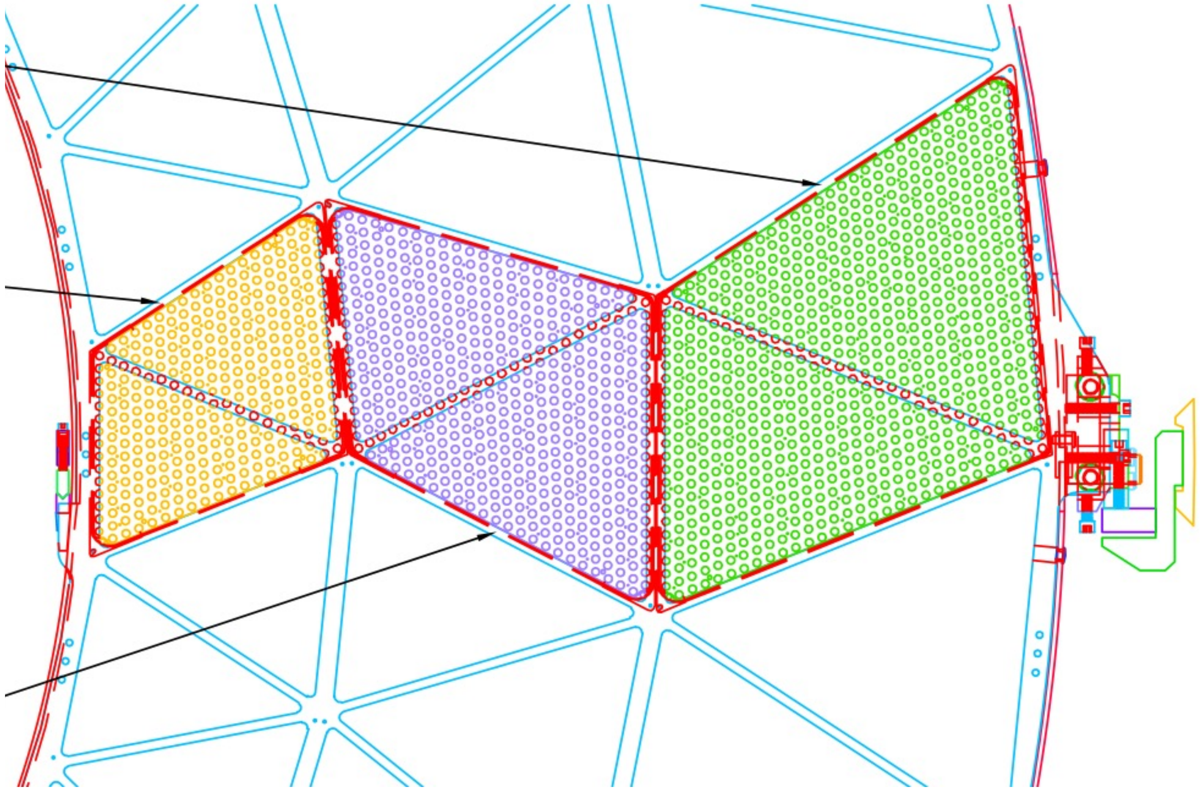

- A positive improvement of 1.5% is observed for the Arb8 representation, whereas the BRep solid exhibits a minor degradation with respect to the reference boolean solids.

Optimizing ROOT storage settings for DAOD PHYSLITE.

Optimizing ROOT storage settings for DAOD PHYSLITE: motivations

- ATLAS must also address the problem of data derivation storage. During Run 2 ATLAS deployed an analysis model that produced large data duplication:
  - Primary Analysis Object Data (AOD) taking up to 30% of total storage;
  - Derived AOD occupying another 40%.

- From Run 3 more compact analysis data formats able to reduce data duplication will be considered:
  - DAOD_PHYS: with a maximum size of 50 kB/event;
  - DAOD_PHYSLITE: with a size of 10/15 kB/event.

Optimizing ROOT storage settings for DAOD PHYSLITE: purpose

- Estimate the impact that different file compression algorithms and storage options provided by ROOT have on:
  - The primary xAOD files;
  - The derivative formats (DxAOD files);
  - The DAOD_PHYS and DAOD_PHYSLITE new data formats (in progress).

- Evaluate the impact on:
  - file size;
  - reading speeds from disk.

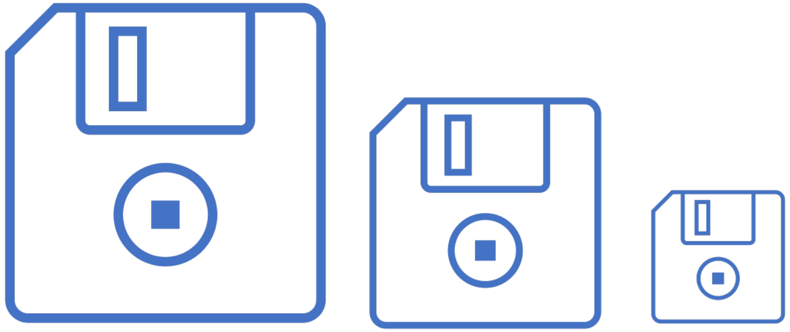

Optimizing ROOT storage settings for DAOD PHYSLITE: methods

- So far, the measurements are done for a primary xAOD file and three different derivation samples:
  - Primary xAOD (ttbar events);
  - Big DxAOD (TOPQ1, ~170kB per event on average);
  - Medium DxAOD (SUSY5, ~60kB per event on average);
  - Small DxAOD (TRUTH3, ~5kB per event on average).

- The measurements consider different compression algorithms provided by ROOT:
  - LZ4 (9 different compression levels);
  - ZLIB (9 different compression levels);
  - LZMA (9 different compression levels).

- All tests in this study collect I/O performance metrics while reading the test files from disk.

Optimizing ROOT storage settings for DAOD PHYSLITE: preliminary results

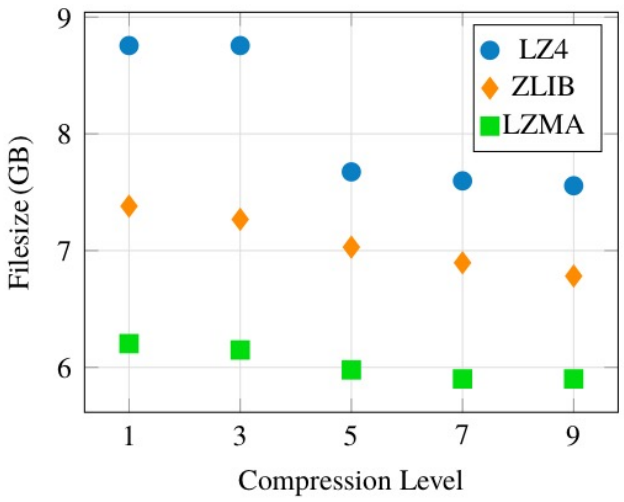

(a) xAOD

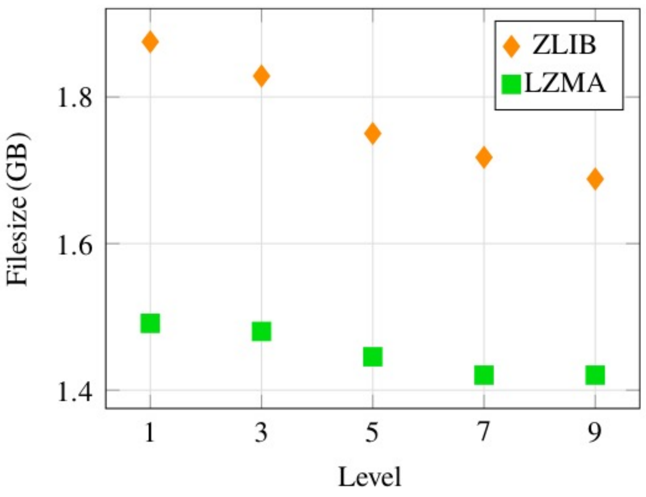

(b) TOPQ1

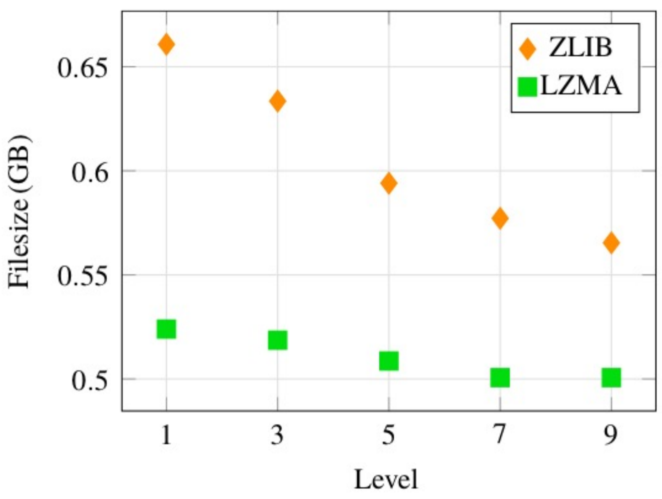

(c) SUSY5

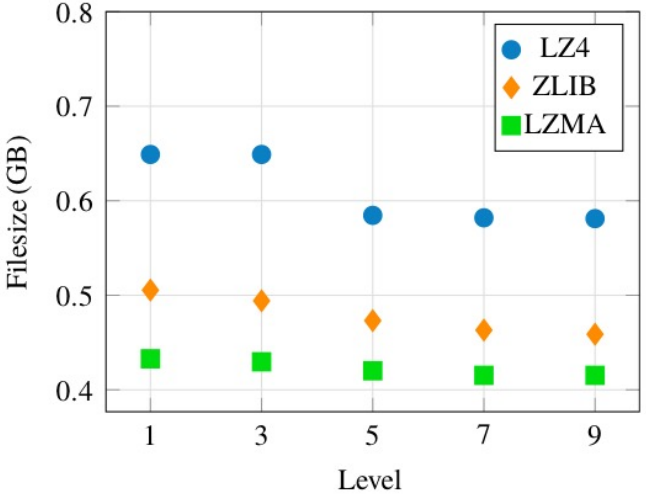

(d) TRUTH3

- All three algorithms have the same characteristic behaviour, independent of the event size;
- LZ4 is the fastest in reading but results in largest files;
- LZMA provides much better compression, at the cost of slower reading speeds;
- ZLIB represents an interesting option for very small events such as DAOD_PHYS and DAOD_PHYSLITE.

Optimizing ROOT storage settings for DAOD PHYSLITE: preliminary results

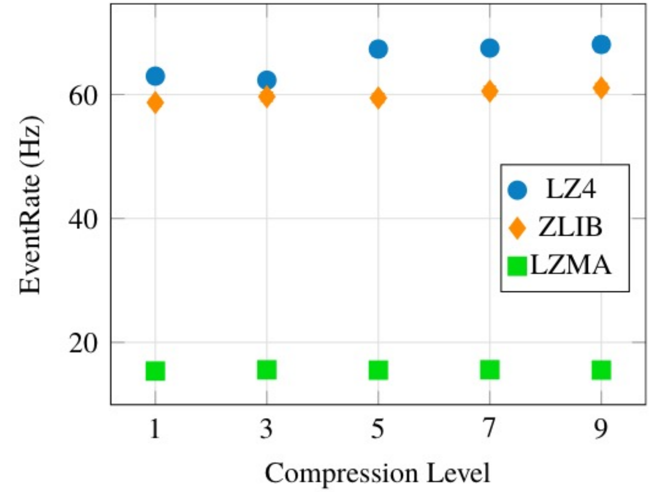

(a) xAOD

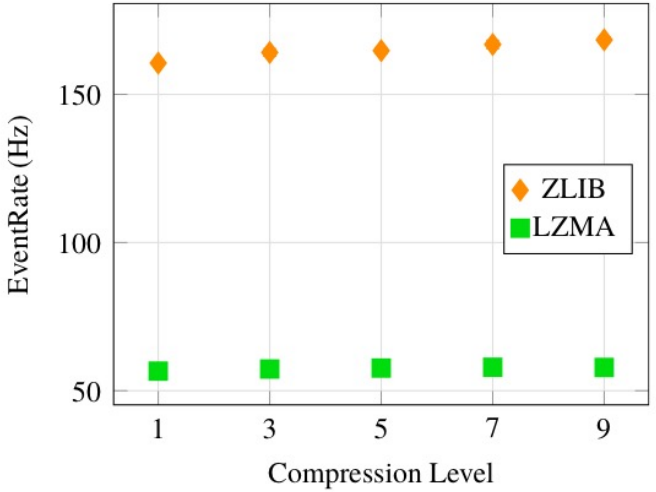

(b) TOPQ1

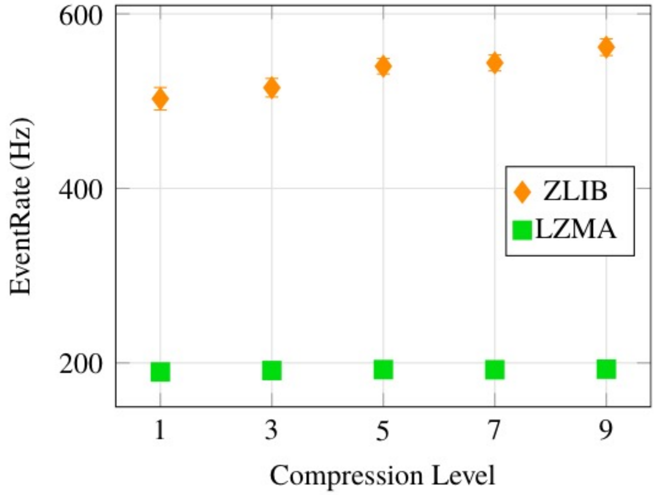

(c) SUSY5

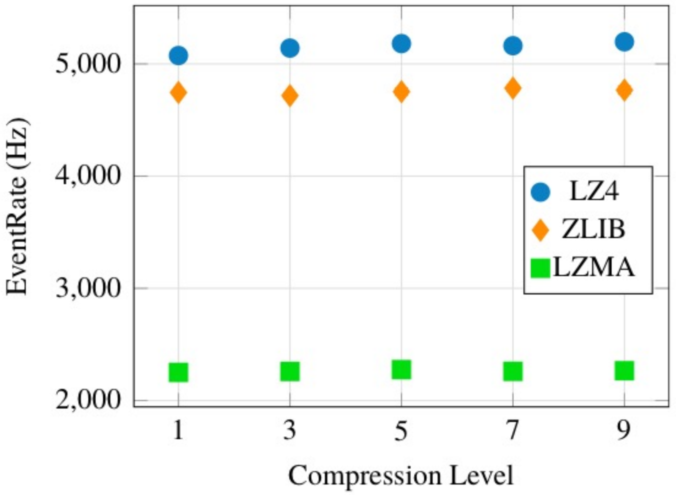

(d) TRUTH3

- All three algorithms have the same characteristic behaviour, independent of the event size;
- LZ4 is the fastest in reading but results in largest files;
- LZMA provides much better compression, at the cost of slower reading speeds;
- ZLIB represents an interesting option for very small events such as DAOD_PHYS and DAOD_PHYSLITE.

ATLAS analysis workflow on cloud resources

ATLAS analysis workflow on cloud resources

- Implement and test an ATLAS analysis workflow on cloud resources (for example using CNAF's resources):
- Caveat: ATLAS is not new to this type of approach: some tools have already been developed (such as batch jobs scripts, VM provisioning);
- Aspects that can/should be investigated:
  - distribution of the ATLAS code -> code/functions modularity;
  - resource allocation;
  - data access.

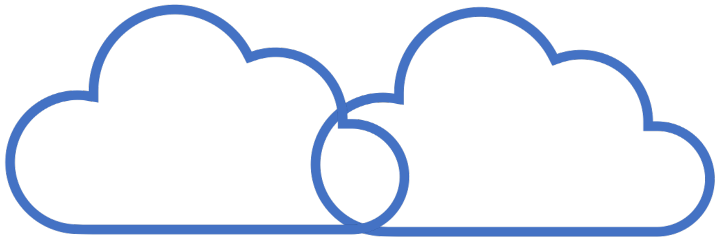

Thank you for the attention