ALICE Run 3 Analysis Framework

Anton Alkin, Giulio Eulisse, Jan Fiete Grosse-Oetringhaus, Peter Hristov, Maja Kabus

25th International Conference on Computing in High-Energy and Nuclear Physics, 17-21 May 2021

Run 3 challenges

- ALICE: heavy ion physics, study of strong interaction and QGP properties
- Continuous 50 kHz readout and Pb–Pb interaction rate in Run 3/4 [1]
- Tracking, detection of low-momentum particles, PID, vertexing
- Target luminosity of 6-7 nb^{-1}

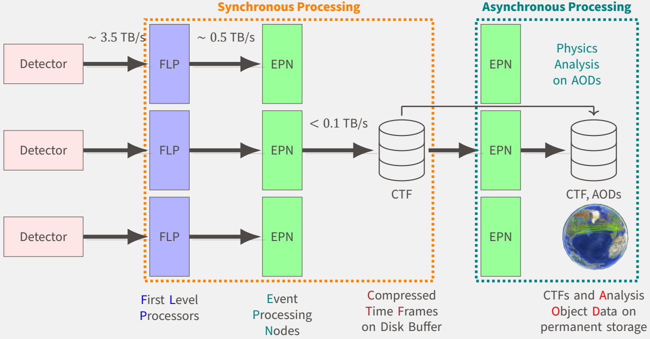

- 1 month of Pb–Pb data taking would create ~ 4 PB of AODs

ALICE O^2 Framework[2]

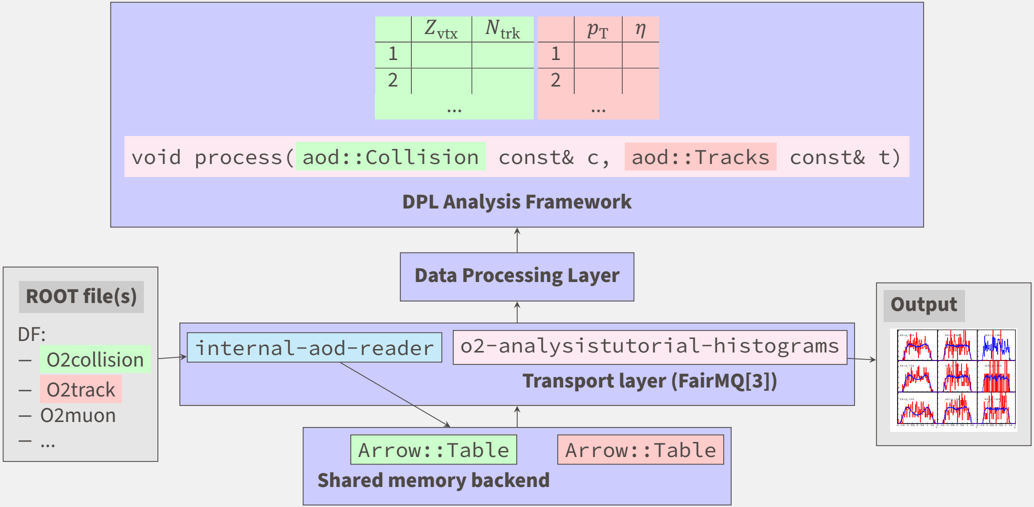

Analysis Data Model

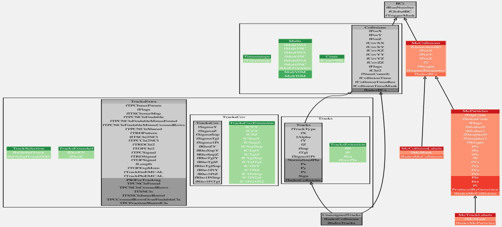

Columns

	X	α	f(X, Z, m)	Index	Z = X sin α
1				2	
2				3	
	Static		Dynamic	Index	Expression
	Arrow::Array		lambda function	Arrow::Array	Arrow::Array created in memory with Gandiva[4]

	A	B
1		
2		
3		

Interconnected tables

- Self-contained (Tables), as collections of Columns, connected by indices passed through shared memory
- Represented as ROOT `TTree` [5] on disk and as Apache Arrow `Table` [6] in memory
- Hierarchy of indices represents logical connections among data Tables (Tracks →Collisions →BCs)
- Columns and Tables are represented by C++ types for the end user resulting in negligible performance overhead

Computing Model

Workflows

- `Workflow` is a collection of connected data processors representing a particular analysis or group of analyses
- They can be dynamically formed, topology is computed automatically, based on known data subscriptions
- Data subscriptions are set by analyzers simply by stating in the code which Tables are required for a particular analysis

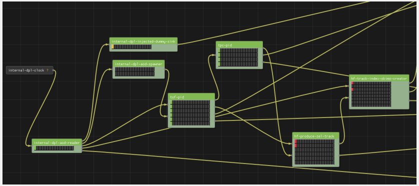

ALICE Hyperloop train system

- Web-based system for managing the running of workflows
- Implements benchmarking, configuration bookkeeping, job submission to the GRID or Analysis Facility[7] and production of derived datasets (skims)

Table Manipulation

Database-like operations

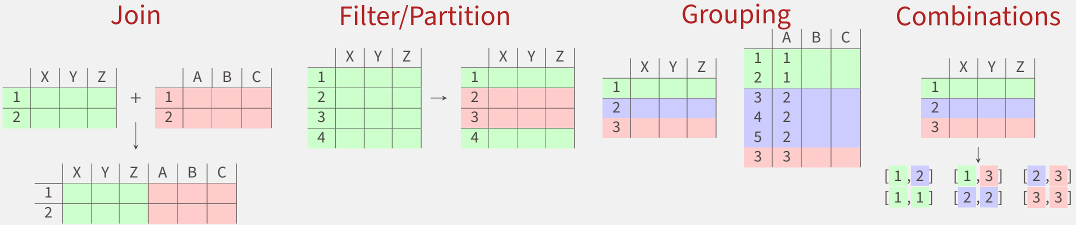

- All operations are zero-copy due to Apache Arrow backend
- Analyzer can directly request joined, grouped, partitioned of filtered table as an input to their task, combining all four operations if needed
- It is possible to inspect 2-, 3- and more rows combinations of a particular table without nested loops or memory caches, by using combinations generator
- A traditional "event loop" interface is also provided

Interoperability

ML applications

- Apache Arrow Tables are already supported or planned to be supported as an input in the popular ML toolkits
- O^2 framework allows for integration of external C++ and Python toolkits directly into processing

External tools

- Derived data and analysis results are stored as ROOT trees and histograms and are readily available for post-processing

Current Status

- More than 30 physics analyses were ported to O^2 using converted Run 1 and 2 data
- The code is tested locally, in the GRID and on Analysis Facility at GSI
- O^2 implementations show on average 3 to 10 times better throughput in terms of events/s compared to legacy framework
- More optimizations are foreseen (for example bulk reading in ROOT)

Benchmark	AliPhysics	O^2
Histograms	3.9	38
Correlations	2	21
Spectra	4.6	14
	(events/s)	

Future Development

- Work is ongoing to provide ONNX runtime[8] and Python tools with a convenient interface for end users, operating directly on data model tables. In principle, DPL allows to plug into a workflow any kind of FairMQ-based data processor, permitting direct integration of almost any external toolkit
- Extension of expression DSL to support more mathematical functions and conditional expressions
- Extension of combinations generator so that it can be used together with other table manipulation features (required, for example, in mixed event analyses)
- Sharing of configuration between tasks to avoid duplication or inconsistencies in the long processing chains
- Current goal is to be able to fully reproduce most of Run 1/2 analyses with O^2 framework directly on converted data and perform preliminary Run 3 analyses on with simulated data by autumn of 2021

References

[1] P. Buncic, M. Krzewicki and P. Vande Vyvre, CERN-LHCC-2015-006.

[2] *ALICE Online-Offline software framework, O^2, main repository*, https://github.com/AliceO2Group/AliceO2 (2021), accessed: 2021-05-10

[3] M Al-Turany et al, J. Phys.: Conf. Ser. **513** 022001 (2014)

[4] *Gandiva: A LLVM-based Analytical Expression Compiler for Apache Arrow*, https://arrow.apache.org/blog/2018/12/05/gandiva-donation/ (2018), accessed: 2021-02-24

[5] R. Brun, F. Rademakers, Nucl. Inst. Meth. **A389**, 81 (1997)

[6] *Apache Arrow, a cross-language development platform for in-memory analytics*, https://arrow.apache.org/ (2021), accessed: 2021-02-24

[7] K.Schwarz, EPJ Web Conf. **214**, 08027 (2019) doi:10.1051/epjconf/201921408027

[8] *ONNX Runtime is an open source project that is designed to accelerate machine learning across a wide range of frameworks, operating systems, and hardware platforms*, https://www.onnxruntime.ai/ (2021), accessed: 2021-05-10