Hyperloop – The ALICE analysis train system for Run 3

Raluca Cruceru[+], Costin Grigoras[+], Jan Fiete Grosse-Oetringhaus[+], **Raquel Quishpe*** on behalf of the ALICE Collaboration
[+]CERN, *University of Houston

LHCP 2021 Online

ALICE from Run 2/LEGO trains to Run 3/Hyperloop trains

The ALICE analysis framework allows the creation and submission of analysis jobs to the Grid. These jobs are monitored and resubmitted using MonALISA. In Run 1 and 2, ALICE developed a system of analysis trains to optimize the usage of computing resources.

LEGO trains

- Tool to run and manage analysis trains on AliEn providing a friendly user interface
- Analysis code is contained in an AliEn package, **AliPhysics**
- Analysis tasks (wagons) using the same dataset are run together
- Main workhorse for Run 2 analysis
  2020: 16 000 trains, 172 million Grid jobs, 39 000 average job slots

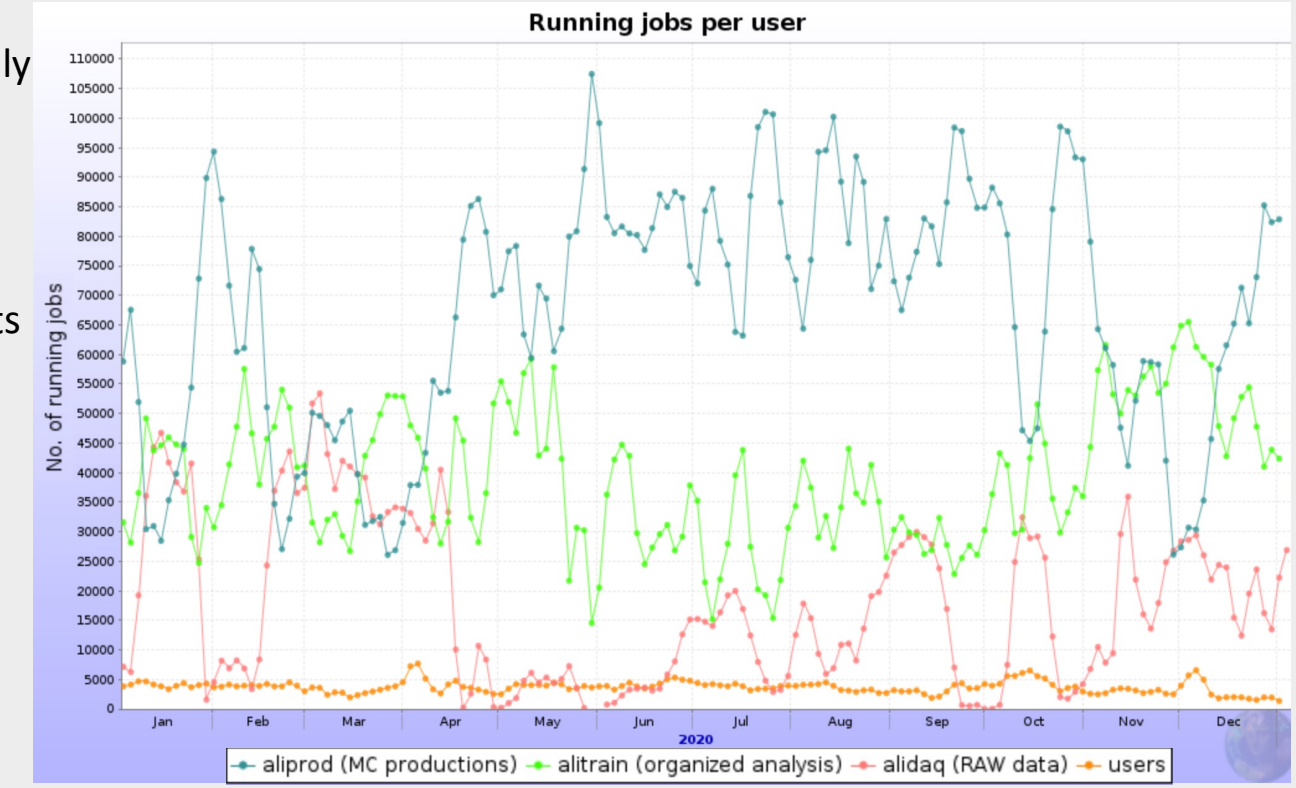

Long-term plan

- During LS2, ALICE went through upgrades allowing to collect in Run 3 roughly 2 orders of magnitude more data than in Run 2
- Run 3 data analysis uses **ALICE Online-Offline** (**O^2**) framework
- Run 2 data converted to Run 3 format, analysis moves to **O^2**
- **Hyperloop trains** have been developed based on the solid LEGO trains concept and integrated in the **O^2 workflow**
- LEGO trains phased out and replaced by **Hyperloop trains**

LEGO trains

- Analysis train framework for Run 2
- Database: PostgreSQL, backend: MonALISA, frontend: static HTML
- Trains are divided per Physics Working Group (PWG), data type and collision system (87 active)
- Same user interface for all users and train operators
- Requires train operators (per PWG) to test, compose and submit train runs

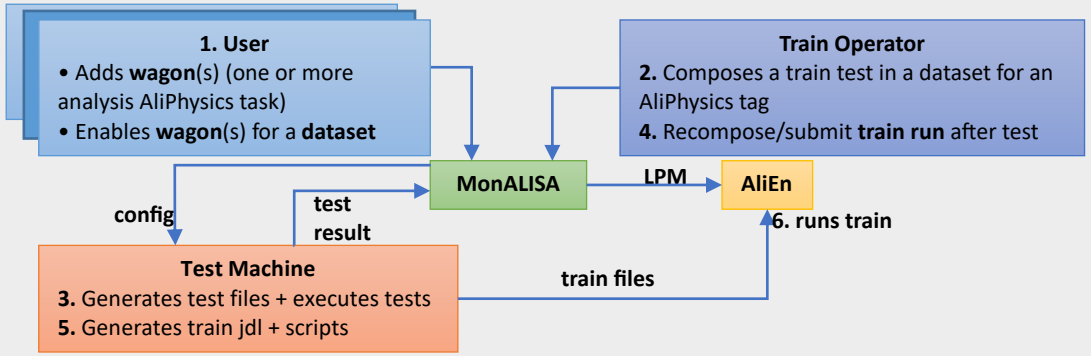

Wagon configuration

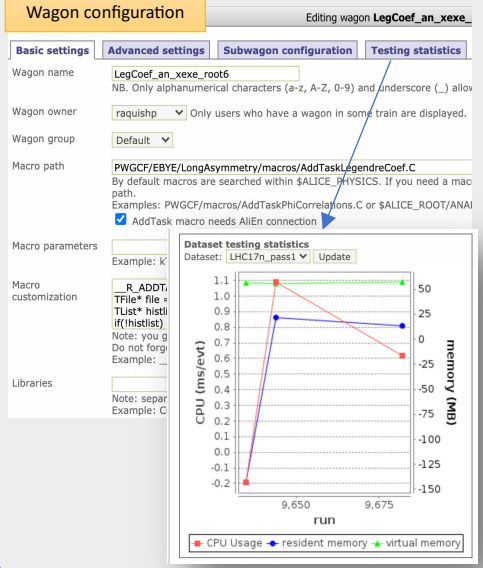

Train run result

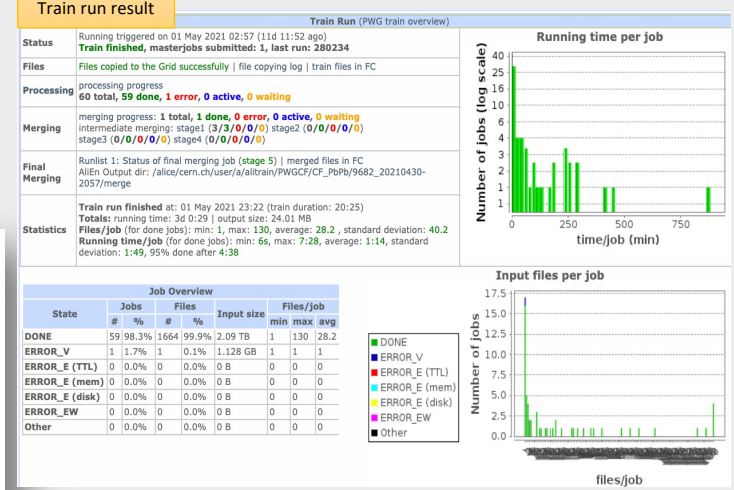

Hyperloop trains

- Analysis train framework for Run 3
- Database: PostgreSQL, backend: MonALISA + java-based model, frontend: **React.js**
- Unified trains throughout PWGs
- Personalized user and operator interfaces, saves page configuration
- **Analysis** defined in **JIRA**, wagons are defined per analysis
- Wagons (O^2 workflow) can be enabled for an available O^2 tag or **pull request**
- Immediate and **automatic wagon test** ⏳→ ⏳→ (🌟, ❗,💣)
- Wagon and dataset **bookkeeping**, allows wagons and train runs comparison

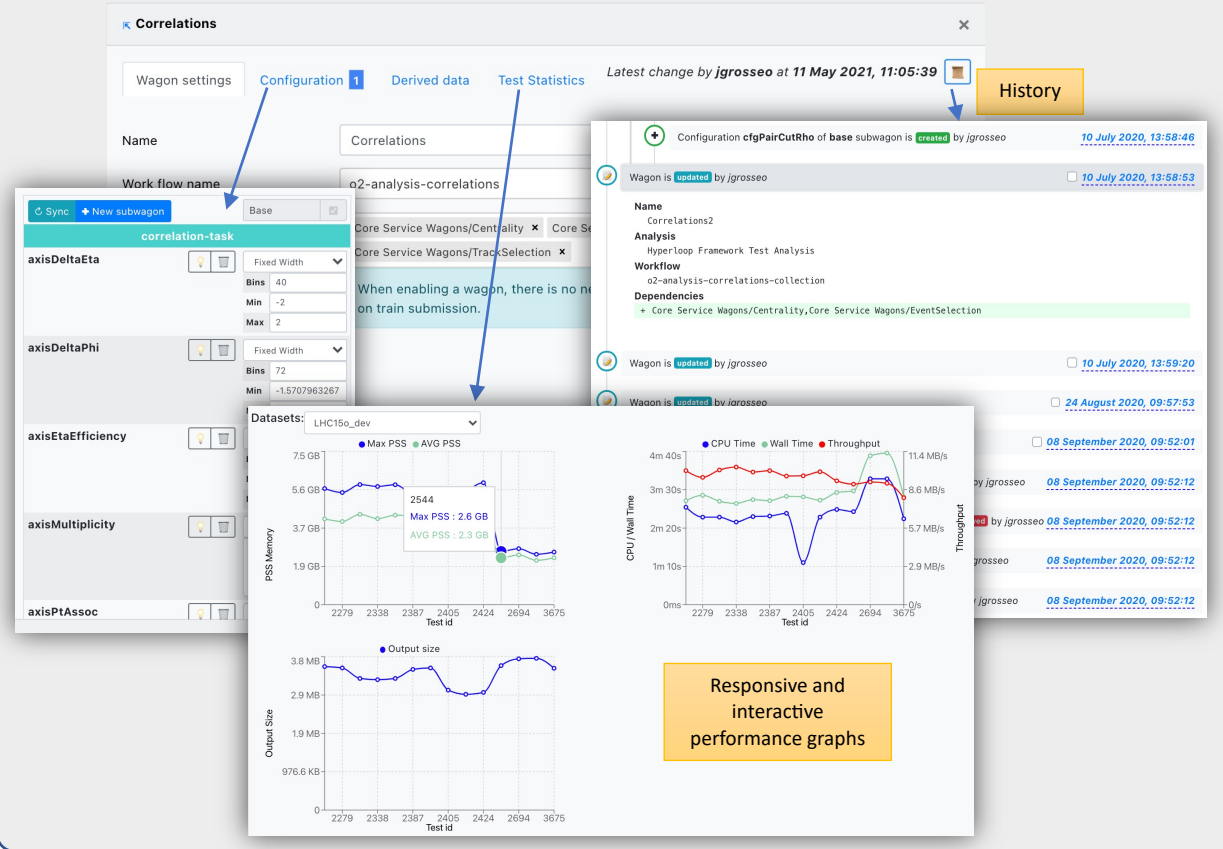

Hyperloop trains

- Most train operations are automatic based on available target memory and wagon configuration compatibility
- Requires less and **global train operators**

Dataset definition

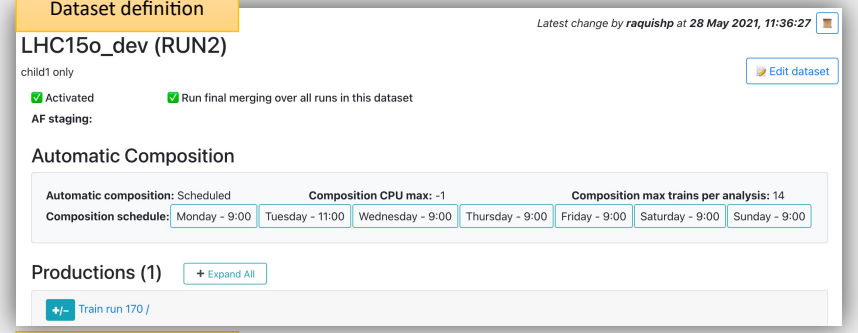

Train submission

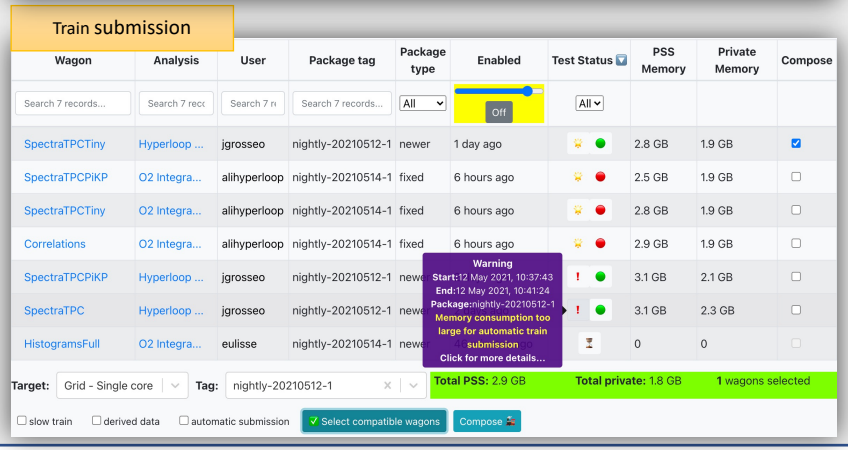

Train run result

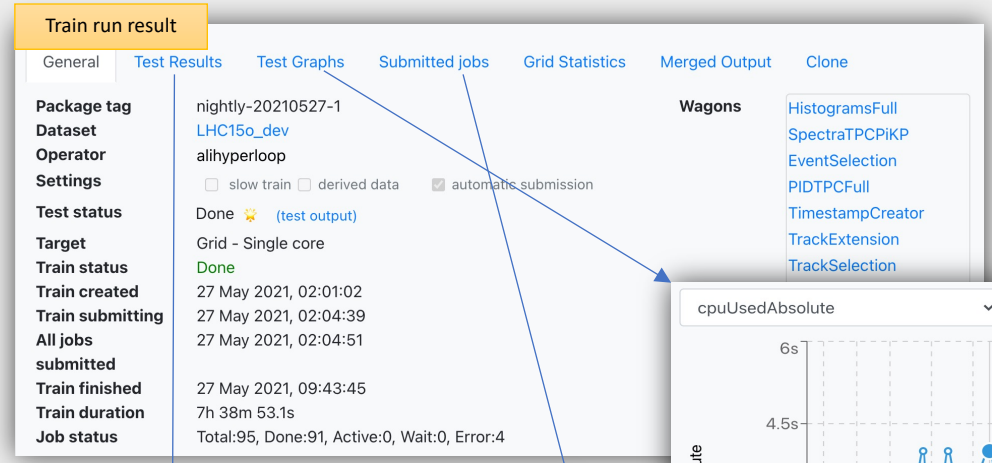

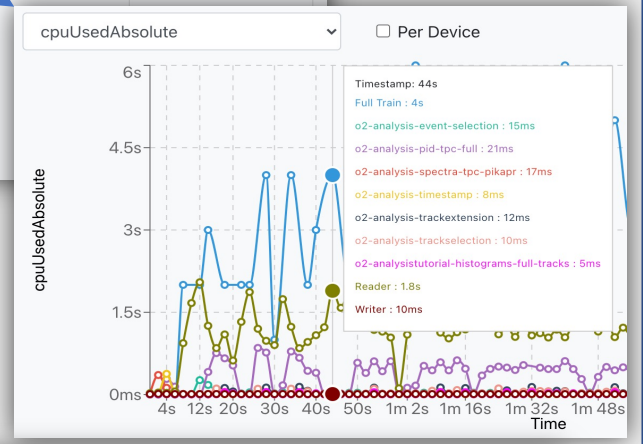

Wagon	PSS Memory		Private Memory	CPU Time	CCDB	
					Calls	Transfer size
Reader	Max	1.6 GB	1.1 GB	1m 51s		
	Avg	1.2 GB	794.8 MB			
	Slope	3.9 MB/s	1.8 MB/s			
EventSelection	Max	308.7 MB	120.5 MB	3s		
	Avg	150.4 MB	114.7 MB			
	Slope	-360.1 KB/s	179.5 KB/s			
HistogramsFull	Max	25.7 MB	9.6 MB	2s		
	Avg	19.6 MB	9.5 MB			
	Slope	87.7 KB/s	157.4 B/s			
PIDTPCFull	Max	226.3 MB	158.3 MB	38s		
	Avg	185.5 MB	120.2 MB			
	Slope	368.5 KB/s	11.7 KB/s			
SpectraTPCPiKP	Max	165.4 MB	90.5 MB	3s		
	Avg	142.9 MB	88.6 MB			
	Slope	357.1 KB/s	53.9 KB/s			
TimestampCreator	Max	182.4 MB	71.7 MB	1s		
	Avg	74.4 MB	71.7 MB			
	Slope	-36.6 KB/s	921.1 B/s			

PID	Run no.	Output directory	Total	Done	Active	Wait	Error	Merging
2272622805	246392	/alice/data/2015/LHC15o/000246392/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8456	20	18	0	0	2	Done
2272622803	246391	/alice/data/2015/LHC15o/000246391/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8455	16	15	0	0	1	Done
2272622802	245068	/alice/data/2015/LHC15o/000245068/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8454	5	5	0	0	0	Done
2272622800	245066	/alice/data/2015/LHC15o/000245066/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8453	6	6	0	0	0	Done
2272622797	245064	/alice/data/2015/LHC15o/000245064/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8452	20	19	0	0	1	Done
2272622795	244983	/alice/data/2015/LHC15o/000244983/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8451	5	5	0	0	0	Done
2272622793	244982	/alice/data/2015/LHC15o/000244982/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8450	8	8	0	0	0	Done
2272622790	244980	/alice/data/2015/LHC15o/000244980/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8449	5	5	0	0	0	Done
2272622786	244975	/alice/data/2015/LHC15o/000244975/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8448	4	4	0	0	0	Done
2272622784	244918	/alice/data/2015/LHC15o/000244918/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8447	5	5	0	0	0	Done
2272622782	244917	/alice/data/2015/LHC15o/000244917/pass5_lowIR/PWGZZ/Run3_Conversion/170_20210506-0955_child_1/hy_8446	1	1	0	0	0	Done

Summary

- Based on the organized analysis trains (LEGO) success, a modern and reactive train framework (Hyperloop) was built which supports the Run 3 software suite
- The Hyperloop train framework has been used during the **ALICE Analysis Challenge** since June 2020
- More than 600 trains successfully submitted to the Grid and Analysis Facilities
- Factor 3-10 higher event throughput, currently being optimized