Status of the LEGO framework

Costin Grigoras,
Jan Fiete Grosse-Oetringhaus,
Markus Zimmermann

19/03/2015

Status

- LEGO framework is the system for organized analysis in ALICE
- https://alimonitor.cern.ch/trains/
- key elements:
  - MonALISA
  - LEGO backend
  - Lightweight Production Manager (LPM)
  - AliEn
- twiki:
  https://twiki.cern.ch/twiki/bin/viewauth/ALICE/AnalysisTrains

```
active = run since the last Offline week
         presentation on the 20th of November
```

- 58 active trains (out of 89 in the system) (for all PWGs)
- 120 active users (out of 204 in the system)

Very well used system

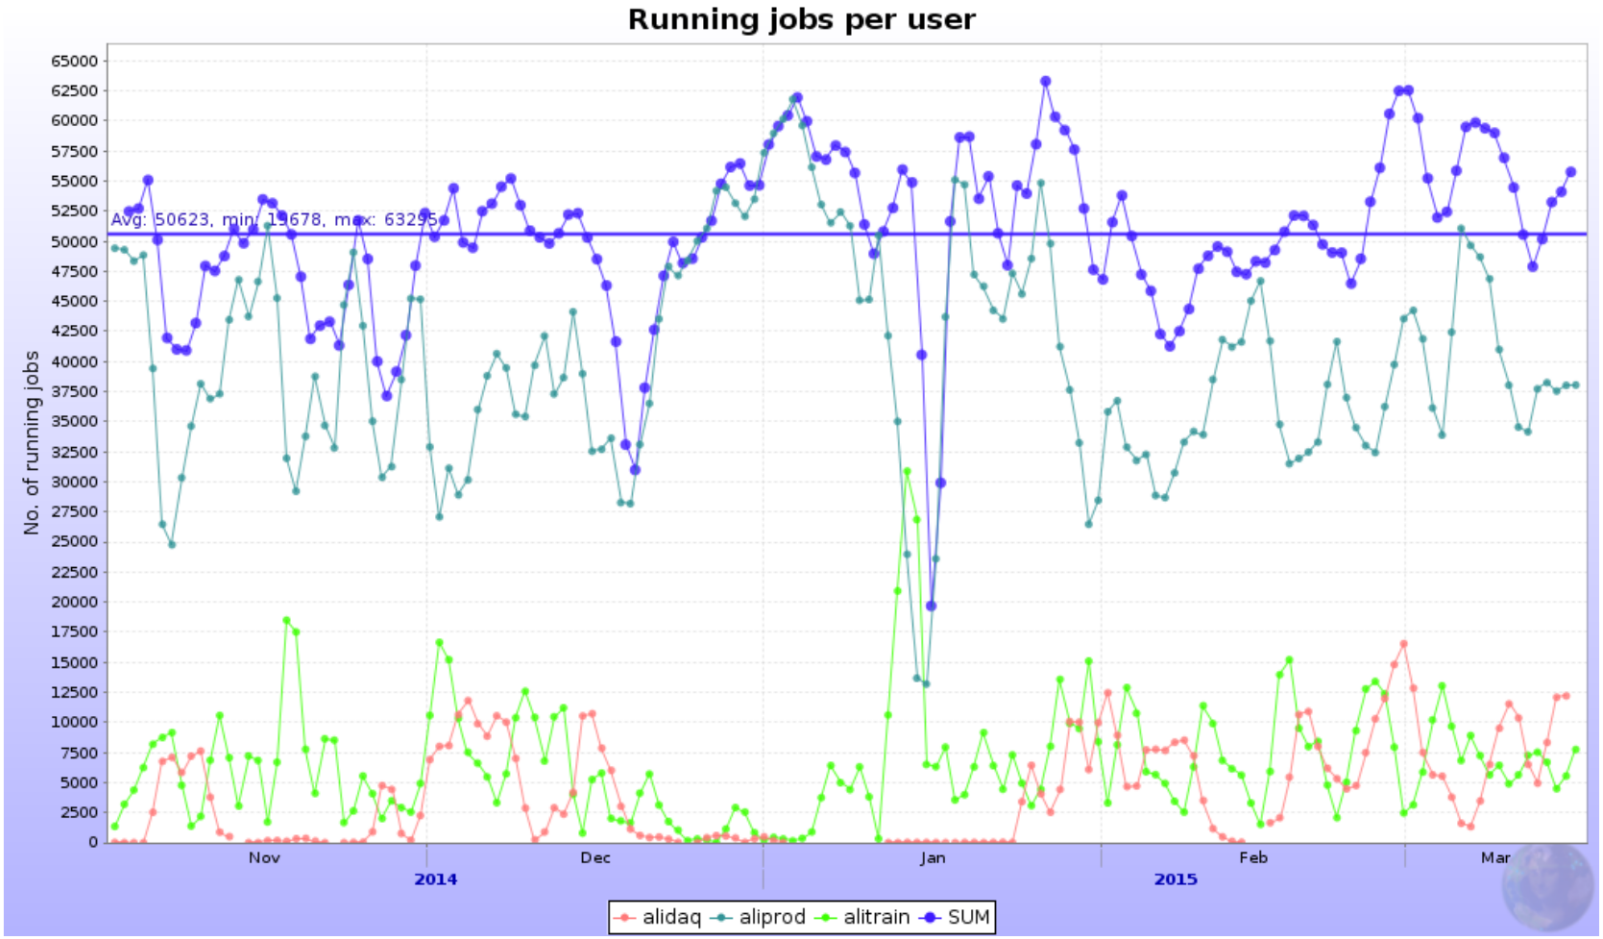

Figure 1: Performance of the GRID in the last 6 months.

user	avg number of jobs	relative to the sum
sum	54700 (up from 42600)	100%
alitrain	6600 (up from 5200)	12% (12)
users	3500 (up from 2700)	6% (6)

Per month statistics of the train system

	07.2014 - 10.2014	11.2014 - 02.2015
Total wall time	335y	398y
Total number of train runs	670	575
Wagons per train run	7.2	7.1
Number of Grid jobs	3.1 Million	3.0 Million
Train duration	13:07	13:39
Submission	1:49	1:13
Analysis & per Run merging	10:20	11:26
Final merging	0:58	0:59
ESD	50y (15%)	55y (14%)
AOD	225y (67%)	269y (68%)
MC ESD	39y (12%)	51y (13%)
MC AOD	19y (6%)	17y (4%)

Table 2: Per month normalized status values of the train system.

Presentation of the LEGO trains

- the proceedings of the ACAT conference are on arXiv: http://arxiv.org/abs/1502.06381

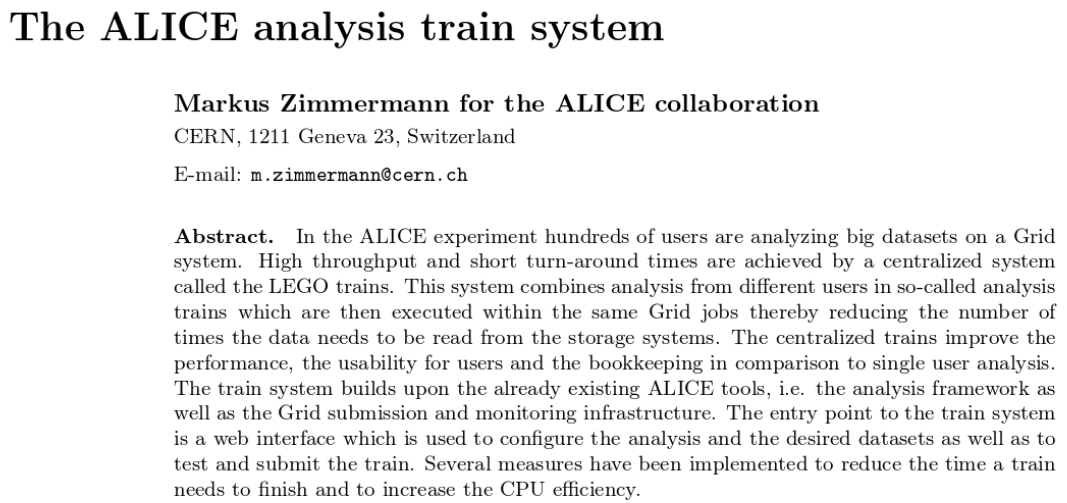

The ALICE analysis train system

Markus Zimmermann for the ALICE collaboration
CERN, 1211 Geneva 23, Switzerland
E-mail: m.zimmermann@cern.ch

Abstract. In the ALICE experiment hundreds of users are analyzing big datasets on a Grid system. High throughput and short turn-around times are achieved by a centralized system called the LEGO trains. This system combines analysis from different users in so-called analysis trains which are then executed within the same Grid jobs thereby reducing the number of times the data needs to be read from the storage systems. The centralized trains improve the performance, the usability for users and the bookkeeping in comparison to single user analysis. The train system builds upon the already existing ALICE tools, i.e. the analysis framework as well as the Grid submission and monitoring infrastructure. The entry point to the train system is a web interface which is used to configure the analysis and the desired datasets as well as to test and submit the train. Several measures have been implemented to reduce the time a train needs to finish and to increase the CPU efficiency.

Figure 2: Proceedings of the talk at the ACAT conference last September.

Subwagon feature

- half a year ago the subwagon feature was introduced
- one wagon can run with several configurations simultaneously
- 7 wagons are using this feature (out of 1936)
- increase acceptance of this feature with
  - a presentation at the next analysis tutorial
  - contacting individual users who could have use for this feature
- feature is very useful for systematic studies
  - of different particle species
  - of different track cuts
  - many more

Memory allocation with jemalloc

- memory allocation was improved by using jemalloc
- put into production on 18th of February
- compare rate of errors in the month before and afterwards

time jobs started	memory error	successful	rel(%)
18.01 - 18.02	29 939	3 052 384	0.98%
18.02 - 18.03	15 201	2 152 230	0.71%

Table 3: Number of memory errors in January and February dependent on the time the jobs were started.

- the memory error rate went down by 30%
- the overall rate of the second time period is smaller because not all jobs were finished yesterday

Updates

- adjusted the train system for the AliRoot splitting
- maintenance of the dataset file distribution
  - one 'clean up' was done,
    this means files were copied to improve the job splitting
- under development
  - provide files for local train test directly from the wagon page

Summary

- the train system is still well used by the users
- continuous support
- adjusted the train system for the AliRoot splitting
- no major updates have been put into production since the last offline week