



Status: LEGO Framework

Costin Grigoras
Jan Fiete Grosse-Oetringhaus
Markus Zimmermann

Offline Week
07.11.13

Status

- The LEGO framework is the system for organized analysis in ALICE
 - Key elements: MonALISA + LEGO backend + LPM + AliEn
Details: see e.g. <https://indico.cern.ch/getFile.py/access?contribId=2&resId=1&materialId=slides&confId=119073>
and <https://twiki.cern.ch/twiki/bin/viewauth/ALICE/AnalysisTrains>
- 57 active trains (from all PWGs)
- 142 users
- Per month (based on last 6 months):
 - 350 train runs
 - 2 million jobs
 - 178 years wall time
 - 9 wagons/train run on average
(spread depending on PWG: 3 – 22)
 - Average train duration: 21h
(spread depending on PWG: 18 – 24h)

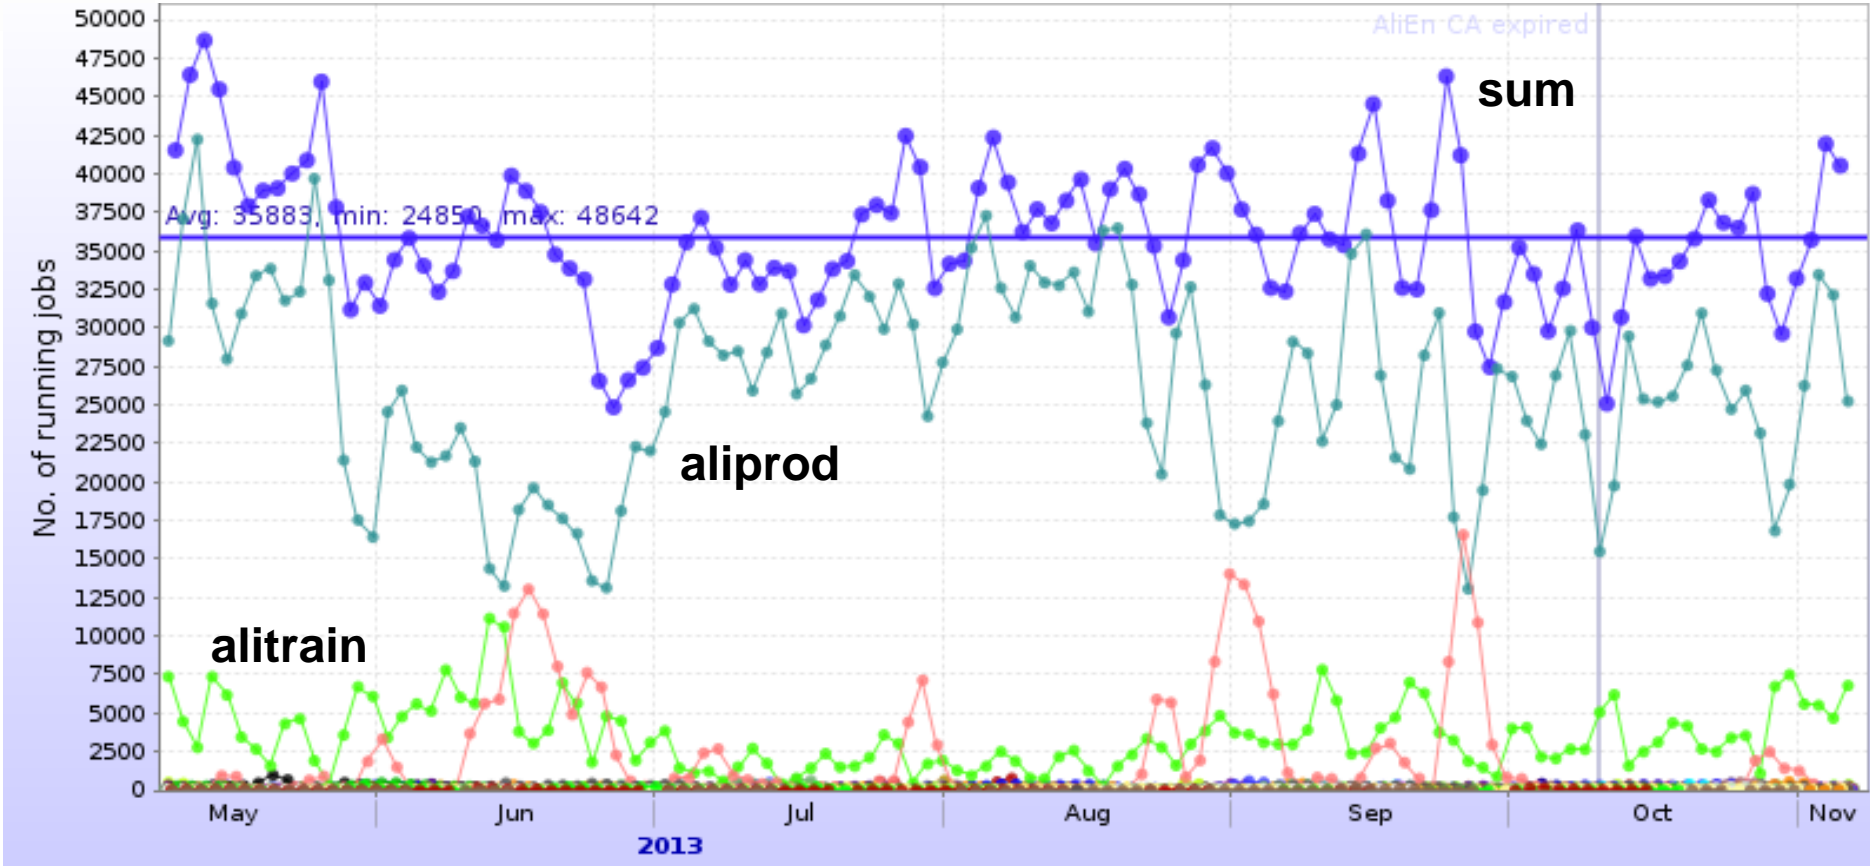
Very well used system



ALICE

Running Jobs

(last 6 months)



average number of jobs			
sum	36000	alitrain	3500
aliproduct, daq	28800	users	3700

Per PWG Statistics

(last 6 months)

PWG	Total Wall Time	Total Number of train runs	Wall per train	Total train duration	Submission	Analysis + per run merging	Final merging
CF	352y	696	184d	21:31	5:47	15:01	1:33
DQ	115y	105	1y 37d	19:58	2:35	15:49	1:34
GA	79y	323	90d	21:36	2:58	16:59	1:42
HF	268y	453	216d	1d 0:01	5:06	17:23	1:30
JE	189y	296	233d	18:39	3:09	14:13	1:20
LF	58y	218	97d	20:41	2:55	16:44	1:05
PP	5y	27	72d	1d 2:58	1:18	23:39	2:00
Sum	1070y	2124	183d	21:36	4:16	16:03	1:33

AOD vs ESD Usage

(last 6 months)

PWG	ESD	% ESD	AOD	% AOD	MC ESD	% MC ESD	MC AOD	% MC AOD
CF	14d	0	339y	96	215d	0	13y	4
DQ	20y	18	92y	80	19d	0	2y	2
GA	34y	43	1y 161d	2	33y 24d	41	11y	14
HF	76y	29	133y	50	34y	13	23y	9
JE	51y	27	86y	46	49y 17d	26	3y	1
LF	4y	8	33y	57	17y	30	3y	5
PP	5y	100	0	0	0	0	0	0
Sum	192y	18	687y	64	135y	13	55y	5

2/3 of wall time is AOD based analysis




Per Week Statistics

Analysis +
per run merging


Final merging

Waiting for
submission



Dataset Consolidation

- Moving of a fraction of replicas in a dataset to optimize job splitting
- Reduction of number of jobs by factor 2 – 5 on certain datasets
 - No difference on train turn around within our uncertainties



→ Details in Markus talk



Dataset Consolidation (2)

- What improvement can be expected if applied to all datasets?
- 37 datasets used in last 2 months (few times)
 - 3.8 million jobs
- 38% have sub-optimal splitting (on average per job less than half of the desired number of files)
 - 1.2 million jobs
 - Optimizing those: 1.2M \rightarrow 360k jobs
- Optimizing all: 3.8M \rightarrow 2.6M jobs

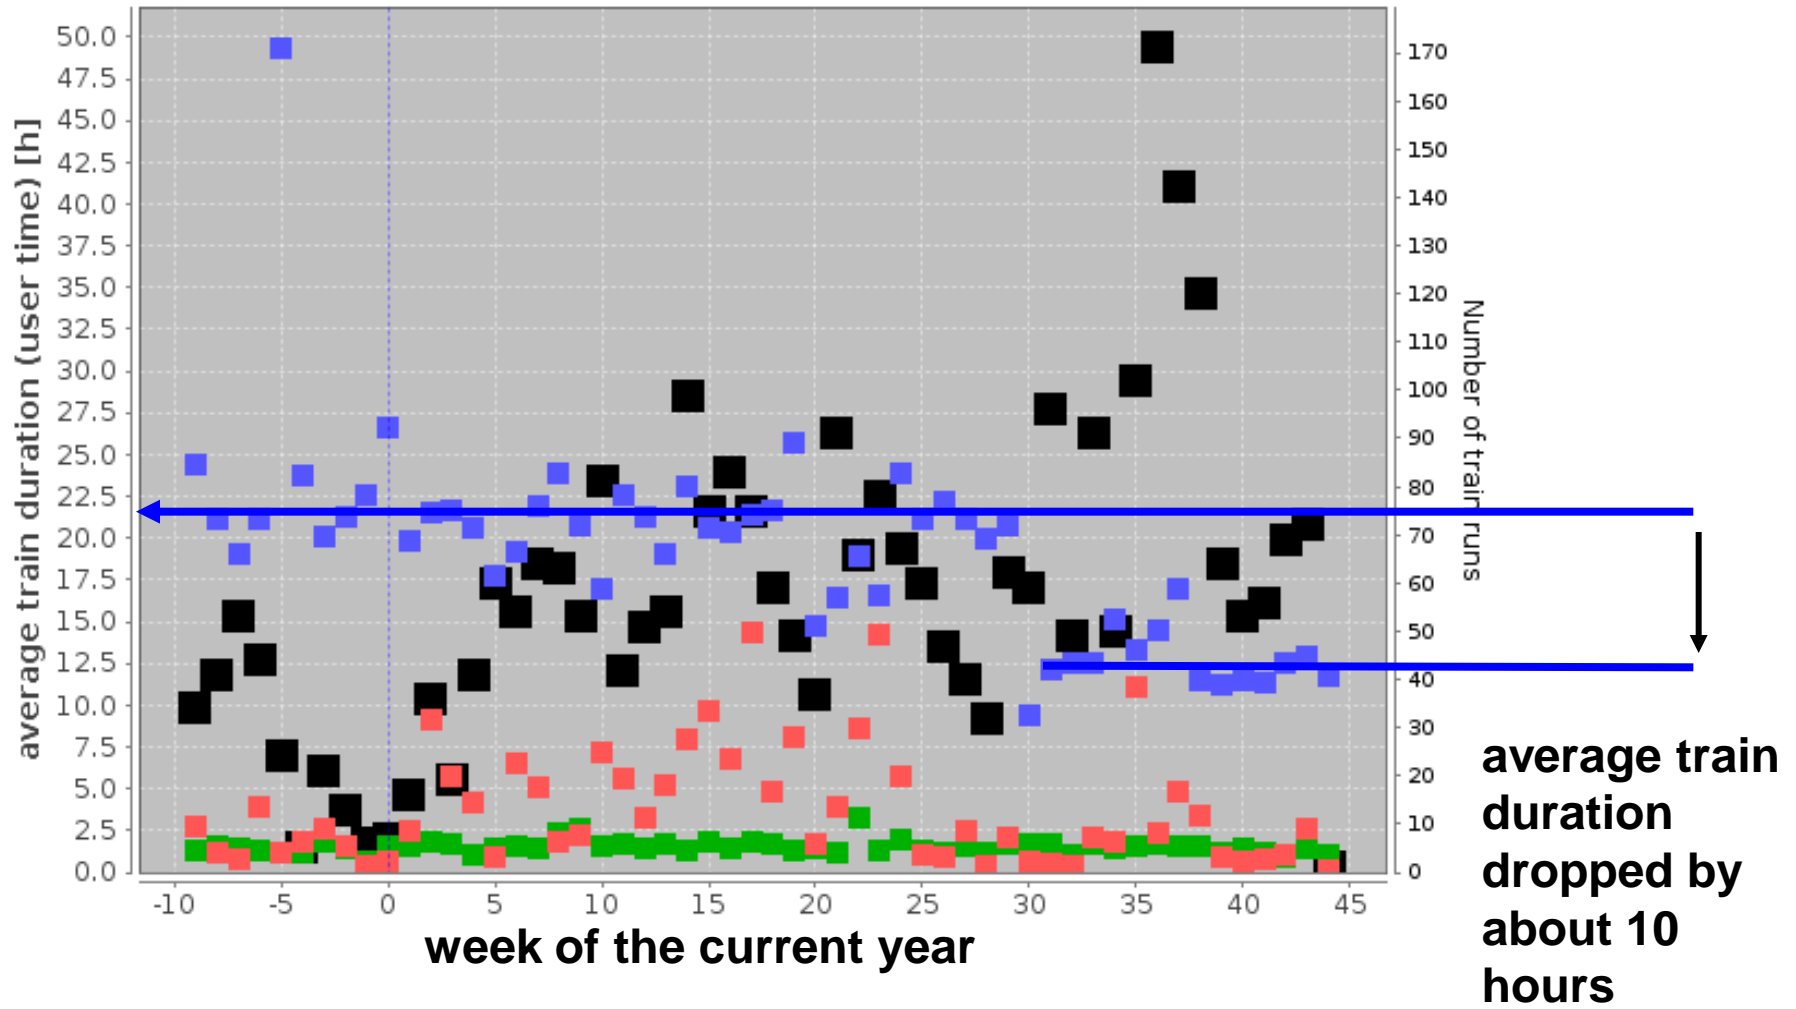


The Tail of the Distribution

- Usually a small fraction of long jobs delay the train completion significantly
- Dealing with the tail of the distribution
 - Drop SE requirements for last 10% of jobs of a given train
 - Kill last 2% of jobs after grace period
- Reduced turn around time significantly

→ Details in Markus talk

The Tail of the Distribution (2)



Plans

- Reduce turn-around time to 12 hours
- Interest in daily analysis tags
 - Requires split of AliRoot in core part and PWG part
 - Spreads out submission over the week (at the moment two peaks after the tag)
- Meta datasets containing more than one production
 - Avoids many operator actions for productions which for technical reasons have been split into several tags
 - Allows to group data together, e.g. LHC10b,c,d
- Support creation & analysis of microAODs



Backup



ALICE

- wg_no | wagons
- -----+-----
- GA | 5.3324808184143223
- CF | 4.2343750000000000
- PP | 2.3142857142857143
- DQ | 3.1521739130434783
- HF | 13.7364185110663984
- LF | 5.5887850467289720
- JE | 21.3303964757709251