

Advanced Analytics service to enhance workflow control at the ATLAS Production System

Mikhail Titov^{1,*}, Mikhail Borodin², Dmitry Golubkov^{1,3}, and Alexei Klimentov^{1,4}
on behalf of the ATLAS Collaboration**

¹National Research Centre “Kurchatov Institute”, Moscow, Russia

²University of Iowa, Iowa City, IA, USA

³Institute for High Energy Physics of NRC “Kurchatov Institute”, Protvino, Russia

⁴Brookhaven National Laboratory, Upton, NY, USA

Abstract. Modern workload management systems that are responsible for central data production and processing in High Energy and Nuclear Physics experiments have highly complicated architectures and require a specialized control service for resource and processing components balancing. Such a service represents a comprehensive set of analytical tools, management utilities and monitoring views aimed at providing a deep understanding of internal processes, and is considered as an extension for situational awareness analytic service. Its key points are analysis of task processing, e.g., selection and regulation of key task features that affect its processing the most; modeling of processed data life-cycles for further analysis, e.g., generate guidelines for particular stage of data processing; and forecasting processes with focus on data and tasks states as well as on the management system itself, e.g., to detect the source of any potential malfunction. The prototype of the advanced analytics service will be an essential part of the analytical service of the ATLAS Production System (ProdSys2). Advanced analytics service uses such tools as Time-To-Complete (TTC) estimation towards units of the processing (i.e., tasks and chains of tasks) to control the processing state and to be able to highlight abnormal operations and executions. Obtained metrics are used in decision making processes to regulate the system behaviour and resources consumption.

1 Introduction

The second generation of the Production System (ProdSys2) [1] of the ATLAS experiment (LHC, CERN) [2], in conjunction with the workload management system - Production and Distributed Analysis system (PanDA) [3], represents a complex set of computing components that are responsible for organizing, planning, starting and executing distributed computing tasks and jobs. Computing task represents a logical grouping of computing jobs and includes general description with corresponding requested parameters (e.g., type of production, campaign, software version, etc.). Computing job is assigned to the computing resource to be processed and to generate output data according to the defined program/transformation

*e-mail: mikhail.titov@cern.ch

**copyright 2018 CERN for the benefit of the ATLAS Collaboration. CC-BY-4.0 license.

over the input data. ProdSys2/PanDA is responsible for all stages of (re)processing, analysis and modeling of raw data obtained from the detector and corresponding derived data, as well as simulating of physical processes and functioning of the detector using Monte Carlo methods. Using the ProdSys2/PanDA software, the ATLAS scientific community, individual physics groups and scientists have access to hundreds of the Worldwide LHC Computing Grid (WLCG)¹ computing centers, supercomputers (HPC), cloud computing resources and volunteer computing resources (ATLAS@Home [4]).

ProdSys2 layers and components are presented in Figure 1. Each layer brings an additional set of parameters and metrics to control the processing workflow. Key sources of data for analytics processes are represented by the following core components: i) Database Engine for Tasks (DEFT), which is a main component of ProdSys2 that is responsible for forming computing tasks (task chains and group of tasks, i.e., production request) based on the set of parameters and processing conditions; and ii) Job Execution and Definition Interface (JEDI), which is a part of ProdSys2/PanDA and is responsible for managing the payload at the task level (i.e., brokerage and execution) and for the dynamic job definition and execution based on corresponding earlier defined tasks (optimization of the resources usage).




Figure 1. ProdSys2 layers and components

The enhancement of the task processing workflow requires a deep understanding of internal processes of ProdSys2. Thus a comprehensive analytical and management tools and utilities are designed to extend possibilities for intelligent task management, and to be considered as an extension for situational awareness analytics service [5].

2 Advanced analytics service

As one of the approaches for task analysis that considers the extraction of the operational metrics and its usage in decision making system, the advanced analytics service development was started. The goal of the service is to provide the analysis of the task processing lifecycle (e.g., selection and regulation of key task features that affect its processing the most), modeling of processed data lifecycles for further deep analysis (e.g., generation of guidelines for particular stage of data processing), and forecasting processes with emphasis on data and task states as well as on the management system itself (e.g., detection of the source of a particular potential malfunction).

The prototype of the analytics service for the ProdSys2 is designed to focus on predictive analytics approach that is aimed to increase the efficiency and the awareness of the operational processes. The design includes in itself the following key components:

- the set of independent tools that form the predictive model handling package which is responsible for predictive model creation and its usage for the process of TTC predictions generation. This package is adjusted and integrated into the prototype of the service.

¹Worldwide LHC Computing Grid, <http://wlcg.web.cern.ch/> [accessed on 2018-10-14]

- the web application which represents a central operational hub and is responsible for the next procedures:
 - monitor tasks execution process and their forecasted time to be completed (e.g., exec processes that are historic data of runs, evaluation of estimated durations of task executions);
 - control the parameters for the predictions generation (e.g., selection parameters for training and input data collections, method / technology and set of features for prediction process).

2.1 Infrastructure

Predictive model handling package is built to run within the parallel processing framework Apache Spark² and it uses Spark.MLlib [6] for forecasting process. Package’s components are presented in Figure 2 (these components and the proposed architecture of the service were originally introduced in ref. [5]): i) collector - extracts corresponding task parameters from DEFT/JEDI that affect its execution process the most (uses such tools as Apache Sqoop³ and Pig⁴ for extraction and normalization processes); ii) predictor - creates predictive model (the "training" process) and uses this model to generate predictions of new tasks durations based on machine learning methods (Gradient-Boosted Trees and Random Forests regression methods); iii) distributor - applies statistical analysis as post-processing before the delivery obtained results back to DEFT and analytical service itself by corresponding APIs. The full description of the Hadoop service components (including mentioned earlier Spark, Sqoop, and Pig) that are provided by CERN-IT is presented in ref. [7].

Infrastructure for web application includes: virtual machine that is based on CERN Open-Stack IaaS [8] with operating system CERN CentOS7 (x86_64); MySQL database that is provided by CERN service Database On Demand [9]; nginx - high-performance HTTP server and reverse proxy; gunicorn - python WSGI HTTP server for UNIX; django - high-level web application framework.

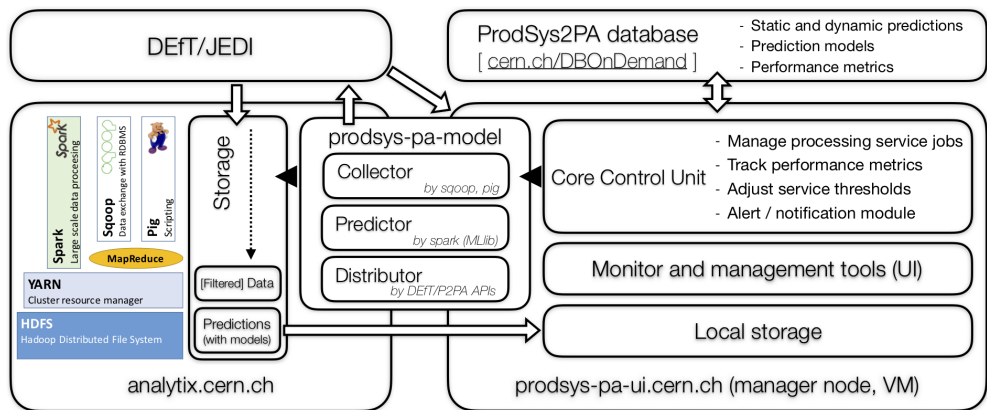


Figure 2. The architecture of the prototype of the ProdSys2 Predictive Analytics service (the Hadoop cluster “analytix” with highlighted service components as provided by CERN-IT)

²Apache Spark, <https://spark.apache.org> [accessed on 2018-10-14]

³Apache Sqoop, <https://sqoop.apache.org> [accessed on 2018-10-14]

⁴Apache Pig, <https://pig.apache.org> [accessed on 2018-10-14]

2.2 Analysis processes

The current implementation of the service prototype provides the following analysis processes [5], which are responsible for estimations and calculations of corresponding analytical metrics.

Threshold definition

The denoted process uses statistical analysis, it calculates the upper limit of the duration of tasks execution process in such a way that 95% of all tasks of the corresponding type and for the defined time period (the last 180 days) are executed not longer than the calculated value. Tasks are grouped into a certain type according to the set of parameters: *project* (under which processed data is derived), *productionStep* (the stage within the data processing and derivation chain), *workingGroup* (physics group that initiated the corresponding task).

Cold-prediction generation

This process represents a predictive modeling, it estimates task duration (i.e., prediction generation) during the task formation process. There is also a prior step - creation of predictive model, which is used for prediction generation process. It is based on task definition parameters that categorize the average execution process for the defined task type (with particular conditions). It represents an initial prediction and uses descriptive data. As it was mentioned earlier, the current implementation is built using Spark.MLlib Random Forests regression method.

2.3 Experiments

One of the performance evaluation of the presented analysis processes is the visualization of the comparison between real task durations and estimated ones (with corresponding quantitative metrics) that is presented in Figure 3 for ATLAS tasks which were created in period of from 12 to 26 August, 2018, and the total number of tasks is 16,880.

Figure 3 is characterized by the following metrics of the absolute error distribution between real and “cold”-predicted tasks durations: mean value of the absolute error is 0.5 days, standard deviation of the absolute error is 4.4 days, min/max values of the absolute error are -20.5/38.5 days respectively, and $3\sigma = 13.2$ days. Also, the estimation of the prediction quality is characterized by RMSE (root-mean-square error, measure of the differences between values predicted by a model and the values observed) = 4.4. The current estimated durations are considered as rough estimations that give an overview of the potential task execution duration, its further improvement (e.g., extension of the parameters set for the analysis, including dynamic parameters) will increase the accuracy and the quality of the analytics service.

3 Conclusions

The designed prototype of the service is aimed to enhance workflow control at the ATLAS Production System and to be able to detect and highlight abnormal operations and executions. Obtained metrics are used in decision making processes to regulate the system behaviour and resources consumption. Techniques and methods of predictive analytics would benefit the monitoring and control processes. The advanced analytics service (based on predictive analytics techniques) would also optimize the whole management process.

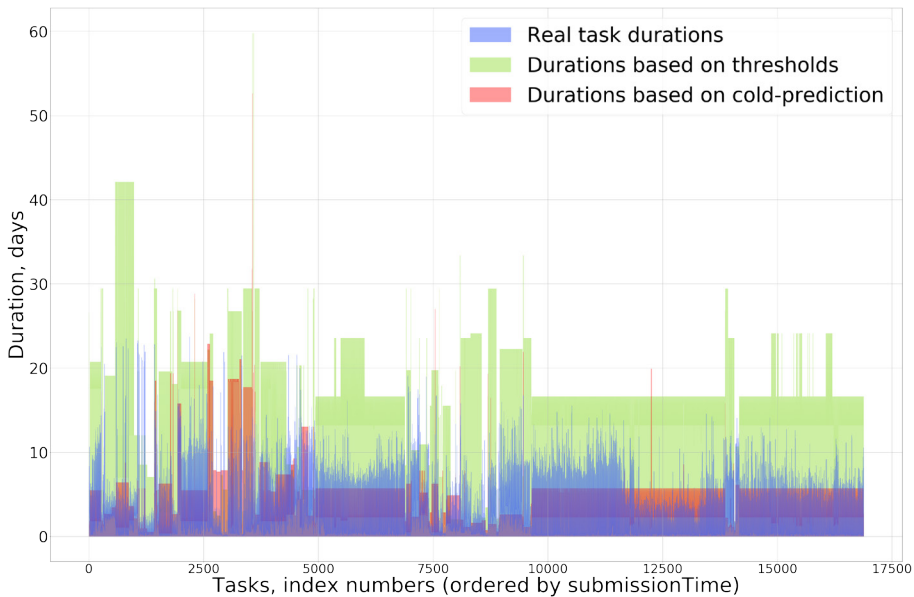


Figure 3. Real and estimated durations of ATLAS tasks execution (group of tasks for 2 weeks, 16,880 tasks): blue color - real duration for every task in the range, green color - estimated duration based on threshold definition per task type for every task in the range, red color - estimated duration based on prediction generation of “cold” type for every task in the range

This work has been carried out using computing resources of the federal collective usage center Complex for Simulation and Data Processing for Mega-science Facilities at NRC “Kurchatov Institute”, <http://ckp.nrcki.ru/>. NRC KI researchers have been funded by the Russian Ministry of Science and High Level Education under the contract No. 14.Z50.31.0024.

References

- [1] F.H. Barreiro *et al.*, J. Phys.: Conf. Ser. **898**, 052016 (2017)
- [2] ATLAS Collaboration, JINST **3**, S08003 (2008)
- [3] F.H. Barreiro *et al.*, J. Phys.: Conf. Ser. **898**, 052002 (2017)
- [4] C. Adam-Bourdarios *et al.*, J. Phys.: Conf. Ser. **664**, 022009 (2015)
- [5] M. Titov *et al.*, CEUR Workshop Proc. **2023**, 61-67 (2017)
- [6] X. Meng *et al.*, J. Mach. Learn. Res. **17**, 1235-1241 (2016)
- [7] D. Duellmann *et al.*, J. Phys.: Conf. Ser. **898**, 072034 (2017)
- [8] T. Bell *et al.*, J. Phys.: Conf. Ser. **664**, 022003 (2015)
- [9] R.G. Aparicio, I.C. Coz, J. Phys.: Conf. Ser. **664**, 042021 (2015)