

The ATLAS Production System Predictive Analytics service: an approach for intelligent task analysis

**M.A. Titov^{1,a}, M.S. Borodin², D.V. Golubkov^{1,3}, A.A. Klimentov^{1,4}
on behalf of the ATLAS Collaboration***

¹ *National Research Centre «Kurchatov Institute», 1 pl. Akademika Kurchatova, Moscow, 123182, Russia*

² *University of Iowa, 108 Calvin Hall, Iowa City, IA, 52242, USA*

³ *Institute for High Energy Physics of NRC «Kurchatov Institute», 1 pl. Nauki, Protvino, Moscow region, 142281, Russia*

⁴ *Brookhaven National Laboratory, P.O. Box 5000, Upton, NY, 11973, USA*

E-mail: ^a mikhail.titov@cern.ch

The second generation of the Production System (ProdSys2) of the ATLAS experiment (LHC, CERN), in conjunction with the workload management system PanDA (Production and Distributed Analysis), represents a complex set of computing components that are responsible for defining, organizing, scheduling, starting and executing payloads in a distributed computing infrastructure. ProdSys2/PanDA are responsible for all stages of (re)processing, analysis and modeling of raw and derived data, as well as simulation of physical processes and functioning of the detector using Monte Carlo methods. The prototype of the ProdSys2 Predictive Analytics (P2PA) service is an essential part of the growing analytical service for the ProdSys2 and it will play a key role in the ATLAS distributed computing. P2PA uses such tools as Time-To-Complete (TTC) estimation towards units of the processing (i.e., tasks, chains and groups of tasks) to control the processing state and rate, and to be able to highlight abnormal operations and executions (e.g., to discover stalled processes). It uses methods and techniques of machine learning to obtain corresponding predictive models and metrics that are aimed to characterize the current system's state and its changes over a short period of time.

Keywords: predictive analytics, production system, Apache Spark

© 2018 Mikhail A. Titov, Mikhail S. Borodin, Dmitry V. Golubkov, Alexei A. Klimentov

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



1. Introduction

Evolution of the Production System (ProdSys2) [1] of the ATLAS experiment [2] leads to extension of its possibilities by using not just technical and engineering solutions but techniques and methods of intelligent analysis based on data mining and machine learning. Such analysis is applied towards management and execution of computing tasks, as well as towards operational management processes. New components and services are designed to enhance the task processing workflow and to increase the automation in decision making processes [3,4].

The current key components of ProdSys2, such as the Database Engine for Tasks (DEfT) and the Job Execution and Definition Interface (JEDI), are used as main sources for information about computing tasks (set of parameters per task or chain of tasks) and their processing states. A computing task, in terms of ATLAS, represents a logical grouping of computing jobs that are responsible for the execution of algorithm/transformation on input files and generation of output files (dynamic jobs definition and execution are performed by JEDI). Profound understanding of a task lifecycle will improve its processing workflow and optimize usage of computing resources.

2. Problem statement

The new and advanced analytical service, that is aimed to collect and process information about tasks for their deep analysis, and to provide operational metrics for ProdSys2, is based on predictive modeling and analysis, and is called the Predictive Analytics service. The ultimate goal of this service is to solve the following problems/questions: i) discover and handle key task features that impact the workflow; ii) regulate the task processing/execution at a given stage; iii) predict task metrics and its next state (e.g., normal execution, stalled, etc.).

The next step in automation of the task processing management raises questions that are expected to be solved by the decision-making system, which will be a part of the Predictive Analytics service and will use service core tools for deep analysis of computing tasks. All that includes estimation of the correlation between task parameters and descriptive parameters of computing resources (e.g., selection and reservation of available computing capacities, determination of resources of the particular type for clustered group of tasks), mining of sequences of task reassignments (e.g., keep full track of task lifecycle stages and states, and task progress).

3. ProdSys2 Predictive Analytics service

The current implementation of the service includes two packages which represent key components (figure 1) [4].

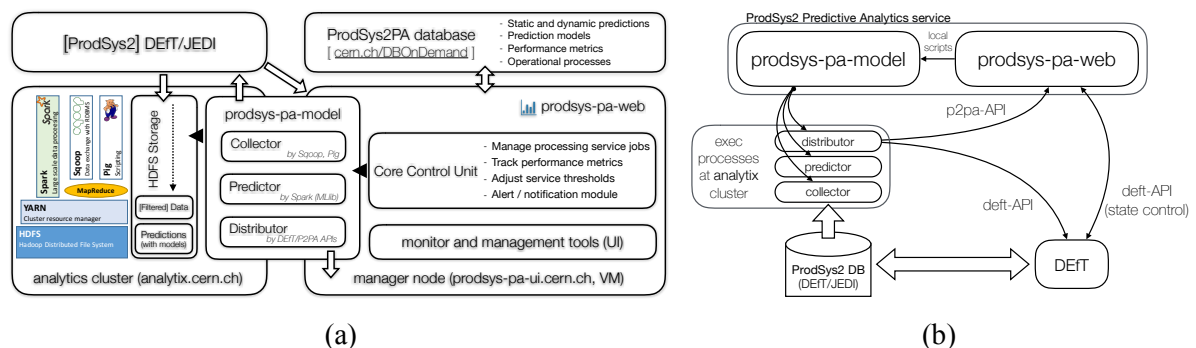


Figure 1. The architecture (a) and the communication (b) schemes of the P2PA service (analytics cluster “analytix” with highlighted services as provided by CERN-IT [5]).

Predictive model handling package (*prodsys-pa-model*) [6] is designed as an independent set of tools for task analysis: task information collector (extracts requested task parameters from DEFT/JEDI); analysis of the task operational parameters (creates predictive model and uses it to generate predictions of time-to-complete/TTC per new task); delivery of the obtained results (uses DEFT and P2PA APIs for predictions distribution). This package runs on an analytics cluster that provides HDFS and the parallel processing framework Apache Spark¹ (e.g., cluster analytix at the CERN Computing Center), and it is adapted to be a part of the service. Web application package (*prodsys-pa-web*) consolidates in itself monitor and management tools, and provides an interface to interact with task analysis process. It is built using the Django² web framework and related service python libraries (django REST framework³, celery⁴).

P2PA service also collects certain task timing parameters to evaluate applied prediction methods and chosen set of parameters (i.e., the quality of feature selection process). The current implementation of the prediction generation processes uses Random Forests regression method by Spark.MLlib, but it is in long-term plan to add other libraries with new methods. Evaluation of prediction models uses the following metrics as basic: mean squared error (MSE) and root mean squared error (RMSE). The quality/accuracy of generated predictions is characterized by the corresponding metric, i.e., a confidence coefficient. Its evaluation uses the tracking of the task execution progress ("state control" process in figure 1b).

The full set of generated data along with the evaluation metrics (including operational metrics to estimate the performance of applied methods) are presented to the user as the monitoring part of the service. Figures 2,3 show screenshots of web application that provides information about operation processes (e.g., predictive model creation, predictions generation) and task profiles with extracted parameters and estimated metrics, such as TTC predicted, description of the block of generated predictions with corresponding confidence coefficient and MSE.

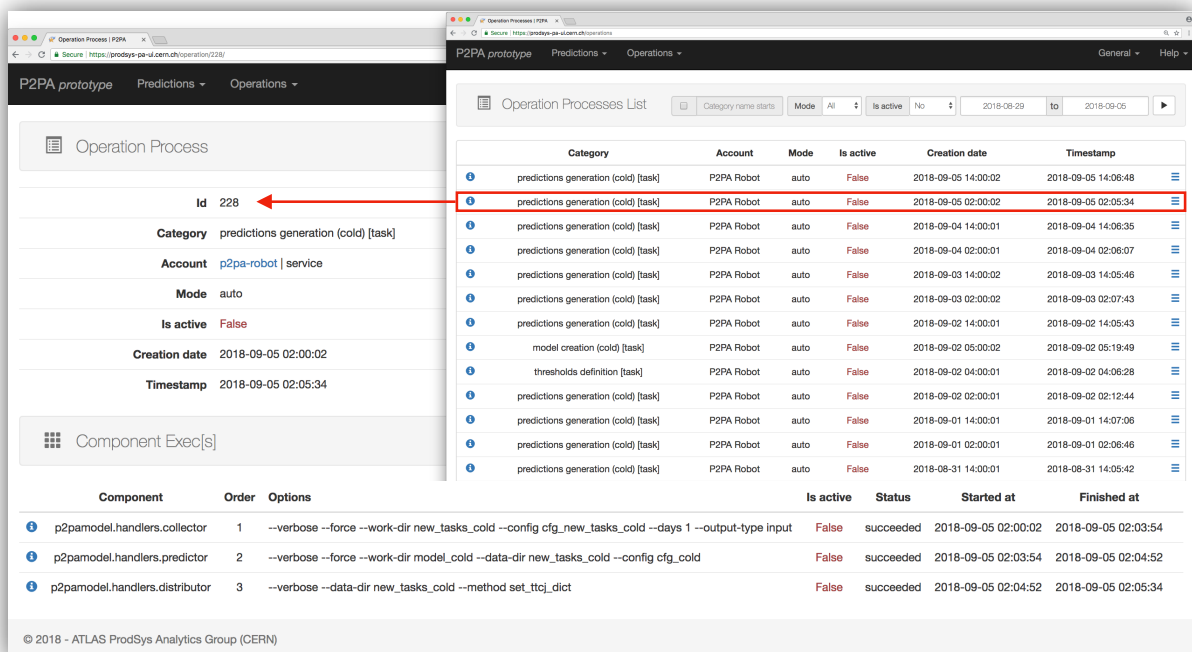


Figure 2. P2PA web application (UI) screenshots for operation process(es).

¹ Apache Spark, <https://spark.apache.org> [accessed on 2018-10-25]

² Django project (version 1.11), <https://www.djangoproject.com/> [accessed on 2018-10-25]

³ Django REST framework, <https://www.django-rest-framework.org/> [accessed on 2018-10-25]

⁴ Celery: Distributed Task Queue, <http://www.celeryproject.org/> [accessed on 2018-10-25]

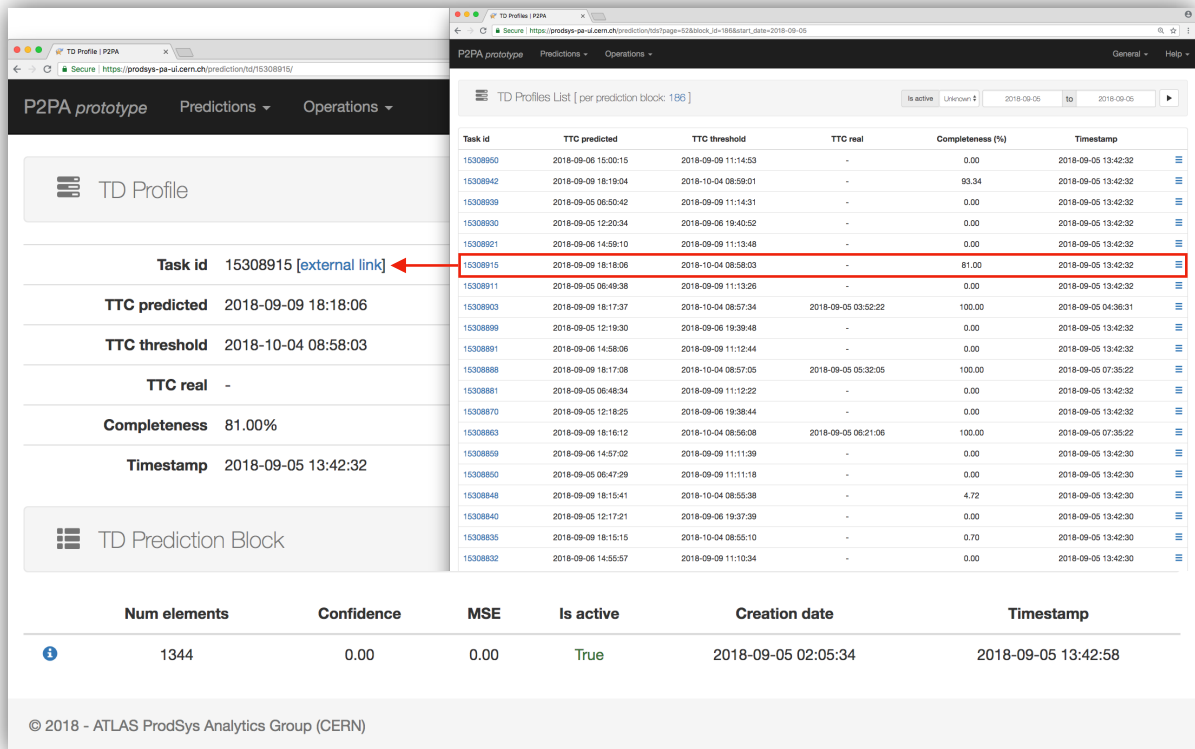


Figure 3. P2PA web application (UI) screenshots for task profile(s) with generated TTC estimation(s).

4. Analysis of a computing task

Generated predictions and obtained metrics are planned to be used in decision-making processes to regulate ProdSys2 behaviour and resources consumption. Thus, it is important to identify essential system influenced features that reflect the system behaviour. The key metric per computing task in ProdSys2 is TTC, which is used as an indicator for task condition (e.g., faster than average, longer than average, etc.) and its further exploration will let revealing reasons for its processing deviation that is important in forecasting the state of ProdSys2 in general. It is not yet planned to use it as a pre-task-definition check for optimal parameters finding (e.g., computing center), but such possibility will be considered with the service improvements.

There are several steps in estimation of task TTC, where each of them enhances the quality of obtained results from the previous ones.

Steps for task TTC estimation:

- definition of the value range - 95th percentile of task duration is used per group of tasks that are distinguished by a set of features. Current implementation uses the following features: *projectName*, *productionStep*, *workingGroup*;
- prediction of task duration based on descriptive / initial parameters of the task;
- periodically repeated step that uses dynamic parameters (of the task and computing environment including computing site for processing) for the adjustment of earlier predicted TTC and task duration eventually.

The current choice of controlled parameters/metrics is due to their correlation with corresponding ProdSys2 possible failure states [3], but with a further increase of their quality/accuracy and the need to introduce new ones, their comparative analysis will be conducted.

5. Acknowledgement

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.

6. Conclusion

ProdSys2 Predictive Analytics service is designed to enhance workflow control at the ATLAS Production System and to be able to detect and highlight abnormal operations and executions. Its prototype demonstrates the usefulness of the provided metrics and state control mechanism. It still lacks most evaluation metrics to fine-tune the prediction process that would increase the quality of generated predictions and operational metrics.

Furthermore, the future decision-making system should rely on generated quality metrics, since it is responsible for regulation of the resource consumption. The quality of obtained metrics (estimated values of controlled parameters) is constantly improving and new evaluation parameters and metrics will be introduced for task analysis and mining processes.

References

- [1] Borodin M. et al. The ATLAS Production System Evolution: New Data Processing and Analysis Paradigm for the LHC Run2 and High-Luminosity // *Journal of Physics: Conference Series* 898 (2017) 052016
- [2] ATLAS Collaboration, 2008 JINST 3 S08003
- [3] Titov M. et al. Predictive analytics as an essential mechanism for situational awareness at the ATLAS Production System // *CEUR Workshop Proceedings 2023* (2017) pp.61--67
- [4] Titov M. et al. Advanced Analytics service to enhance workflow control at the ATLAS Production System // *Proceedings of the 23rd International Conference on Computing in High Energy and Nuclear Physics (CHEP)*, Sofia, Bulgaria, 9-13 July 2018
- [5] Duellmann D. et al. Hadoop and friends - first experience at CERN with a new platform for high throughput analysis steps // *Journal of Physics: Conference Series* 898 (2017) 072034
- [6] Predictive model handling package, <https://github.com/XDatum/prodsys-pa-model> [accessed on 2018-10-25]