SEARCH FOR ANOMALIES IN THE COMPUTING JOBS EXECUTION OF THE ATLAS EXPERIMENT WITH THE USE OF VISUAL ANALYTICS

M. Grigorieva[1], T. Galkin[2], T. Korchuganova[3], I. Milman[2], V. Piluygin[2], M. Titov[1]

[1]NRC “Kurchatov Institute”
[2]NRNU “MEPHI”
[3]NR TPU

Outline

- Computing challenges in the ATLAS experiment
- Problem statement
- ATLAS data sources
- Proposed approach of visual analytics
- Analysis of ATLAS jobs execution
- Interpretation of the results
- Conclusions and future work

Computing Challenges in ATLAS

From ATLAS Dashboard http://dashb-atlas-job.cern.ch/

Statistics for time period: 2017-01-01 – 2018-08-01

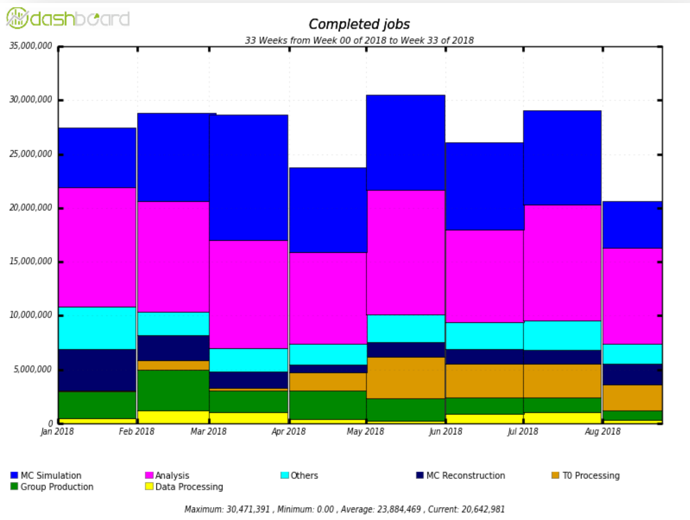

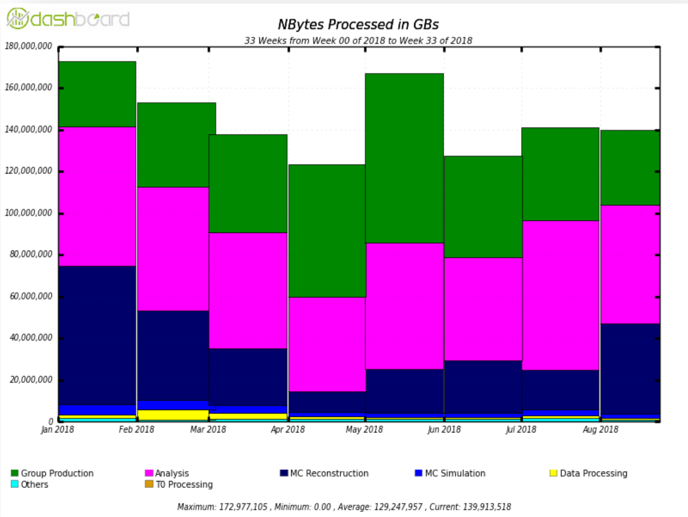

ATLAS STATS

Computing Resources:

- WLCG
- Opportunistic resources
  - HPC, Academic clouds, University clusters, Volunteer computers
- Total for more than decade:
  - 10 millions of tasks
  - 3 billions of jobs

- **TASK** is an activity that needs to be accomplished within a defined period of time. It contains execution code and input/output files, corresponding to underlying physics process and initial conditions.
- Each task is fragmented in **JOBS** which correspond to a fixed number of events.

Problem Statement

- **Large-scale distributed system in ATLAS faces the following challenges:**
  - Big diversity and complexity
  - Highly dynamic computing environments
  - Ongoing competition between different threads of computing jobs
  - Complex workflow of jobs execution
  - Uncountable possible reasons of failures and unstable behavior
- **A fundamental goal** is to increase the stability and efficiency of the distributed data processing and analysis systems.
  - **Current task:** The analysis of jobs execution process
- **What should be done?**
  - Development of analysis algorithms and tools using ML and statistics methods for detection of disruption of the operational process of workload management systems in ATLAS
  - Development of interactive methods and tools of visual analytics, providing the use of dynamic and static spatial interpretations of the analyzed data.

Visual analytics of multidimensional data

Visual Analytics combines data visualization with ML and other automated techniques to create systems that help people make sense of data.

Typically, domain experts have limited involvement in the process of data analysis.
Traditional machine-learning workflow: practitioners collect data, select features, preprocess and transform the data, choose a representation and learning algorithm to construct the model, tune parameters of the algorithm, and finally assess the quality of the resulting model.
This assessment often leads to further iterations on many of the previous steps.
Domain experts involvement in this process is mediated by the practitioners and is limited to providing data, answering domain-related questions, or giving feedback about the learned model.
This results in a design process with lengthy and asynchronous iterations and limits the end users' ability to affect the resulting models [1].

By integrating ML algorithms with interactive visualization, visual analytics aims at providing visual platforms for analysts to interact directly with data and models.

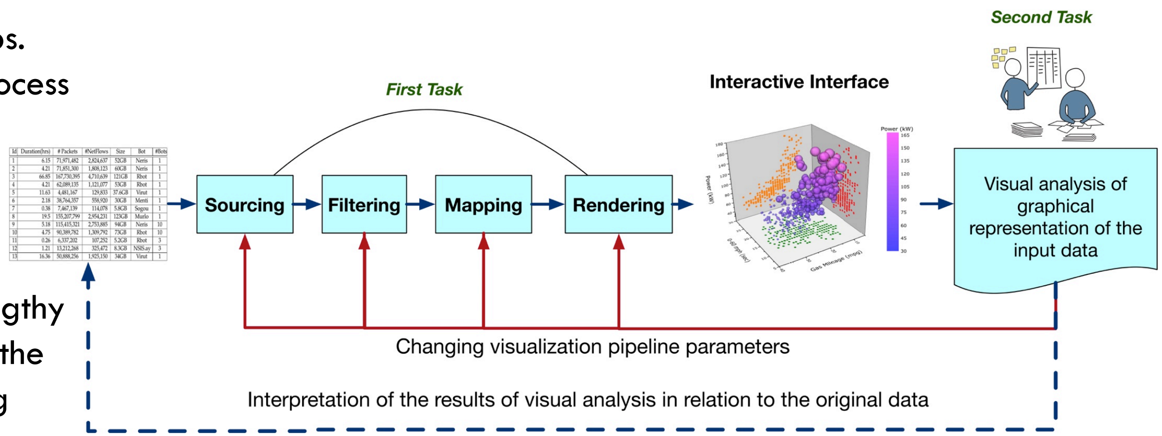

[1] D. Sacha. What you see is what you can change: Human-centered machine learning by interactive visualization // Neurocomputing 268 (2017) 164-175

ATLAS Data Sources and Job Execution Metrics

- ATLAS Data Sources
  - Rucio (Distributed Data Management System)
  - NWS (Network Weather Service)
  - AGIS (ATLAS Grid Information system)
  - DEFT (Database Engine for Tasks)
  - JEDI (Job Execution and Definition Interface)
  - PanDA MemoryMonitor

ATLAS-Kibana
(ElasticSearch)

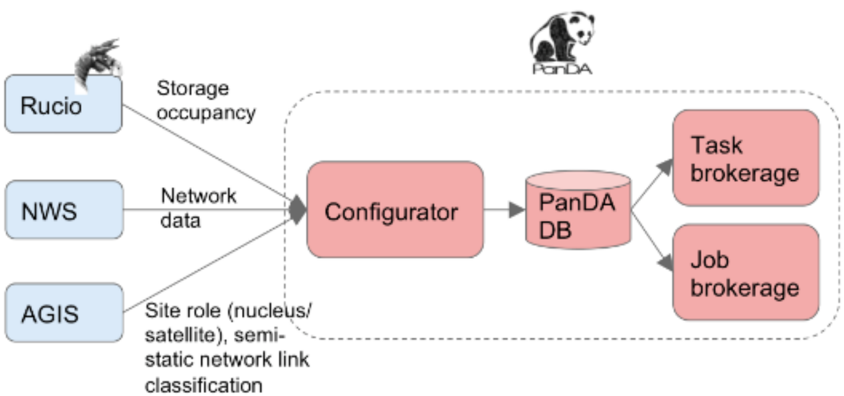

- Job Execution Metrics
  - Application-level
    - Job description, status, input/output files, start/end time
  - Middleware-level
    - Data transfer service, scheduler(queues), computing element, storage management
  - Network-level
    - Network connection status, data transfer rate between sites
  - Resource-level
    - CPU utilization, memory/storage usage, packet I/O rates

The proposed method of multidimensional data visual representation

- *Jobs = multidimensional points, parameters values = coordinates*
- Euclidean or Mahalanobis distances is calculated between all pairs of points
- *Points* are projected to a 3-dimensional space and drawn as spheres
- If distance between points is less than the threshold, then a cylinder is constructed to connect the spheres
  - Threshold can be changed interactively, allowing to observe the changes in the cluster structure
- The color of the cylinder simulates the distance between the points from red (small distance) to blue (close to the threshold):

$$p_i = (p_i^1, p_i^2, \dots, p_i^n) \xrightarrow[SphC]{p_i^{l_1}, p_i^{l_2}, p_i^{l_3}, SphR}$$

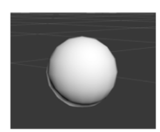

$$p_{i_1}, p_{i_2}, d\colon \rho(p_{i_1}, p_{i_2}) \le d \xrightarrow[CylC]{CylR}$$

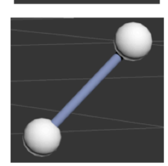

Representation of computing jobs execution (multidimensional tabular data)

	Parameter 1	Parameter 2	...	Parameter n
Job 1	x_{11}^j	x_{12}^j	...	x_{1n}^j
...	...	...	...	
Job i	x_{i1}^j	x_{i2}^j	...	x_{in}^j
...	...	...	...	...
Job m	x_{m1}^j	x_{m2}^j	...	x_{mn}^j

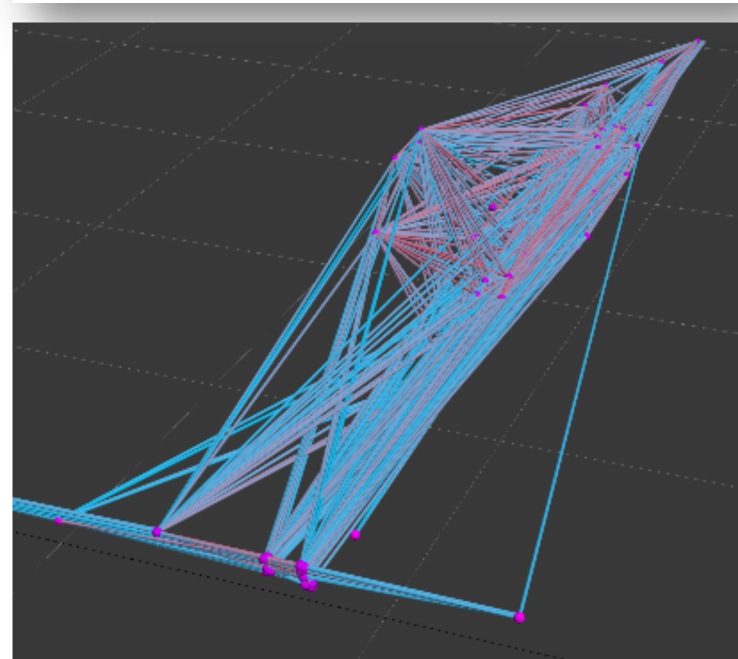

Multidimensional visual analysis software prototype

- First prototype is created based on **Autodesk 3ds Max**
- A combination of MAXScript scripts and C# module was used
- Depending on the amount of RAM, the software can handle up to a couple of hundreds of objects
- Points (Spheres) are coded with colors
- Results can be exported to excel (xlsx) files

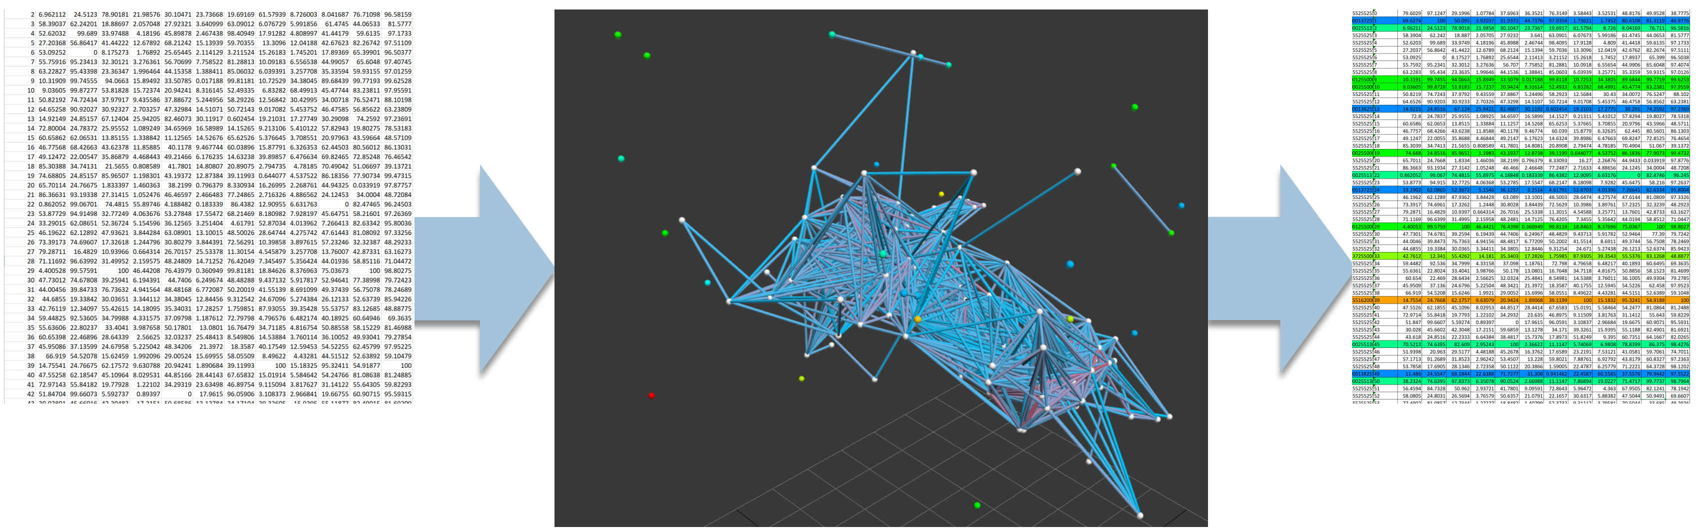

General Case Hypothesis

- **<u>Task №14138001</u>** (may-june 2018) ~50 000 jobs, 108 computing sites

Distributions of jobs' execution time and CPU time on computing sites should match

TimeExe Distributions

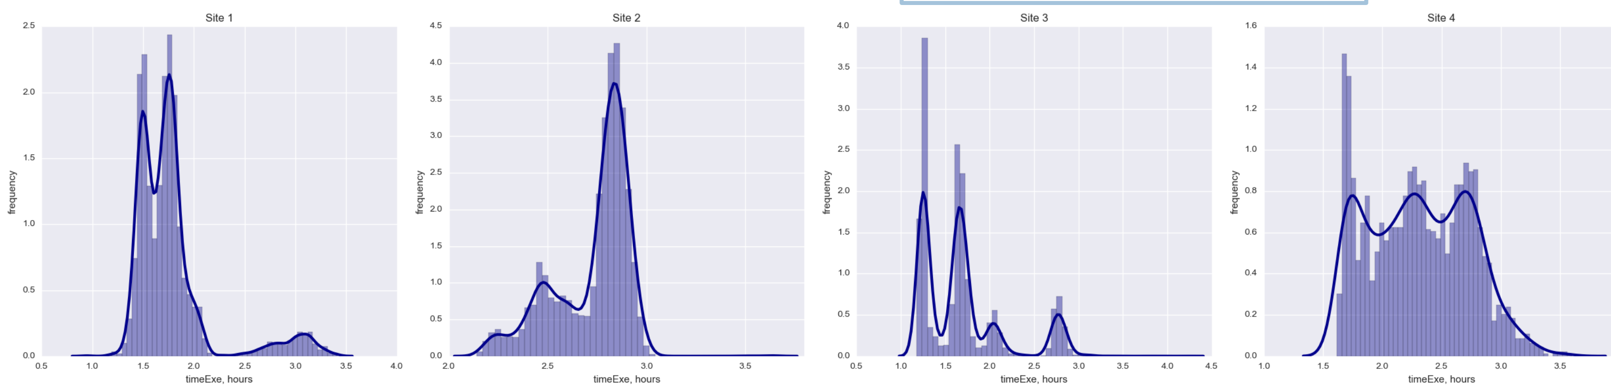

Bars indicates the number of jobs, having the same execution time.

Lines – univariate distribution of observations.

CPUTime Distributions

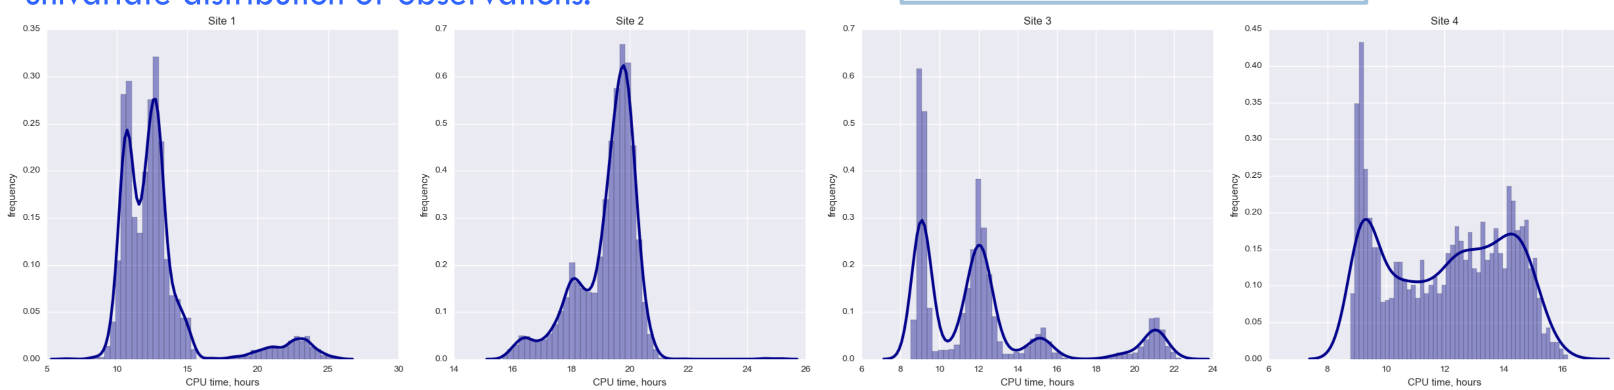

Non-trivial jobs execution on computing site

- **<u>Task №14296407</u>** (may-june 2018), number of jobs ~ 9000, 18 computing sites

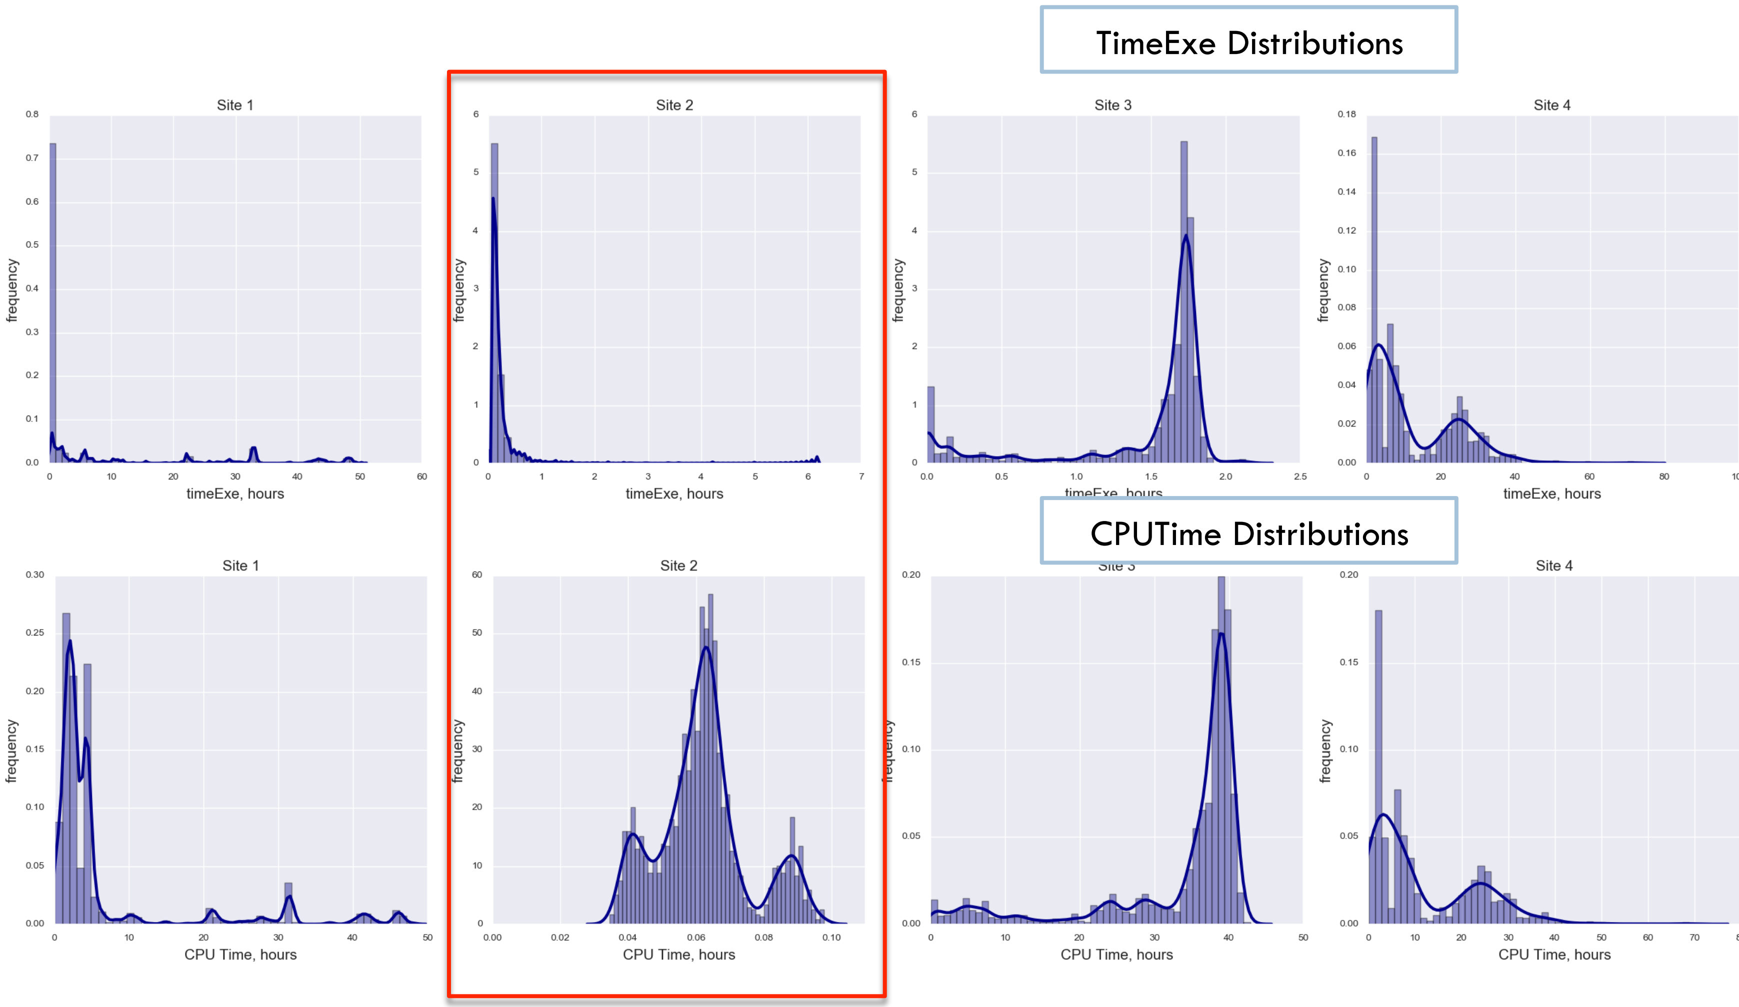

Site 2 analysis

Task №14296407

- **Remove unnecessary features**
  - Features with a high percentage of missing values
  - Collinear (highly correlated) features
  - Features with a single unique value
- Dependent feature – **WallTime**
- **RandomForestRegressor** - to retrieve features importances for the dependent feature
- **K-means clustering** - to split initial data sample into 200 data clusters
- Visual analytics prototype was applied to this data sample to build 3D spatial scene
- Distance threshold between spheres, calculated as **Mahalanobis distance,** was tuned using interactive interface

attributes	importance
cpu_eff	0.712740
IObytesReadRate	0.184547
cpuconsumptiontime	0.034778
timeSetup	0.016238
timeStageIn	0.009818
IObytesRead	0.009795
workDirSize	0.008707
avgswap	0.008375
inputfilebytes	0.003622
timeStageOut	0.001935

Clusters and Anomalies

Task №14296407. Site 2

2 clusters were found:

- Large cluster with the average wall time of **25 minutes (normal)**
- Small cluster with small wall time of **10 minutes (best)**

Irregular points were found with a very long wall time **(227 minutes)!**

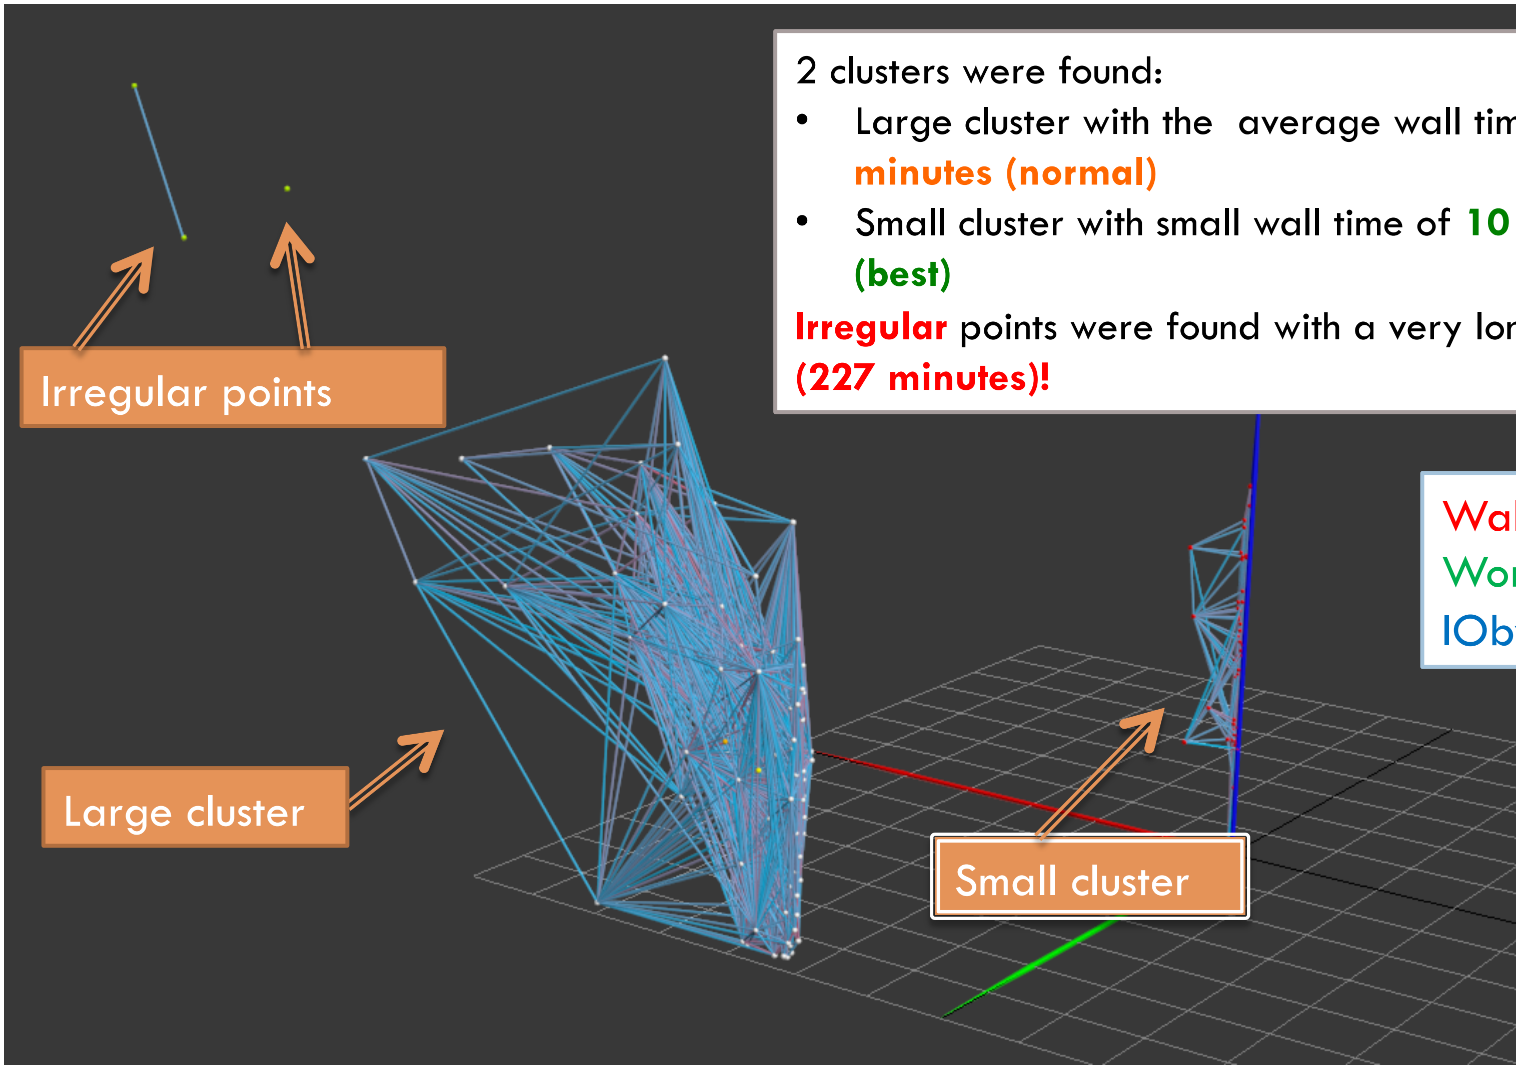

Results of cluster analysis. Time & Memory

Task №14296407. Site 2

Features	Large cluster	Small cluster	Irregular points
WallTime	25 min	**10 min**	**227 minutes**
CPUTime	3,8 min	3,2 min	3,5 min
TimeSetup	36 sec	**13 sec**	**120 sec**
TimeStageIn	370 sec	**110 sec**	356 sec
TimeStageOut	59 sec	**33 sec**	**680 sec**
DBTime	20 sec	**7 sec**	**266 sec**
QueueTime	32 min	64 min	58 min
AvgRSS	544	544	530
MaxRSS	825	817	814
AvgVmem	1 813	1 731	2 014
MaxVmem	3 041	2 767	3 056

Megabytes

Results of cluster analysis. IO Memory Metrics

Task №14296407. Site 2

Features	Large Cluster	Small cluster	Irregular points
Inputfilebytes	300	300	300
Outputfilebytes	600	600	600
IObytesWritten	531	488	568
IObytesRead	1 957	1 704	**3 029**
IOcharRead	1 088	976	**1 581**
IOcharWriten	526	484	561
WorkDirSize	600	**8**	600
IObytesReadRate	3,868	**6,716**	**0,690**
IObytesWriteRate	1,068	**1,945**	**0,180**
IOcharReadRate	2,146	**3,877**	**0,457**
IOcharWriteRate	1,059	**1,931**	**0,180**

Megabytes

MB/sec

Analysis of clusters in the initial dataset

Task №14296407. Site 2

- Metrics of 2 clusters and irregular points were analyzed. The following were found:
  - The CPU time is in the expected range in all clusters and irregular points
  - The amounts of RAM and virtual memory are almost the same in all clusters and irregular points
  - Input and output files sizes are 300 and 600 Mb respectively for all clusters and points
  - Observed input data read much larger than single input file sizes
    - 6 times larger for normal cluster
    - 10 times larger for irregular points
  - Written data is close to the output file size

Analysis of clusters in the initial datasample

Task №14296407. Site 2

- Irregular points have the following peculiar properties (in comparison with the large cluster):
  - timeSetup - 3 times longer
  - timeStageOut – 10 times longer
  - dbTime – 13 times longer
  - QueueTime – twice longer
  - ReadRate - 5 times slower
  - WriteRate - 6 times slower
- The small cluster has the highest rates of data read/write (twice larger than in large cluster) and the shortest wall time
  - Probably it can be connected with the workDirSize, which is the smallest for this cluster

- The amount of read information is a subject to further investigation and analysis. Hypothesis of the reasons could be:
  - Failed jobs on the same site lead to overload of the data streams
  - Input of the data failed and had to start from the beginning

Conclusions

- The metodology of data analysis with the combined usage of statistics, machine learning and visual analytics methods was proposed
- The first prototype of interactive visual analytics platform was developed on the basis of 3dsMax
- The developed metodology and visual prototype were applied to the analysis of ATLAS jobs execution

Future plans

- Increase the size of investigated data samples
- Enhance numerical features with categorical
- Apply new ATLAS data sources (AGIS, NWS) providing the information about sites and network status during jobs execution
- The development of the visual analytics tools
  - Extend the application to dynamic data analysis
  - Use open-source platform (e.g., VTK)

Acknowledgements

- This work was supported by RSCF grant №18-71-10003
- Thanks to I. Vukotic, A. Anisenkov, J. Elmsheuser
- Special thanks to A. Klimentov