Laboratory for
Parallel Programming

TECHNISCHE UNIVERSITÄT DARMSTADT

Bachelor's / Master's Thesis

KPI Visualization for Performance Modeling of Cloud Applications and Systems

Motivation

Cloud computing plays a crucial role in the digitalization. Therefore, the number of cloud services and providers is constantly rising. To define the performance characteristics of a service cloud providers make service-level-agreements (SLAs) with their customers determining the level of service quality e.g. the tolerable failure rate. To ensure that the providers can always meet these SLA agreements, a thorough performance analysis of each cloud service is required. Only when a provider knows the breaking points of each service, he can adjust his system and resource allocations accordingly, allowing an increased system efficiency and reduced operating costs. Existing methods can be used to analyze the performance characteristics of cloud services, however, the outcomes are often difficult to understand and require expert knowledge to gain insights into possible service or system optimization's.

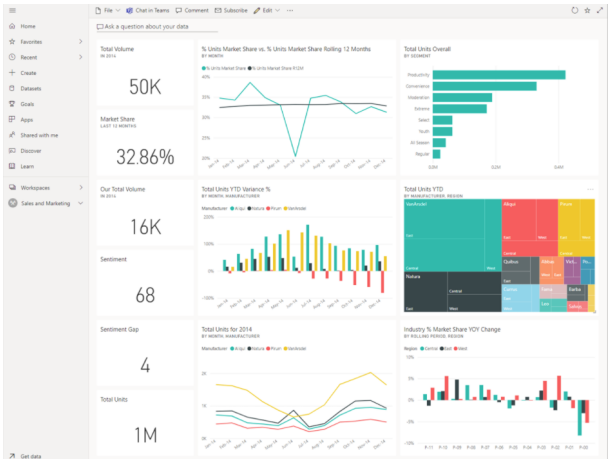

Example for the visualization of different performance metrics in a dashboard report.

Task

The target of this thesis project is to implement a tool that uses the output of Extra-P to automatically create a performance analysis report and dashboard for High-Performance-Computing (HPC) and Cloud applications. Furthermore, we want to analyze the interaction of several Cloud applications and create entire System reports and dashboards. In order to achieve this goal the student will work on the following tasks:

- Research the state of the art on visualization techniques for scientific data using Python or other tools, business reports and dashboard creation
- Identify the Key-Performance-Indicators (KPIs) for the performance analysis of HPC and Cloud applications
- Design and develop a visualization concept for the KPIs and the results of our performance analysis
- Develop an approach that uses the output of Extra-P to automatically create a business report and dashboard containing the most important performance data/insights of an application
- Evaluate your solution with literature, surveys, and user trials

Requirements

- Python (matplotlib or other visualization tools), Parallel Programming
- Experience with the visualization of scientific data
- Optional: Experience with Cloud Systems

Contact

Marcus Ritter <marcus.ritter@tu-darmstadt.de>

References

[1] Alexandru Calotoiu, Torsten Hoefler, Marius Poke, Felix Wolf: Using Automated Performance Modeling to Find Scalability Bugs in Complex Codes. In Proc. of the ACM/IEEE Conference on Supercomputing (SC13), Denver, CO, USA, pages 1–12, ACM, November 2013. PDF

Issue date: October 2021