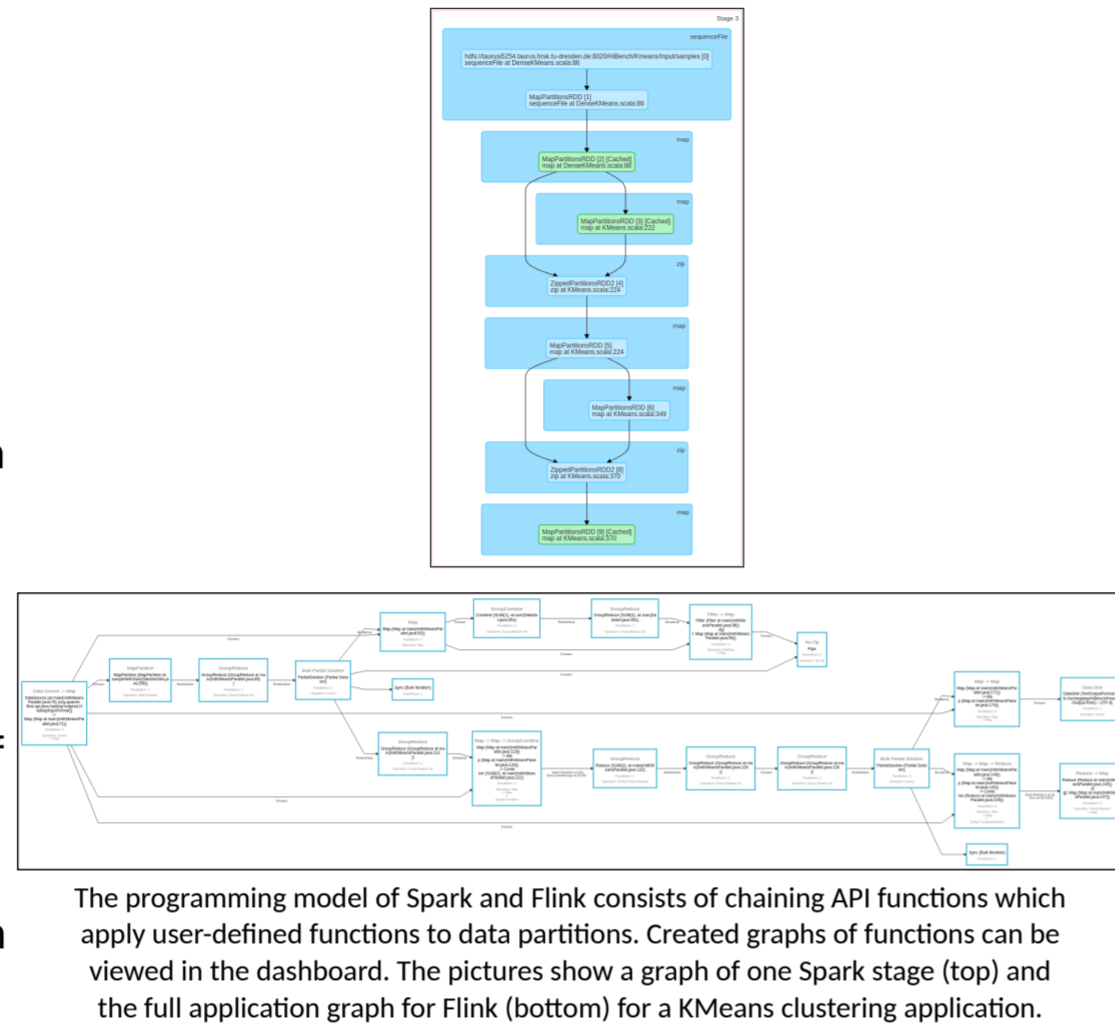


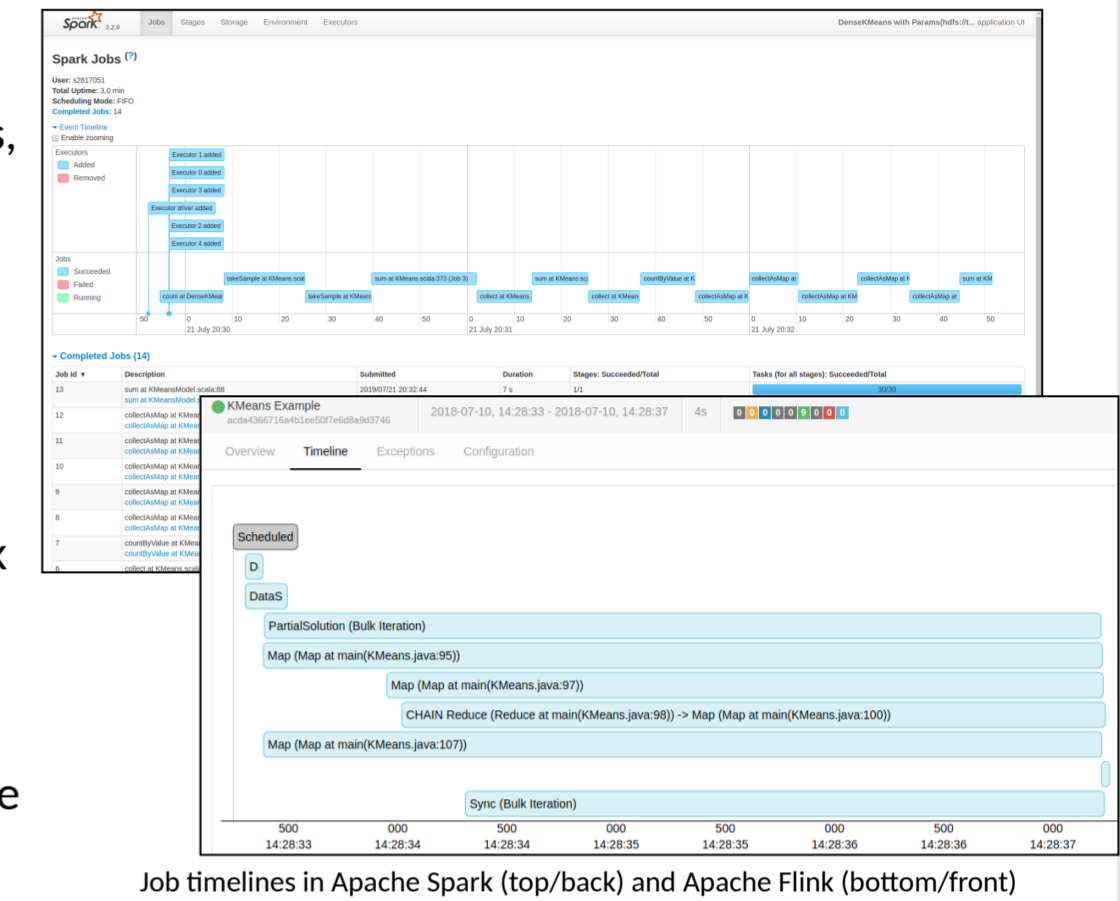
1) Introduction

- Apache Spark and Apache Flink are two Big Data frameworks used for fast data exploration and analysis.
- Applications are implemented by chaining API functions which are customized by user-defined functions.
- Performance investigation and optimization is limited by the integrated dashboards and monitoring systems.
- The interaction of the user code and the framework is encapsulated in the programming model, i. e. the abstraction of second-order API functions.
- Due to this hidden information, performance investigation and optimization of workflows is hard.



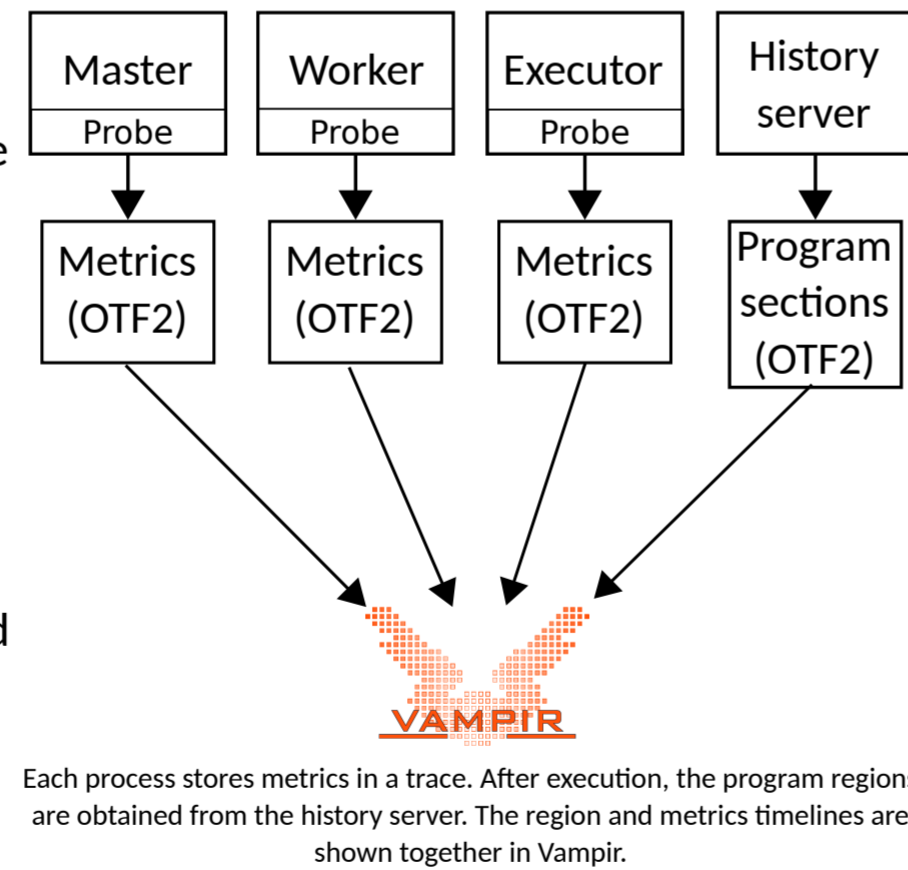
2) Challenges

- Monitoring the execution with the integrated dashboards gives only aggregated data indicating performance of the workflow steps, e. g. numbers of records.
- The separation of timelines (e. g. jobs, stages and tasks in Apache Spark) and metrics (aggregated for e. g. jobs or tasks) in the dashboard limits insights into the workflow characteristics over time.
- A performance comparison of Spark and Flink requires to compare multiple views visually.
- Established performance investigation tools from the HPC area, such as Vampir, are not supported and cannot be directly applied, due to missing programming language support, e. g. Java or Scala.



3) Proposed approach

- Idea: Performance data sampled at runtime enhances data collected by in-built framework monitoring (history server).
- The framework-provided monitoring interfaces for the Java management extensions (JMX) system are enabled in order to be able to collect framework and JVM-related performance metrics.
- Probes are inserted in the processes to sample the JMX metrics in regular time intervals.
- The probes store data in Open Trace Format 2 (OTF2).
- After the execution, a separate process converts start and end times of program sections, such as stages and tasks, from the history server to OTF2.
- Trace files with exact time information and values of metrics are combined with timely information of program sections and displayed by Vampir.



4) Results and future work

- Our proposed work integrates the overview of program sections and related performance metrics.
- The view contains data of all levels of the hierarchical execution management model, e. g. stages and tasks.
- The overhead of the metrics sampling depends on the time interval between samples (<10% for 20 metrics @100Hz).
- Data for performance investigation is stored in an established format independent from Spark or Flink versions, so that state-of-the-art performance tools, i. e. Vampir, can be used for analysis.
- The system allows to enrich the traces with further data, such as call stack data gathered via Java bytecode instrumentation and to compare different implementations for a problem.

