DIAS: Data-Intensive Applications and Systems Laboratory
School of Computer and Communication Sciences
Ecole Polytechnique Fédérale de Lausanne
Building BC, Station 14
CH-1015 Lausanne
URL: http://dias.epfl.ch/

Optimizing memory allocation in streaming systems

Keywords: stream processing, memory allocation, resource management

Problem: Streaming engines are facing ever-growing pressure to handle queries with larger and larger memories. This happens because of two reasons: First, data streams grow exponentially in volume and velocity, and as a result, the state of stateful operators like joins and aggregates grows as well. Second, data analysts seek to extract deeper insights with increasingly sophisticated queries that have more and more stateful operators. The current approach to memory management in most streaming engines relies on static and naive allocation strategies that fail to account for the specific needs of different operators or queries [1]. This inflexible allocation leads to significant inefficiencies where some operators are allocated excess memory while others struggle with insufficient resources, forcing them to spill state to disk and introducing substantial latency penalties.

Project: This project will begin with a systematic investigation of the factors influencing memory demands across different streaming operators and query types, quantifying the impact of various parameters such as window size, key cardinality, and data distribution. Building on these insights, the project will evaluate the inefficiencies of current memory allocation strategies and their performance implications. Finally, the student will design and implement a memory allocation mechanism that optimizes resource distribution across concurrent streaming tasks, with the goal of maximizing overall system performance.

Scope: Master's semester project

Plan:

1. Familiarize oneself with background work on stream processing and resource allocation
2. Profile memory usage patterns of different streaming operators (e.g., joins, aggregations, windows) under varying workload conditions and parameters
3. Analyze and quantify memory waste in current static allocation approaches through experimental evaluation
4. Design and implement a memory allocation mechanism that distributes memory resources across operators based on their demands

References:

1. Marios Fragkoulis, Paris Carbone & Vasiliki Kalavri, Asterios Katsifodimos. (2023). A survey on the evolution of stream processing systems.

Supervisor: Prof. Anastasia Ailamaki, anastasia.ailamaki@epfl.ch
Responsible collaborator(s): Eleni Zapridou