

Performance Optimizations and Operator Semantics for Streaming Data Flow Programs

Abstract:

Global scale and Internet native companies like Google collect more data than ever and require insights from it quicker than before. The database research community started to address these trends in the early 2000th and two new research topics gained major interest: large-scale non-relational data processing as well as low latency data stream processing.

Both research areas started to overlap with the development of MapReduce-like large scale distributed data stream processing systems. While those systems gain more and more attention in industry there are still mayor challenges to operator those system at large scale.

Provisioning those systems and runtime tuning of queries is still a manual, error prone, and time consuming process carried out by experts. Furthermore, there is still no agreement for semantics of continuous data stream processing. Different systems offer different semantics and often suffer from non-deterministic query execution.

The goal of this thesis is two fold. First, we investigate runtime characteristics of large scale distributed streaming system to better understand system and query runtime behavior with the aim to provision queries automatically. We introduce a cost model for streaming data flow programs and introduce different optimization algorithms based on our cost model to provision queries in a cost based manner. Second, we suggest the "Dual Streaming Model" to express semantics of continuous queries over data streams and tables. In our model, inconsistencies of logical and physical record order are handled within the model itself allowing for deterministic semantics as well as low latency query execution. We formally define the "Dual Streaming Model" and discuss the differences and advantages compared to existing approaches.