

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE
UNIVERSITY OF SOUTHERN DENMARK, ODENSE

COMPUTER SCIENCE COLLOQUIUM

Dynamic Load Balancing within Parallelized Stream Processing Operators

Zhenjie Zhang
Advanced Digital Sciences Center
University of Illinois at Urbana Champaign

Tuesday, 10 May, 2016 at 14:15
IMADA's Seminar Room

Abstract:

Real-time streaming analytics is arising in recent years as an important computation scheme to support fast response to big data with large volume and high velocity. While mainstream systems, e.g. Storm and Spark Streaming, may adopt different programming paradigms for streaming analytical applications, intra-operator parallelism is commonly employed to improve processing throughput and exploit the computation resource. Dynamic loading balancing turns out to be the key enabler of high efficiency for parallelized operators in stream processing, in the sense that processing performance is optimized when the workload is evenly partitioned within the operator. In this paper, we address the algorithmic complexity behind the dynamic load balancing problem for intra-operator parallelism, with a new mixed partitioning scheme taking advantage of both hash-based and key-based workload partition strategies. Instead of minimizing variance of the assignment, our dynamic load balancing algorithm aims to minimize the data migration cost under load balancing constraint, when the distribution and scale of incoming data stream change. We present robust algorithms with strong theoretical guarantee under bicriteria approximation framework, to tackle the NP-hard problem. We also discuss the communication complexity when implementing the algorithm on real distributed cluster or cloud computing platform.

Host: Yongluan Zhou