

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

COMPUTER SCIENCE COLLOQUIUM

S3: An Efficient Shared-Scan Scheduler on MapReduce Framework

Lei Shi

Department of Computer Science
National University of Singapore

Tuesday, 29 November, 2011 at 14:15

IMADA's Seminar Room

Abstract:

Hadoop, an open-source implementation of MapReduce, has been widely used for data-intensive computing. In order to improve performance, multiple jobs operating on a common data file can be processed as a batch to share the cost of scanning the file. However, in practice, jobs do not arrive at the same time, and batching them means longer waiting time for jobs that arrive earlier.

To handle such problem, we propose S3 - a novel Shared Scan Scheduler for Hadoop - that shares the scanning of a common file for multiple jobs that may arrive at different times. Essentially, a file is organized into, say k , segments. A job accessing the file is then processed in k iterations - in each iteration, a sub-job is initiated to operate on one segment. Moreover, S3 allows a job to be scheduled for processing from any segment. In other words, instead of having to wait for a long time for the processing of an entire file to complete, we can now initiate the processing as soon as the processing of the current segment ends. In addition, sub-jobs from multiple jobs accessing the same segment can be processed collectively by a single access to the segment. Thus, under S3, as jobs are submitted and completed, the number of sub-jobs per batch changes. We have implemented our S3 approach in Hadoop, and our experimental results on a local cluster show that S3 outperforms the naïve no-sharing scheme and the file-based shared-scan approach.

Host: Yongluan Zhou