



Spring 2008 CIS Colloquium Series

Fast and Highly-Available Stream Processing

Jeong-Hyon Hwang

February 13, 2008

Abstract

Recently, there has been significant interest in applications where high-volume, continuous data streams need to be processed with low latency. These applications include financial market monitoring, network monitoring, sensor-based environment monitoring, call analysis, battlefield monitoring, asset tracking, and Web feed analysis. To facilitate the applications, several stream-processing systems have been developed. In these systems, queries are represented as a network of operators that transform the data streaming through them. This talk focuses on both fast and reliable processing of real-time data streams despite slow or failed servers and network links.

In this talk, I will first introduce our basic recovery models while considering recovery semantics, recovery speed, and resource usage. Next, I will present a fast recovery technique for commodity server clusters. In this technique, operators on each server are backed up on different servers and thus can be recovered in parallel. This technique assigns backup servers and schedules checkpoints in a manner that maximizes the recovery speed. Finally, I will discuss an approach for Internet-scale stream processing. In this approach, multiple operator replicas send outputs to downstream replicas, allowing each replica to use whichever data arrives first. To further reduce latency, replicas run without coordination, possibly processing data in different orders. Despite this, the approach always delivers applications the same results as in the non-replicated, failure-free case. It also deploys replicas at locations that effectively improve performance as well as availability.

This work was done in the context of Aurora/Borealis, one of the first distributed stream-processing systems (<http://www.cs.brown.edu/research/borealis>). The experimental results were obtained from two real testbeds, a server cluster at Brown University and a worldwide network platform called PlanetLab.

Bio

Jeong-Hyon Hwang is a Ph.D. candidate in the Department of Computer Science at Brown University. His recent work has centered on fast, fault-tolerant, and scalable processing of real-time data streams. His research interests include database systems, distributed systems, and networking. He received M.S. and B.S. degrees in Computer Science from Brown University in 2003 and Korea University in 1998, respectively.

Location: 4th Floor Conference Room (Wachman 447)

Time: TBD, Wednesday, February 13, 2008

Refreshments will be served!