The University of Chicago
Department of Computer Science
Seminar

PRESENTS:

Sailesh Krishnamurthy
Amazon Web Services

Title: " Amazon Aurora: Design considerations for high throughput cloud–Native relational databases."

Abstract: Amazon Aurora is a relational database service for OLTP workloads offered as part of Amazon Web Services (AWS). In this talk, we describe the architecture of Aurora and the design considerations leading to that architecture. We believe the central constraint in high throughput data processing has moved from compute and storage to the network. Aurora brings a novel architecture to the relational database to address this constraint, most notably by pushing redo processing to a multi–tenant scale–out storage service, purpose–built for Aurora. We describe how doing so not only reduces network traffic, but also allows for fast crash recovery, failovers to replicas without loss of data, and fault–tolerant, self–healing storage. Traditional implementations that leverage distributed storage would use distributed consensus algorithms for commits, reads, replication, and membership changes and amplify the cost of underlying storage. We will describe how Aurora avoids distributed consensus under most circumstances by establishing invariants and leveraging local transient state. These techniques improve performance, reduce variability, and lowers costs.

Bio:

Sailesh Krishnamurthy is a General Manager at Amazon Web Services where he runs engineering, go–to–market and operations for RDS Aurora, MySQL and MariaDB. He is an innovator and entrepreneur with 20+ years of experience in databases and the cloud. Prior to Amazon, Sailesh was at Cisco Systems via the acquisition of Truviso, a real–time streaming data analytics software company that he co–founded. Sailesh is an authority in data management and is an author of over a dozen academic papers and several issued U.S. patents. He has a Ph.D. in Computer Science from UC Berkeley in 2006.

Friday, June 21, 2019

JCL 390 @ 11:00 am