

TAPAN SRIVASTAVA

tapansriv@uchicago.edu · tapansrivastava.com

EDUCATION

Ph.D., Computer Science

University of Chicago

Advised by Prof. Raul Castro Fernandez

2020 – Present

M.S., Computer Science

University of Chicago

Advised by Prof. Raul Castro Fernandez

Master's Paper: *Saving Money For Analytical Workloads in the Cloud*

2020 – 2023

B.S. Honors, Computer Science

University of Chicago

Bachelor's Thesis Advised by Prof. Hank Hoffmann

2016 – 2020

GPA: 3.90/4.0

B.S., Mathematics

University of Chicago

2016 – 2020

GPA: 3.88/4.0

SELECTED PROJECTS

Saving Money for Analytical Workloads in the Cloud

University of Chicago

As users migrate their analytical workloads to cloud databases, it is becoming just as important to reduce monetary cost as it is to improve query performance. In the cloud, users choose to pay based on either the compute time or the amount of data a query reads. We observe that analytical queries are compute- or IO-intensive and each query type executes cheaper in a different pricing model. We exploit this opportunity and propose methods to build cheaper execution plans across pricing models. We implement these methods, observe how SQL text and query optimizers impact cost, and consider a transparent deployment of DuckDB on infrastructure-as-a-service to avoid the cost premiums of platform-as-a-service without placing the burden of deployment or maintenance on users. We reduce workload costs by as much as 80% and produce execution plans spanning multiple pricing models that save as much as 47%. We also reduce individual query costs by as much as 90%. We simulate the effect of different cloud prices on real query performance and observe that there are still significant multi-cloud opportunities if cloud vendors change their prices. These results indicate the massive opportunity to save money by building execution plans across multiple pricing models.

Penelope: Peer-to-peer Power Management

University of Chicago

Large scale distributed computing setups rely on power management systems to enforce tight power budgets. Existing systems use a central authority that redistributes excess power to power-hungry nodes. This central authority, however, is both a single point of failure and a critical bottleneck—especially at large scale. To address these limitations we propose *Penelope*, a distributed power management system which shifts power through peer-to-peer transactions, ensuring that it remains robust in faulty environments and at large scale. We implement *Penelope* and compare its achieved performance to *SLURM*, a centralized power manager, under a variety of power budgets. We find that under normal conditions *SLURM* and *Penelope* achieve almost equivalent performance; however in faulty environments, *Penelope* achieves 8–15% mean application performance gains over *SLURM*. At large scale and with increasing frequency of messages, *Penelope* maintains its performance in contrast to centralized approaches which degrade and become unusable. This work was published at ICPP'22.

PUBLICATIONS

Penelope: Peer-to-peer Power Management

ICPP '22: Proceedings of the 51st International Conference on Parallel Processing

Tapan Srivastava, Huazhe Zhang, Henry Hoffmann
10.1145/3545008.3545047

POSTERS

Saving Money For Analytical Workloads in the Cloud

GCASR: Greater Chicago Area Systems Research Workshop

Tapan Srivastava, Raul Castro Fernandez

Illinois Institute of Technology, April 2023

Saving Money For Analytical Workloads in the Cloud

CERES Summit

Tapan Srivastava, Raul Castro Fernandez

University of Chicago, May 2023

PROFESSIONAL
EXPERIENCE

Software Engineering Intern

Pittsburgh, PA

Uber ATG

- Built backend services to create a debugging tool for the tablets onboard the self-driving vehicle.
- Utilized gRPC servers and protobuf to effectively log data for playback during debugging.

Software Engineering Intern

Pittsburgh, PA

Uber ATG

- Migrated an existing service running onboard the self-driving car from using an HTTP server to using a gRPC server.
- Designed, implemented, and deployed a new backend service for a feature on the self-driving car.
- Implemented new service in Go and Thrift, logged via Kafka streams, and utilized a distributed task orchestrator to manage workflows in a distributed production environment.
- Deployed the service into production.

HONORS AND
AWARDS

University of Chicago Dean's List

2017-2020

University of Chicago Presidential Scholarship

2016

University of Chicago University Scholarship

2016

RELEVANT
COURSES

Computer Science: • Databases • Operating Systems • Complexity Theory • Computer Architecture • Advanced Distributed Systems

Mathematics: • Analysis • Abstract Algebra • Linear Algebra
• Discrete Mathematics • Statistical Models and Methods

TECHNICAL
SKILLS

Languages: C/C++, Java, Python, Bash, L^AT_EX

General: Data Structures, Algorithms, Object Oriented Programming

ADDITIONAL
ACTIVITIES

- Created iPhone app to analyze Jazz chord changes
- First Chair Alto Sax, Overlake High School Jazz Band
- Intramural Spades Finalist 2019