

AARATI KAKARAPARTHY

Email: aaratik@cs.wisc.edu
Mobile: +1 (608) 960 0341
Website: <http://pages.cs.wisc.edu/~aaratik>
Github: <https://github.com/aarati-K>



EDUCATION

Program	Institution	GPA	Year
PhD in Computer Science	University of Wisconsin, Madison	3.89/4	2019 – present
Master's in Computer Science	University of Wisconsin, Madison	3.88/4	2017 – 2019
Bachelor's in Computer Science and Engineering (Minor: Microelectronics)	IIT Madras	9.13/10	2011 – 2015

RESEARCH INTERESTS

Using Machine Learning to make database management systems adaptive to workload properties, execution environment, and emerging hardware technologies.

PUBLICATIONS

- Optimizing Databases by Learning Hidden Parameters of Solid State Drives** (VLDB 2020)
Aarati Kakaraparthi, Jignesh M. Patel, Kwanghyun Park, Brian P. Kroth
- The Case for Unifying Data Loading in Machine Learning Clusters** (HotCloud 2019)
Aarati Kakaraparthi, Abhay Venkatesh, Amar Phanishayee, Shivaram Venkataraman

PROFESSIONAL EXPERIENCE

- Research Intern with DMX Group at Microsoft, Redmond** (June – August 2020)
advised by Vivek Narasayya and Christian König
- Predicting multiple resource usage metrics of cloud databases.
 - Exploring scheduling policies to improve the resource utilization of cloud databases.
- Research Assistant at Microsoft Gray Systems Lab** (since January 2018)
advised by Prof. Jignesh Patel, UW Madison
- Improving the performance of database operations by studying the characteristics of commercial SSDs, and investigating applicability to database systems.
 - Studying the implementation of common database operations on FPGAs using Microsoft's Catapult architecture.
- Full Stack Developer at Zenefits in Vancouver, Canada & Bangalore, India** (October 2015 – July 2017)
- Developed RESTful API services to integrate third-party applications with the Zenefits platform, as part of the platform engineering team.
 - Managed end-to-end lifecycle of the services, from development and testing, to deployment and monitoring.
 - Redesigned the Zenefits dashboard codebase, while collaborating across teams to integrate multiple sub-products.
 - Won second prize in an internal hackathon with around 50 participating teams.
- Summer Intern at Samsung Electronics in Suwon, South Korea** (May – July 2014)
- Implemented an SMS spam detector as a combination of a generalized global model and a personalized local model.
 - Designed a personalized lightweight entropy-based classification model for local classification.
 - Implemented a global classifier trained on input from multiple users.
- Summer Intern at Google in Bangalore, India** (May – July 2013)
- Designed a topic recommendation system for Google Baraza (a question-answer website similar to Quora), depending on the users' community and past activity.
 - Implemented a prototype recommendation system using Flume and BigTable.

SCHOLASTIC ACHIEVEMENTS

- **All India Rank 108** in IIT-JEE (*from among ~500,000 students, top 0.02% nationwide*) 2011
- **All India Rank 947** in AIEEE (*from among ~1.1 million students, top 0.1% nationwide*) 2011
- **Top 1% nationwide** in Indian National Physics, Chemistry and Astronomy Olympiads 2009

COURSES

- Topics in DBMS
- Big Data Systems
- Advanced Computer Architecture I
- Advanced Operating Systems
- Mathematical Foundations of ML
- Data Mining
- Advanced Computer Networks
- Cryptography and Network Security
- Advanced Theory of Computation

GRADUATE-LEVEL PROJECTS (*at UW Madison*)

Optimizing Database Operations by Learning Hidden Parameters of SSDs (*May 2018 – October 2019*)
with Microsoft Gray Systems Lab and Prof. Jignesh Patel, UW Madison

- Research methods to reduce latency of requests on SSDs and understanding their internal operations.
- Proposed new techniques and implemented them in SQLite3 and MariaDB, which led to a significant improvement of 29% in the read throughput.
- Proposed methods to deduce important parameters of SSDs, and their applications to database systems.

OneAccess – A Unified Data Access Layer for Machine Learning (*October 2018 – March 2019*)
with Prof. Shivaram Venkataraman, UW Madison

- Attempt to develop synergy between ML frameworks and the storage layer for better performance.
- Developed OneAccess, which improves the performance of data loading by emphasizing sequential accesses through reservoir sampling.
- > 3.5x improvement in loading time over PyTorch's in-built data loader, for the MS-COCO Detection dataset.

Coflow Network Scheduling in Datacenters *with Prof. Aditya Akella, UW Madison* (*March – May 2018*)

- Researched the problem of network scheduling in datacenter settings, and identified drawbacks in existing techniques.
- Proposed a heuristic for network coflow scheduling which improved the job completion time by 10% for synthetic workloads.

Operator Fusion in Quickstep DBMS *with Prof. Jignesh Patel, UW Madison* (*October – December 2017*)

- Developed an in-depth knowledge of Apache Quickstep, an open source OLTP database engine.
- Implemented pipelining of output across join operators during query execution, and studied its impact.
- Identified limitations in the design of Quickstep which resulted in rigidity towards pipelining.

UNDERGRADUATE PROJECTS (*at IIT Madras*)

A Study of the Disjoint Cover and Protocol Cover in Communication Complexity (*Jan – May 2015*)
with Prof. Jayalal Sarma, IIT Madras, India

- Researched the problem of the gap between disjoint cover and protocol cover in Communication Complexity.
- Established the equivalence of the disjoint cover and protocol cover to a unique constrained set system, and attempted to solve the equivalent combinatorial problem.

Finding Anomalous Trajectories in Road Network Data (*April 2014*)

- Implemented an Isolation-based approach to detect anomalous trajectories.
- Used a forest-based approach to calculate anomaly score.
- Tested the model on the T-Drive data-set published by Microsoft.

SKILLS

- **Programming Languages:** *Proficient* in C, C++, Python
Familiar with Java, Javascript, HTML, CSS
- **Open Source:** *Familiar with the source code of* SQLite3, MariaDB, PyTorch

COMMUNITY INVOLVEMENT

Bi-weekly Activities Chair of Women's ACM (WACM) *in UW Madison* (*September 2018 – May 2019*)

- Responsible for organizing events directed towards forming a community of women in computing, that provides social, professional, advisory, and mentorship opportunities.
- Conducted a session teaching basic Linux commands to over 100 beginners and undergraduate students.

Participation in Computer Sciences Education activities *in UW Madison* (*May 2018 – May 2019*)

- Introduced programming in Scratch to a group of 4th and 5th graders, as part of an 8-week summer club.
- Tutored undergraduates as part of the Computer Sciences Learning Center, to gain experience in general teaching practices.

Table Tennis Team Captain, Sharavati Hostel *in IIT Madras* (*August 2012 – May 2013*)

- Managed hostel table tennis budget and organized intra-hostel table tennis tournaments.