

Sujaya Maiyya

+1-805-837-9626 • smaiyya@uwaterloo.ca
<https://cs.uwaterloo.ca/smaiyya/>

My interdisciplinary research spans distributed systems, databases, and data security. **I design, prototype, and evaluate protocols for managing geo-distributed data in a performance efficient and secure manner.** My research focuses on integrating cryptographic techniques with databases to mitigate data integrity and data privacy vulnerabilities in traditional database systems. The increasing popularity of blockchains and the use of third-party cloud – scenarios in which the servers storing and processing the data are not trusted – makes my research highly relevant to today's problems.

Work

- **University of Waterloo** **Waterloo, Canada**
Assistant Professor *September 2022 - Present*

Education

- **PhD in Computer Science** **Santa Barbara, USA**
UC Santa Barbara *January 2018 - June 2022*
- **MSc in Computer Science** **Santa Barbara, USA**
UC Santa Barbara *Sept 2016 - December 2017*
- **B.E. in Information Science** **Bangalore, India**
P.E.S Institute of Technology *August 2010 - May 2014*

Awards and Accomplishments

- **MIT EECS Rising Stars** *August 2021*
- **IBM PhD Fellowship 2020** award *March 2020*
- **Google PhD Fellowship 2020** award (declined) *April 2020*
- **Outstanding Graduate Student 2018-19** award in the UCSB CS dept. *June 2019*
- Elected **Graduate Student Representative** (thrice) at UCSB *2018 - 2020*
- **Spot Bonus** award for leadership qualities at Citrix *May 2016*
- **SAP Labs Scholarship** (4 times) for highest CGPA in the dept. at PESIT *2012-2014*
- **Second Prize at Intel CCF Innovation Contest** *August 2013*
- **Summer Research Fellowship** from Human Resource Development Ministry *March 2013*
- **Second Prize** at Ayana Hackathon *May 2012*

Publications

In Progress Publications

Sujaya Maiyya, P. Ananth, D. Agrawal, and A. El Abbadi. ORTOA: One Round Trip Oblivious Access. <https://eprint.iacr.org/2022/1506>. Accessed: 2022-11-11.

Refereed Conference Publications

MJ. Amiri, D. Shu, **Sujaya Maiyya**, D. Agrawal, and A. El Abbadi. Ziziphus: Scalable Data Management Across Byzantine Edge Servers. *IEEE ICDE*, 2023.

Sujaya Maiyya, S. Ibrahim, C. Scarberry, D. Agrawal, A. El Abbadi, H. Lin, S. Tessaro, and V. Zakhary. Quoram: A quorum-replicated fault tolerant oram datastore. *USENIX Security*, 2022.

F. Zhao, **Sujaya Maiyya**, R. Wiener, D. Agrawal, and A. El Abbadi. KII[±]: Approximate quantile sketches

over dynamic datasets. *Proceedings of the VLDB Endowment*, 2021.

Sujaya Maiyya, I. Ahmad, D. Agrawal, and A. El Abbadi. Samya: Geo-distributed data system for high contention data aggregates. *International Conference on Data Engineering (ICDE)*, 2021.

Sujaya Maiyya, DBH. Cho, D. Agrawal, and A. El Abbadi. Fides: Managing data on untrusted infrastructure. *International Conference on Distributed Computing Systems (ICDCS)*, 2020.

MJ. Amiri, **Sujaya Maiyya**, D. Agrawal, and A. El Abbadi. Seemore: A fault-tolerant protocol for hybrid cloud environments. *International Conference on Data Engineering (ICDE)*, 2020.

Sujaya Maiyya, F. Nawab, D. Agrawal, and A. El Abbadi. Unifying consensus and atomic commitment for effective cloud data management. *Proceedings of the VLDB Endowment*, 12(5):611–623, 2019.

V. Arora, RKS. Babu, **Sujaya Maiyya**, D. Agrawal, A. El Abbadi, X. Xue, et al. Dynamic timestamp allocation for reducing transaction aborts. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*, pages 269–276. IEEE, 2018.

Refereed Conference Tutorials

Sujaya Maiyya, V. Zakhary, MJ. Amiri, D. Agrawal, and A. El Abbadi. Database and distributed computing foundations of blockchains. In *SIGMOD*, 2019.

Sujaya Maiyya, V. Zakhary, D. Agrawal, and A. El Abbadi. Database and distributed computing fundamentals for scalable, fault-tolerant, and consistent maintenance of blockchains. *Proceedings of the VLDB Endowment*, 11(12):2098–2101, 2018.

Refereed Conference Workshops

V. Zakhary, MJ Amiri, **Sujaya Maiyya**, D. Agrawal, and A. El Abbadi. Towards global asset management in blockchain systems. *BCDL co-located with VLDB*, 2019.

Research and Industrial Work Experience

- **Cornell University** **Ithaca, USA**
Postdoctoral Researcher *July 2022 - August 2022*
- **IBM T.J. Watson Research Center** **Yorktown Heights, USA**
Research Intern *June - Sept 2021*
 - Interned at the Hybrid Cloud division on IBM Research, working on KAR runtime, a system similar in concept to a serverless platform.
 - Added distributed transactional semantics within KAR to guarantee ACID properties. Implemented TPC-C benchmarking to evaluate the proposed transactional framework and to highlight the expansiveness of KAR.
- **Google** **Sunnyvale, USA**
Software Engineer Intern *June - Sept 2018, June - Sept 2019*
 - Interned twice with Quotaserver – a distributed bin counting - team. Identified the performance bottleneck of Quotaserver which was due to the backend database Spanner.
 - Optimized the performance of Quotaserver by sharding a single key in Spanner across different cells. The optimization resulted in 2.3x performance gain and 51% reduction in latency of the system.
- **Google** **New York, USA**
Software Engineer Intern *June 2017 - Sept 2017*
 - Worked in Google Ads group to introduce a new type of scheduler for Ads report generation that produced timely reports for customers across the globe.
 - The project went in production and impacted many ad publishers, redefining the existing client SLAs.
- **Citrix R&D** **Bangalore, India**
Software Developer *Feb 2014 - July 2016*

- Worked with Cloud Orchestration team in developing various features for NetScaler Control Center (NCC).
- Successfully drove the project of adding a new feature in NCC which enabled major CloudStack tenants to use load-balancing features of NetScaler. Have committed fixes into OpenSource Apache CloudStack.
- Extended admin partition feature of NetScaler to various cloud users eliminating hardware constraints required for loadbalancing.

Teaching Experience

Teaching Assistant

- CS171: **Introduction to Distributed Systems** *Spring 2020*
- CS130A: **Data Structures and Algorithms** *Fall 2018*
- CS171: **Introduction to Distributed Systems** *Spring 2018*
- CS130A: **Data Structures and Algorithms** *Winter 2018*
- CS32: **Object-Oriented Design and Implementation** *Fall 2017*
- CS16: **Problem Solving With Computers** *Spring 2017*
- CS56: **Advanced Application Development** *Winter 2017*

Guest Lectures

- **CS171: Introduction to Distributed Systems**
Distributed Atomic Commitment *April 2020*
- **CS271: Advanced Distributed Systems (Graduate level)**
Byzantine Agreement and Paxos *February 2020*
- **CS271: Advanced Distributed Systems (Graduate level)**
Replicated Dictionary and Leader Election *January 2020*
- **CS271: Advanced Distributed Systems (Graduate level)**
Database and Distributed Computing Fundamentals of Blockchain *February 2019*
- **CS130A: Data Structures and Algorithms**
Graph Traversal Algorithms *November 2018*
- **CS171: Introduction to Distributed Systems**
Replicated Dictionary and Log *April 2018*

Community Service

- **Volunteer Tutor** at School On Wheels *2020-present*
- **SB Hacks Judge** *2020 and 2021*
- Elected Graduate Representative: **Graduate Admissions Committee** *2019-2020*
- Elected Graduate Representative: **Faculty Recruitment Committee** *2018-2019*
- Elected Graduate Representative: **Talks&Events Committee** *2017-2018*

Professional Service

- Reviewer: SIGMOD 2023 Demo Track
- Reviewer: ICDE 2023
- Reviewer: SoCC 2022
- Reviewer: Future Generation Computer Systems Journal 2022
- Reviewer: DKE Journal 2021
- Sub-reviewer: FAB 2021
- Sub-reviewer: VLDB 2019
- Sub-reviewer: EDBT 2019
- Sub-reviewer: NETYS 2021, 2020, 2019, 2018

Invited Talks, Tutorials, and Posters

- **Salesforce** - *Managing Data Fault-tolerantly in Trusted and Untrusted Infrastructures* *November 2021*
- **MIT EECS Rising Stars** Poster - *Managing Data in Trusted and Untrusted Infrastructures*

October 2021

- **UC Berkeley** Talk - *Managing Data in Trusted and Untrusted Infrastructures*
August 2021
- **IBM Research** Talk - *Managing Data in Trusted and Untrusted Infrastructures*
August 2021
- **SBBD 2020** Tutorial - *Database and Distributed Computing Foundations of Blockchains*
September 2020
- **HPTS'19** Poster - *Atomic Commitment in Trusted and Untrusted Infrastructures*
November 2019
- **Google** Poster - *Atomic Commitment in Trusted and Untrusted Infrastructures*
July 2019
- **CRA-W Grad Cohort** Poster - *Unifying Consensus And Atomic Commitment*
April 2019
- **SoCal DB** Poster - *Fault-Tolerant Global Scale Data Management*
October 2018
- **VLDB 2018** Tutorial - *Database and Distributed Computing Foundations of Blockchains*
August 2018