
Research Statement

SUJAYA MAIYYA
UNIVERSITY OF CALIFORNIA SANTA BARBARA

I am a researcher and a second year PhD student in Computer Science. The topics I work on include, but are not limited to, Large Scale Data Management, Cloud Computing and Blockchain. I design, prototype and evaluate protocols for managing copious amounts of data with the aim of building practical large scale systems that can serve end users with minimal latency. My research also focuses on building fault tolerant systems that can handle multiple failure modes.

1 ONGOING RESEARCH

1.1 UNIFYING CONSENSUS AND ATOMIC COMMITMENT PROTOCOLS

One of my ongoing research projects is to unify *consensus* and *atomic commitment* protocols, which are seemingly disparate problems, into a single framework. The framework is developed to consolidate a myriad of transactional protocols developed in both the distributed systems and the database community dis-jointly. Cloud enterprises need a way to *replicate* their data on multiple servers to provide fault-tolerance and a way to *shard* the data across different servers to provide scalability. When transactions update data, all the involved shards and replicas need to be updated. Databases such as Spanner, update the shards independent of updating replicas, and this obliviousness incurs additional rounds of message communication for transaction commitment. In our work, we develop a new protocol that combines the efforts for committing a distributed transaction with the process of replicating the data. The amalgamation of the two efforts significantly reduces the latency required to commit transactions and thus, helps in processing more requests. This work is under revision for VLDB 2019.

1.2 HYBRID FAULT-TOLERANT PROTOCOLS

With the increase in small scale enterprises, there is increasing demand to rent servers from public cloud providers. Small enterprises usually own a few trusted servers in their compute fleet but would be insufficient to provide the required fault-tolerance or to cater to geographically distributed customers. Many state machine replication protocols have been developed that cater to either a fully trusted environment (eg. Lamport's Paxos) or a completely untrusted environment (eg. Liskov's Practical Byzantine Fault Tolerance, Nakamoto consensus). In a collaborative project, I helped develop a replication protocol that leverages the trust of a private cloud and the scalability of a public cloud. The protocol tolerates a bounded numbers

of crash and Byzantine failures, while providing the correctness and liveness required by a replication protocol. The use of a hybrid protocol also helps in reducing the number of servers to be rented from external providers. With an ever-growing need for data computation and storage, a practical solution such as ours will help small scale businesses in providing online service geographically with optimal costs. This work is currently under review.

2 FUTURE RESEARCH INTERESTS

2.1 BLOCKCHAIN FOR DATA MANAGEMENT IN A MALICIOUS SETTING

Having studied consensus in a closed setting, the advent of blockchain has opened a new realm of topics for consensus in an open setting where nodes can join and leave the network at will. Blockchain, which is a blend of distributed systems, databases and cryptography, has been a fascinating area of research to me. Having done an extensive literature survey in this area, I, along with my fellow researchers at UCSB, presented a tutorial on blockchain titled 'Database and Distributed Computing Fundamentals of Blockchain' at VLDB 2018. My study so far has focused on some of the most recent academic works, aiming to reduce the size of consensus group to a small subset of nodes that are participating in the blockchain. Although the tutorial was focused on permission-less blockchains, my inclination is towards blockchains in general.

Blockchain essentially provides a way to commit and verify transactions in a fully decentralized, un-trusted network of participants. This view poses many intriguing questions such as - Can blockchain be replaced by a transaction commitment protocol that tolerates byzantine faults? Can blockchain be used in a cloud setting to rent servers from an un-trusted provider at low cost but use blockchain to verify the transactions handled by un-trusted servers? I believe that blockchains, combined with varying failure modes, provide a vast scope of interesting and practically relevant problems and I intend to explore this space during my PhD career.