Xinwen Wang

https://www.cs.cornell.edu/~xinwen Email: xw467@cornell.edu

Education

Computer Science Ph.D.	Cornell University	Aug 2023
Computer Science B.S.	Stony Brook University	May 2017

Research Interests

Distributed Systems, Operating Systems, Resource Disaggregation.

Interns

Software Engineer Intern, Facebook June 2021
Designed and implemented a rate limiter for Shard Manager API Service.

Graduate Research Experience & Projects

Uniservice
Proposed Uniservice, a new programming paradigm for the disaggregated architecture. Uniservice employs the actor model and is specifically designed for a single type of hardware resource within a disaggregated architecture. Applications constructed using these uniservices can communicate via fast interconnects and a shared log. This approach simplifies operating systems while making better use of available hardware resources for the disaggregated architecture.

Ziplog
Ziplog is a one-size-fit-all shared log designed to provide all three of high throughput, low latency, and total order. This is made possible by a new `InsertAfter` API and protocols that separate ordering from storage, as well as by interleaving records among different shards. We implemented Ziplog using RDMA in C++.

BloomBox
Proposed BloomBox, a distributed failure detection protocol for geographic hash table. BloomBox utilizes an innovative data structure, the Mergeable Bloom Filter, to greatly reduce the bandwidth required for regenerating blocks in comparison to traditional heartbeat-based failure detection methods. Furthermore, by placing replicas in geographically diverse locations, BloomBox can offer significantly improved availability compared to heartbeat-based protocols.

Vegvisir IoT Blockchain
Vegvisir is a blockchain system designed for IoT devices that operate within intermittently connected networks. It features a tamper-proof log and a Conflict-free Replicated Data Type (CRDT) to ensure data integrity. I developed and implemented a reconciliation algorithm that leverages vector timestamps to enable fast and efficient data synchronization with minimal overhead.

Heterogeneous Consensus
Heterogeneous Consensus (Hetcons) is a consensus algorithm based on Leslie Lamport's Byzantine Paxos. Hetcons allows different learners to specify their own quorums of acceptors, different failure assumptions, and even mixed failure models. This allows two learners to achieve a consensus as long as there is an intersection between their failure assumptions and quorums.

Optimizing Communication in Distributed Machine Learning with RDMA
To overcome the network overhead introduced by distributed training, we developed a novel asynchronous zero-copy approach by using RDMA as the underlying communication channel. We implemented a proof-of-concept prototype on top of PyTorch.

User Mode Linux X-Container
X-Container is an architecture that leverages Xen paravirtualization (PV) to turn Linux kernel into a LibOS. I was working on porting X-Container to User Mode Linux instead of relying on Xen.

Undergraduate Research Experience & Projects

Student ***BetrFS at Stony Brook University*** **2016 – 2017**
The B^ε-tree File System, or BetrFS, is an in-kernel file system that uses B^ε trees to organize on-disk storage. I was working on testing the file system performance on Ramdisk.

Research Assistant ***ICLab at Stony Brook University*** **2015 –2016**
ICLab is a research platform to enable repeatable and representative studies of a broad class of online information controls, such as traffic differentiation, censorship and content modification. I developed a headless browser by using selenium, which is a web browser automation.

Teaching Assistant:

- CS4411 Operating Systems Practicum (Spring 2020, Spring 2023, Cornell)
- CS3410 Computer System Organization and Programming (Spring 2019, Cornell)
- CS5412 Cloud Computing (Spring 2018, Cornell)
- CS2110 OOP and Data Structure (Fall 2017, Cornell)

Certificates & Extracurricular Activities

Completing Linux Kernel Internals and Development Training	2018
Stony Brook Computer Science Honor Program	2016 – 2017
Academic Achievement Award	2014 – 2015
Dean List	2014 – 2016

Publications

- "Ziplog: A Totally Ordered Log combining Low Latency with Scalable Throughput." *Yu-Ju Huang, Shubham Chaudhary, **Xinwen Wang**, Cong Ding, Lorenzo Alvisi, Robbert van Renesse.* [In Submission]
- "Disaggregating Applications Using Uniservices." ***Xinwen Wang**, Robbert van Renesse.* [In Submission]
- "Charlotte: A Web of Composable Authenticated Distributed Data Structures." *Isaac Sheff, **Xinwen Wang**, Kushal Babel, Haobin Ni, Robbert van Renesse, Andrew C. Myers.* Transactions on Computer Systems (TOCS'23).
- "BloomBox: Improving Availability and Efficiency in Geographic Hash Tables." ***Xinwen Wang**, Robbert van Renesse*. ICDCS'22.
- "Disaggregated Applications Using Nanoservices." ***Xinwen Wang**, Robbert van Renesse*. WORDS'21.
- "Heterogeneous Consensus." *Isaac Sheff, **Xinwen Wang**, Robbert van Renesse, Andrew C. Myers.* OPODIS'20.

Posters

- “An IoT-friendly Blockchain for Coordination and Accountability.” *Danny Adams, Gloire Rubambiza, **Xinwen Wang**, Robbert van Renesse, and Hakim Weatherspoon.* Public Safety Broadband Stakeholder Meeting (PSCR), Chicago, IL, July 2019.