

Yiwen Zhang

(734) 834-4068 | yiwenzhg@umich.edu | 2260 Hayward St, Ann Arbor, MI 48109

SKILLS

C/C++, Python, Computer Networking Systems, RDMA, Quality-of-Service, Operating Systems, Linux Kernel

EDUCATION

University of Michigan	Ann Arbor, MI
Ph.D. Candidate, Computer Science and Engineering	2018 - Present
Masters, Electrical Engineering	April 2018
Bachelors, Electrical Engineering	December 2016

INDUSTRIAL EXPERIENCE

Research Intern, Microsoft	Seattle, Summer 2022
-----------------------------------	----------------------

Azure for Operators. Worked on ML system research for live ML queries.

Research Intern, Google	Sunnyvale, CA, 2020 – 2021 (9 months)
--------------------------------	---------------------------------------

Congestion Control Team, NetInfra. Worked on multi-tenancy support for applications in datacenters.

Software Engineering Intern, Google	Sunnyvale, CA, Summer 2019
--	----------------------------

Congestion Control Team, NetInfra. Worked on datacenter bandwidth allocation for ultra-low RPC latencies.

Performance Modeling Intern, ARM	Chandler, AZ, Summer 2017
---	---------------------------

Performance Modeling Group. Worked on modeling ARM's next-generation CPU.

RESEARCH EXPERIENCE

Research Assistant, Computer Science and Engineering, University of Michigan	2018 - Present
---	----------------

Advisor: Prof. Mosharaf Chowdhury

RDMA Performance Isolation

- Discovered and analyzed performance anomalies in various RDMA implementations across a wide range of RNICs.
- Designed a solution to provide predictable latencies and fair RNIC resource sharing via message-level shaping, initiator-based resource mediation, and passive system-wide latency monitoring.
- Paper accepted at NSDI '22.

Multi-Tenancy Support via Distributed QoS Management

- Analyzed how QoS usage impacts application performance isolation in Google datacenters.
- Designed a distributed QoS enforcement algorithm to achieve desired application performance.
- Paper accepted at SIGCOMM '22.

Automatic Query Planning for Live ML Analytics

- Lead the project on providing automatic query planning for live ML queries at Microsoft.
- Designed a solution to provide guaranteed performance with minimized resource consumption when deploying live ML queries across edge infrastructures.
- Paper accepted at NSDI '24.

TEACHING EXPERIENCE (AT UNIVERSITY OF MICHIGAN)

Graduate Student Instructor, EECS489: Computer Networks	Fall 2018
--	-----------

Wrote assignment auto-grader and led discussion.

Graduate Student Instructor, EECS582: Big Data Systems and Applications	Fall 2017
--	-----------

Designed programming assignments on big-data systems and wrote solutions

PUBLICATIONS

1. Vulcan: Automatic Query Planning for Live ML Analytics. *NSDI*' 24
Yiwen Zhang, Xumiao Zhang, Ganesh Ananthanarayanan, Anand Iyer, Yuanchao Shu, Victor Bahl, Z. Morley Mao, Mosharaf Chowdhury
2. Aequitas: Admission Control for Performance-Critical RPCs in Datacenters. *SIGCOMM*' 22
Yiwen Zhang, Gautam Kumar, Nandita Dukkupati, Xian Wu, Priyaranjan Jha, Mosharaf Chowdhury, Amin Vahdat
3. Justitia: Software Multi-Tenancy in Hardware Kernel-Bypass Networks. *NSDI*' 22
Yiwen Zhang, Yue Tan, Brent Stephens, Mosharaf Chowdhury
4. NetLock: Fast, Centralized Lock Management Using Programmable Switches. *SIGCOMM*' 20
Zhuolong Yu, **Yiwen Zhang**, Vladimir Braverman, Mosharaf Chowdhury, Xin Jin
5. Performance Isolation Anomalies in RDMA. *KBNets*' 17
Yiwen Zhang, Jucheng Gu, Youngmoon Lee, Mosharaf Chowdhury, Shin Kang
6. Efficient Memory Disaggregation with INFINISWAP. *NSDI*' 17
Jucheng Gu, Youngmoon Lee, **Yiwen Zhang**, Mosharaf Chowdhury, Shin Kang.

AWARDS

- **Ph.D. Student Fellowship**, Computer Science and Engineering, University of Michigan 2018 - 2019