

Haoran Qiu

PhD Candidate

Computer Science, University of Illinois Urbana-Champaign

CONTACT INFORMATION

240 Coordinated Science Lab
1308 W Main St., Urbana, IL 61801

Website: <https://haoran-qiu.com>
E-mail: haoranq4@illinois.edu

RESEARCH INTERESTS

Distributed Systems, Cloud Computing, ML for Systems, Systems for ML

EDUCATION

University of Illinois Urbana-Champaign

Ph.D. Candidate, Computer Science, May 2024 (expected)

- Thesis: Cloud Datacenter Management with Efficient and Robust Online Learning
- Advisor: [Prof. Ravishankar K. Iyer](#)
- Thesis Committee: Tamer Başar, Klara Nahrstedt, Indranil Gupta, Onur Mutlu

The University of Hong Kong

B.Eng. in Computer Science with *First Class Honours*, May 2019

- Thesis: Building Fast and Robust Blockchain Systems for AI Applications
- Advisor: [Prof. Heming Cui](#)

HONORS AND AWARDS

- ML and Systems Rising Stars, MLCommons, 2023
- Mavis Future Faculty Fellowship, UIUC, 2023–24
- UIUC CS PhD Fellowship, UIUC, 2023–24
- Best Presentation Award, Workshop on Composable Systems (IPDPS), 2022
- Yunni & Maxine Pao Memorial Fellowship, UIUC, 2021
- Conference Presentation Award, UIUC, 2020
- Travel Grants: USENIX OSDI/ATC 2023, IEEE DSN 2022, ACM SIGMETRICS 2021
- Best Undergraduate Thesis 2nd Runner-up, HKU, 2019
- Dean’s Honour List, HKU, 2016–19
- International Student Academic Excellence Award, University of Wisconsin-Madison, 2018
- Lee Shau Kee Scholarships for Student Enrichment, HKU, 2017
- Honorable Mention in Mathematical Contest In Modeling, COMAP, 2017
- HKU Foundation Scholarships For Outstanding Students, 2015–19

PROFESSIONAL EXPERIENCE

PhD Software Engineer Intern, **Google SRG**, Mountain View, CA. May 2023 – Aug. 2023
Managers: Dr. Phitchaya Mangpo Phothilimthana (DeepMind), Dr. Yawen Wang
Designed and implemented power-aware datacenter capacity management with coordinated inter-cluster scheduling and intra-cluster bin-packing.

Visiting Researcher, **IBM Research**, Yorktown Heights, NY. Sept. 2022 – Dec. 2022
Managers: Dr. Chen Wang, Dr. Hubertus Franke
*Designed and implemented AWARE (published at **ATC 2023 [C9]**), an extensible framework that provides robust and reliable deployment of online learning agents in production cloud platforms.*

PhD Software Engineer Intern, **Google**, Sunnyvale, CA. May 2022 – Aug. 2022
Managers: Dr. Qiaobin Fu, Dr. Joon Kang
*Designed and implemented network-aware scheduling of best-effort jobs for Borg cluster manager by extending FIRM (published at **OSDI 2020 [C7]**).*

Research Intern, **Microsoft Research**, Redmond, WA (Virtual). May 2021 – Aug. 2021
 Managers: Dr. Pulkit Misra, Dr. Íñigo Goiri
*Developed an analytical model to predict performance improvement for workloads with overclocking and an overclocking and power management framework (Under review at **ASPLOS 2024 [P3]**).*

PhD Software Engineer Intern, **Google**, Sunnyvale, CA (Virtual). May 2020 – Aug. 2020
 Managers: Bharath Krishna, Yili Ding
Developed configuration correctness review pipeline for load balancers.

Software Engineer Intern, **Credit Suisse**, Hong Kong. May 2018 – Aug. 2018
 Manager: Sandeep Rakhra
Developed a data processing pipeline for risk information storage and analysis.

Software Engineer Intern, **Hututa Technologies**, Hong Kong. May 2017 – Aug. 2017
 Managers: Dr. Zhiqiang Ma, Dr. Lin Gu
Developed a distributed data processing system to support biological gene mapping algorithms with networking monitoring for fault tolerance.

CONFERENCE &
 JOURNAL
 PUBLICATIONS

[C9] **Haoran Qiu**, Weichao Mao, Chen Wang, Alaa Youssef, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, Ravishankar K. Iyer. *AWARE: Automating Workload Autoscaling with Reinforcement Learning in Production Cloud Systems*, **USENIX ATC 2023**
 🏆 **Artifact Available, Functional, Reproduced**

[C8] **Haoran Qiu**, Weichao Mao, Archit Patke, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, Ravishankar K. Iyer. *SIMPPO: A Scalable Online Learning Framework for Multi-tenant Serverless Resource Management*, **ACM SoCC 2022**

[C7] **Haoran Qiu**, Subho S. Banerjee, Saurabh Jha, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *FIRM: An Intelligent Fine-grained Resource Management Framework for SLO-oriented Microservices*, **USENIX OSDI 2020**
 🏆 **Artifact Available, Functional, Reproduced**

[C6] Weichao Mao, **Haoran Qiu**, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer, Tamer Başar. *Multi-Agent Meta-Reinforcement Learning: Sharper Convergence Rates with Task Similarity*, **NeurIPS 2023**

[C5] Weichao Mao, **Haoran Qiu**, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer, Tamer Başar. *A Mean-Field Game Approach to Cloud Resource Management with Function Approximation*, **NeurIPS 2022**

[C4] Archit Patke, Saurabh Jha, **Haoran Qiu**, Jim Brandt, Ann Gentile, Joe Greenseid, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *Delay Sensitivity-driven Congestion Mitigation for HPC Systems*, **ACM ICS 2021**

[C3] **Haoran Qiu**, Tao Ji, Shixiong Zhao, Xusheng Chen, Ji Qi, Heming Cui, Sen Wang. *A Geography-Based P2P Overlay Network for Fast and Robust Blockchain Systems*, **IEEE Transactions on Services Computing**
 🏆 **Best Undergraduate Thesis 2nd Runner-up, HKU, 2019**

[C2] Shixiong Zhao, Rui Gu, **Haoran Qiu**, Tsz On Li, Yuexuan Wang, Heming Cui, Junfeng Yang. *OWL: Understanding and Detecting Concurrency Attacks*, **IEEE DSN 2018**

[C1] Cheng Wang, Xusheng Chen, **Haoran Qiu**, Boxuan Li, Weiwei Jia, Shixiong Zhao, Heming Cui. *PLOVER: Fast, Multi-core Scalable Virtual Machine Fault-tolerance*, **USENIX NSDI 2018**

WORKSHOP
PUBLICATIONS

[W7] **Haoran Qiu**, Weichao Mao, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, Ravishankar K. Iyer. *On the Promise and Challenges of Foundation Models for Learning-based Cloud Systems Management*, **ML for Systems Workshop @NeurIPS 2023**
🏆 **Selected for Spotlight Presentation**

[W6] **Haoran Qiu**, Linghao Zhang, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *PARM: Adaptive Resource Allocation for Datacenter Power Capping*, **ML for Systems Workshop @NeurIPS 2023**

[W5] **Haoran Qiu**, Weichao Mao, Archit Patke, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, Ravishankar K. Iyer. *Reinforcement Learning for Resource Management in Multi-tenant Serverless Platforms*, **EuroMLSys @EuroSys 2022**

[W4] **Haoran Qiu**, Saurabh Jha, Subho S. Banerjee, Archit Patke, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *Is Function-as-a-Service a Good Fit for Latency-critical Services?*, **WoSC @Middleware 2021**

[W3] Archit Patke, **Haoran Qiu**, Saurabh Jha, Srikumar Venugopal, Michele Gazzetti, Christian Pinto, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *Evaluating Hardware Memory Disaggregation Under Delay and Contention*, **COMPSYS @IPDPS 2022**
🏆 **Best Presentation Award**

[W2] Harshitha Sreejith, Saurabh Jha, **Haoran Qiu**, Larisa Shwartz, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *I-SPOT: Supporting Resilient Cloud Computing Using Counterfactual Reasoning*, **IBM-IEEE CAS/EDS AICS 2022**

[W1] **Haoran Qiu**, Saurabh Jha, Subho S. Banerjee, Archit Patke, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *Automated Distributed System Design and Management with Inflight Analytics*, **IBM-IEEE CAS/EDS AICS 2021**

SELECTED
PRE-PRINTS

[P4] **Haoran Qiu**, Weichao Mao, Archit Patke, Shengkun Cui, Chen Wang, Hubertus Franke, Zbigniew T. Kalbarczyk, Tamer Başar, Ravishankar K. Iyer. *FLASH: Towards Fast Model Adaptation with Meta-learning in ML-Centric Cloud Platforms* (Under Review).

[P3] Jovan Stojkovic, Pulkit Misra, Íñigo Goiri, Sam Whitlock, Esha Choukse, Mayukh Das, Chetan Bansal, Jason Lee, **Haoran Qiu**, Reed Zimmermann, Savyasachi Samal, Brijesh Warriar, Ricardo Bianchini. *SmartOClock: Workload- and Risk-Aware Overclocking in the Cloud* (Under Review).

[P2] Archit Patke, Christian Pinto, **Haoran Qiu**, Shengkun Cui, Saurabh Jha, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *INDIGO: Intelligent Page Migration for Hardware Memory Disaggregation Systems* (Under Review).

[P1] **Haoran Qiu**, Yongzhou Chen, Tianyin Xu, Zbigniew T. Kalbarczyk, Ravishankar K. Iyer. *SLO beyond the Hardware Isolation Limits*, arxiv.org/abs/2109.11666.

INVITED TALKS

Efficient and Robust Online Learning in Cloud Datacenters

- ML for Systems Seminar, Google Brain
- Huawei Cloud Research, Huawei Technologies
- Peisu Xia Young Scholars Forum 2022, Chinese Academy of Sciences
- MLSys Rising Star Workshop 2023

Towards a Foundation Model for Cloud Systems Management

- ML for Systems Workshop at NeurIPS 2023

Sustainable Workload Autoscaling with Predictive AI

- KubeCon/CloudNativeCon North America 2023, CNCF
- ML for Systems Workshop at NeurIPS 2023

Automate Workload Autoscaling in Production Cloud Systems with Meta-Learning

- USENIX ATC 2023
- AI Cloud Platform, IBM Research
- Hybrid Cloud and AI Workshop, IBM-Illinois Discovery Accelerator Institute
- UIUC CS Systems and Networking Spring Retreat, 2023

Scalable and Adaptive Online Learning for Serverless Resource Management

- ACM SoCC 2022
- EuroMLSys 2022 (Co-located with ACM EuroSys 2022)
- Hybrid Cloud and AI Workshop, IBM-Illinois Discovery Accelerator Institute

Intelligent Resource Management for SLO-oriented Cloud Microservices

- USENIX OSDI 2020
- AI Cloud Platform & Red Hat, IBM Research
- CSL Student Conference, 2020
- Sustainable Hybrid Cloud Platform, IBM Research
- Turbonomics AIOps Research, IBM Research
- Applied AI Cloud, Infosys
- 2021 IBM IEEE CAS/EDS AI Compute Symposium

TEACHING EXPERIENCE

CS 598 ML and Data Systems , Teaching Assistant, UIUC	Fall 2023
ECE 598 Dependable AI Systems , Course Assistant, UIUC	Fall 2023
ECE 471 Data Science Analytics , Course Assistant, UIUC	Fall 2022
CS 536 Design of Fault-Tolerant Digital Systems , Course Assistant, UIUC	Fall 2021
COMP 2396 Object-oriented Programming , Teaching Assistant, HKU	Fall 2017

PROFESSIONAL SERVICES

Program Committee / Reviewer

- ML for Systems Workshop at NeurIPS 2023
- DSN 2023 Doctoral Forum
- ACM Transactions on Architecture and Code Optimization, 2023
- IEEE Internet of Things Journal, 2022
- Machine Learning Session, UIUC CSL Student Conference 2021

External Program Committee / Reviewer

- EuroSys 2024
- EuroSys 2022

Artifact Evaluation Committee

- 2023: SOSP, EuroSys, MLSys
- 2022: OSDI, ATC

Community Services

- Mentor, Promoting Undergraduate Research in Engineering (PURE), UIUC, 2023
- Mentor, Illinois Science and Technology Coalition (ISTC), 2023–24
- Member, Institute for Inclusion, Diversity, Equity and Access (IDEA), UIUC, 2023
- Mentor, Undergraduate Research Experience (URE) at IIDAI, 2022–23
- CS Ambassador, UIUC CS Visit Day for Prospective Graduate Students, UIUC, 2020
- Undergraduate Ambassador, HKU, 2016

REFERENCES

References available upon request.