Zixuan Wang

Ph.D. Candidate
University of California, San Diego

www.netadmin.ai
zxwang42@gmail.com
619-356-0954

EDUCATION

University of California, San Diego
Ph.D. candidate in Computer Science.

San Diego, CA, US
Sep. 2018 – Present

Zhejiang University
BS in Computer Science.

Hangzhou, China
Sep. 2014 – July. 2018

INTEREST

My *research* interests are in computer architecture and systems. More specifically, I am interested in cross-stack and technology-driven innovations that enhance the scalability and security of computer systems for emerging applications. My *industrial* efforts focus on deploying the emerging confidential virtual machine technology in real-world systems. And my *open-source* projects facilitate research, industry, and personal usages.

EXPERIENCE

Research Experience

Graduate Research Assistant, STABLE Lab
Advisor: Jishen Zhao, Steven Swanson

UC San Diego
Sep. 2018 – Present

- **Emerging Architecture:**
  * Characterizing and simulating emerging main memory systems [3] [5] [PU4].
  * Reverse engineering and side channel attacks in main memory systems [1].
- **System Integration:**
  * Building scalable distributed AI training infrastructure based on CXL, an emerging memory interconnection protocol [2].
  * Reverse-engineering and attacking CXL-enabled systems [PP1].
  * Developing scalable system support for CXL-based heterogeneous systems [PP2].
- **Emerging Application and Programming Techniques:**
  * Developing generative AI that automatically re-writes legacy code to leverage emerging memory systems [4].
  * Investigating system supprot for autonomous vehicle systems [PU2] [PU3].
  * Characterizing performance of serverless systems based on WebAssembly [PU1].

Research Intern, SOLAB
Mentors: Joonseop Sim, Euicheol Lim

SK Hynix USA
Jun. 2019 – Sep. 2019

- **Emerging Memory:** One of the first performance evaluations of CXL, an emerging memory interconnection protocol.
- **ML Training Acceleration:** Efficient distributed infrastructure to train ML models using CXL [2].

Undergraduate Research Assistant, Computer Architecture Lab
Advisors: Qingsong Shi, Wenzhi Chen

Zhejiang University
Sep. 2015 – Jun. 2018

- **Developed a Full Computer System from Scratch:** Implemented a CPU (with peripherals) on FPGA, a fully functional operating system kernel in C and assembly, and integrated the kernel to run on this CPU.
- **Developed new Undergrad Courses:** Developed two new courses that guide undergrads to develop their own operating systems running on their own CPU.

Industry Experience

Software Engineering Intern
Enhanced cloud user data confidentiality with emerging AMD SEV-SNP SVSM.

Confidential VM, Google Cloud
Jun. 2023 – Sep. 2023

Part-Time Student Researcher
Deployed the confidential VM platform at scale.

Network Infra, Meta
Sep. 2022 – Jan. 2023

Software Engineering Intern
Initiated and developed Meta's first confidential VM platform.

Network Infra, Meta
Jun. 2022 – Sep. 2022

Software Engineering Intern
Modernizing Linux KVM testing with UEFI and AMD SEV confidential VM supports.

Confidential VM, Google Cloud
Jun. 2021 – Sep. 2021

PUBLICATIONS

In Progress & Under Submission

[PP1] *Zixuan Wang*, Milad Esrafilian, Daniel Moghimi, Jishen Zhao, Mohammadkazem Taram. CXLeak: Architectural Attacks via Practical CXL Systems

[PP2] *Zixuan Wang*, Jishen Zhao. Fork is All You Needed in the Era of Heterogeneous Computing

Peer Reviewed

[1] *Zixuan Wang*, Mohammadkazem Taram, Daniel Moghimi, Steven Swanson, Dean Tullsen, Jishen Zhao. NVLeak: Off-Chip Side-Channel Attacks via Non-Volatile Memory Systems, *USENIX Security, 2023*

[2] *Zixuan Wang*, Joonseop Sim, Euicheol Lim, Jishen Zhao. Enabling Efficient Large-Scale Deep Learning Training with Cache Coherent Disaggregated Memory Systems, *HPCA, 2022*

[3] *Zixuan Wang*, Xiao Liu, Jian Yang, Theodore Michailidis, Steven Swanson, Jishen Zhao. Characterizing and Modeling Non-Volatile Memory Systems, *IEEE Micro Top Picks, 2021*

[4] Hanxian Huang, *Zixuan Wang*, Juno Kim, Steven Swanson, Jishen Zhao. Ayudante: A Deep Reinforcement Learning Approach to Assist Persistent Memory Programming, *USENIX ATC, 2021*

[5] *Zixuan Wang*, Xiao Liu, Jian Yang, Theodore Michailidis, Steven Swanson, Jishen Zhao. Characterizing and Modeling Non-Volatile Memory Systems, *MICRO, 2020*

Technical Reports

[PU1] Jamshed Ashurov, *Zixuan Wang*, Jishen Zhao. Characterizing WebAssembly Performance in the Era of Serverless Computing, *ISSTA SRC, 2023*

[PU2] Haolan Liu, *Zixuan Wang*, Jishen Zhao. COLA: Characterizing and Optimizing the Tail Latency for Safe Level-4 Autonomous Vehicle Systems, *ArXiV, 2023*

[PU3] Maximilian Apodaca, Shengye Wang, *Zixuan Wang*, Jishen Zhao. Enabling Fast Recovery for Autonomous Vehicle Systems with Linux Container Checkpointing, *SOSP SRC, 2021*

[PU4] Joseph Izraelevitz, Jian Yang, Lu Zhang, Juno Kim, Xiao Liu, Amirsaman Memaripour,Yun Joon Soh, *Zixuan Wang*, Yi Xu, Subramanya R. Dulloor, Jishen Zhao, Steven Swanson. Basic Performance Measurements of the Intel Optane DC Persistent Memory Module, *ArXiv, 2019*

[PU5] *Zixuan Wang*, Xiao Liu, Jongryool Kim, Hokyoon Lee, Jishen Zhao. Reliable and Flexible Large Scale Memory Network, *NVMW, 2019*

SERVICES

Co-Founder and Organizing Committee   Students@Systems
I'm one of the founders and organizers of Students@Systems: www.students-at-systems.org   *Jan. 2022 – Present*
○ I have hosted multiple online panel discussions on academic job hunting (2022 June, 2023 Oct) and artifact reproducibility (2023 Apr).
○ I helped with organizing more than ten online events, including panels on applying for PhD, and interviews with researchers from underrepresented groups.

Submission Chair   MICRO 2021
I served as a submission chair for the MICRO 2021 conference.   *Mar. 2021 – Jun. 2021*
○ I have developed MightyPC, a recommendation system to match submissions with reviewers.
○ MightyPC has then been used by: MICRO'21, IEEE MICRO TopPicks'22, HPCA'22, MICRO'22, DSN'23, and more.

MENTORSHIPS

Jamshed Ashurov (Undergrad → Master)   UC San Diego
WebAssembly system interface characterization, published on ISSTA'23 SRC.   *2022 – Present*

Haolan Liu (PhD Student)   UC San Diego
Characterizing autonomous vehicle system, under submission.   *2022 – Present*

Maximilian Apodaca (Undergrad → Tesla)   UC San Diego
Container checkpointing, published on SOSP'23 SRC.   *2020 – 2021*

Hanxian Huang (PhD Student)   UC San Diego
Generative AI for programming, published on USENIX ATC'21.   *2020 – 2021*

TEACHING

Teaching Assistant: Introduction to Computer Architecture   University of California, San Diego
Undergrad level computer arch course.   *Jan. 2022 – Mar. 2022*

Associate Instructor: Hardware-Based Computer System Design   Zhejiang University
Developed and instructed a new course that guides students to develop their own CPU to run their OS.   *Mar. 2018 – Jun. 2018*

Associate Instructor: Operating System Course   Zhejiang University
Developed and instructed a new course that guides students to develop their own OS from scratch.   *Sep. 2017 – Feb. 2018*

INVITED TALKS

NVLeak: Off-Chip Side-Channel Attacks via Non-Volatile Memory Systems
NVMW'23, PRISM Center at Semiconductor Research Corporation

Enabling Efficient Large-Scale Deep Learning Training with Cache Coherent Disaggregated Memory Systems
Intel Co., IBM Research, SK hynix Inc., Micron Inc., Alibaba Cloud USA Inc., Foundational Microarchitecture Research (FoMR), CRISP Center at Semiconductor Research Corporation

Characterizing and Modeling Non-Volatile Memory Systems
TECHCON'20, NVMW'21, Foundational Microarchitecture Research (FoMR), CRISP Center at Semiconductor Research Corporation

Trust but Verify: Co-Locating Hypervisor Services with User Code via AMD SEV-SNP SVSM
Google Cloud'23

Securing User Data with Confidential Virtual Machine
Meta Annual Security Summit'22

Modernizing KVM-Unit-Tests with UEFI and AMD Confidential Virtual Machine
Google Cloud'21, AMD'21

HONORS & AWARDS

MICRO PhD Forum Attendee: Selected as one of the presenters on the PhD Forum, 2023 MICRO
Google Peer Bonus: Awarded one peer bonus recognizing the impact of my project, 2023 Google
NVMW Memorable Paper Finalist: Awarded to one of the most impactful paper in persistent memory research, 2023 NVMW
Meta Security Highlight: Highlight presentation on Meta's annual security summit, 2022 Meta
Google Peer Bonus: Awarded two peer bonuses recognizing the impact of my project, 2021 Google
IEEE Micro TopPicks: Annually awarded to 12 best papers in computer architecture area, 2021 IEEE
NVMW Memorable Paper Finalist: Awarded to one of the most impactful paper in persistent memory research, 2021 NVMW
Outstanding Dissertation: Outstanding undergraduate dissertation, 2018 Zhejiang University
He-Zhi-Jun Scholarship: Top 10 outstanding students of the computer science department, 2017 Zhejiang University
Outstanding Prize: Challenge Cup, National Undergraduate Academic Science and Technology Works Competition, 2017 China
Rising Star in Academic: Top 1% of computer science students in academic achievements, 2017 Zhejiang University
Academic Scholarship: Top 10% students of the computer science department
Second Prize: Digilent Design Contest, 2017 China
Third Prize: Advanced Computer Architecture Undergraduate Innovation Competition, 2016 CCF China

INDUSTRY PROJECTS

Trusted Execution of Hypervisor Code within Guest VM   *June, 2023*
Initiated the AMD SEV-SNP SVSM support to enhance Google Cloud's confidential virtual machines.
○ I built the initial SVSM support in Google Cloud's Linux kernel, hypervisor, guest firmware, and guest kernel.

Confidential Virtual Machine Platform   *June, 2022*
Initiated and developed the first confidential VM platform at Meta, highlighted at Meta's Annual Security Summit.
○ I built and deployed the software and operating system support for the first CVM platform at Meta.
○ The project is highlighted at Meta's Annual Security Summit.

Modernizing Linux KVM Testing Infrastructure with Confidential VM   *June, 2021*
Implement the first UEFI and AMD SEV/SEV-ES support in KVM-Unit-Tests, patches merged to upstream Linux KVM.
○ It serves as a solid foundation for the future development of trusted execution in KVM.
○ 19 patches have been merged in upstream Linux KVM, now used by all cloud companies.

REFERENCES

Jishen Zhao   Associate Professor, UC San Diego
Steven Swanson   Professor, UC San Diego
Dean Tullsen   Professor, UC San Diego
Yuan Xie   Chair Professor, HKUST