Wei **Bai**

Member of Technical Staff at Microsoft AI

Redmond, Washington, USA

baiwei0427@gmail.com | baiwei0427.github.io | Google Scholar

About

Networking researcher and engineer with 10+ years across academia and industry, spanning the full stack of AI/data center networking — from GPU collective communication and RDMA transport to NIC and switch architecture. I have co-architected RDMA transports running on hundreds of thousands of GPUs in production (OpenAI, Microsoft), helped build and operate one of the world's largest RoCE deployments (spanning up to 100 km and carrying over 70% of Azure traffic), published at top-tier academic conferences such as SIGCOMM and NSDI, and contributed to open-source projects and open industry standards.

Work Experience

Microsoft *Redmond, USA*

Member of Technical Staff, Microsoft AI *Feb. 2026 – Present*

- Designing networking, RDMA transport, and collective communication mechanisms for large-scale AI clusters.

NVIDIA *Redmond, USA*

Principal Software Research Architect, Networking Software and System Architecture Group *Jan. 2024 – Feb. 2026*

- **Spectrum-X:** Contributed to Spectrum-X, NVIDIA's Ethernet networking platform for AI, focusing on long-distance DC-to-DC and region-to-region networking scenarios.
- **MRC:** Served as a core architect of Multipath Reliable Connection (MRC), a next-generation RDMA transport protocol that distributes a single RDMA connection across multiple network paths. Brought MRC to production across hundreds of thousands of GPUs at OpenAI and Microsoft, and provided hands-on deployment support. MRC has since been published as an Open Compute Project (OCP) open specification.

Microsoft *Redmond, USA & Beijing, China*

Senior Researcher, Microsoft Research Redmond (MSR), Aug. 2019 – Jan. 2024

Researcher, Microsoft Research Asia (MSRA), Jul. 2017 – Aug. 2019 *Jul. 2017 – Jan. 2024*

- **RDMA for Azure:** Core contributor to one of the world's largest RoCE (RDMA over Converged Ethernet) deployments at Microsoft Azure, spanning data centers up to 100 km apart and carrying over 70% of Azure traffic. Worked across four generations of RDMA NICs and more than ten switch ASICs from Broadcom, Cisco, and NVIDIA, covering congestion control, switch buffer optimization, and end-to-end NIC/switch qualification. The resulting RoCE network underpins Azure Storage and enabled the inter-datacenter training of GPT-4.5.
- **SONiC:** Contributed to SONiC (Software for Open Networking in the Cloud), an open-source network operating system now deployed across major cloud and enterprise networks. Delivered disaggregated chassis support, QoS/RDMA features, and a programmable-hardware-based test system using platforms such as Keysight.
- **Next-Generation Cloud and AI Networking Research:** Led research on the next generation of RDMA and AI networking, covering the full spectrum from transport protocol design and NIC testing to switch buffer management and RDMA performance isolation. Uncovered multiple previously unknown bugs in production RDMA NICs through systematic testing. This work directly influenced the design of next-generation NIC and switch architectures.

Education

The Hong Kong University of Science and Technology (HKUST) *Hong Kong SAR, China*

Ph.D. in Computer Science and Engineering *2013 – 2017*

- Advisor: Prof. Kai Chen
- Thesis: Congestion Control Mechanisms for Data Center Networks
- Microsoft Research Asia Fellowship (2015)

Shanghai Jiao Tong University (SJTU) *Shanghai, China*

B.E. in Information Security *2009 – 2013*

Selected Publications

Zhuolong Yu, Bowen Su, **Wei Bai**, Shachar Raindel, Vladimir Braverman, Xin Jin, "Understanding the Micro-Behaviors of Hardware Offloaded Network Stacks with Lumina", in *Proceedings of the ACM SIGCOMM 2023 Conference (SIGCOMM)*, New York City, New York, September 10-14, 2023.

Wei Bai, Shanim Sainul Abdeen, Ankit Agrawal, Krishan Kumar Attre, Paramvir Bahl, Ameya Bhagat, Gowri Bhaskara, Tanya Brokhman, Lei Cao, Ahmad Cheema, Rebecca Chow, Jeff Cohen, Mahmoud Elhaddad, Vivek Ette, Igal Figlin, Daniel Firestone, Mathew George, Ilya German, Lakhmeet Ghai, Eric Green, Albert Greenberg, Manish Gupta, Randy Haagens, Matthew Hendel, Ridwan Howlader, Neetha John, Julia Johnstone, Tom Jolly, Greg Kramer, David Kruse, Ankit Kumar, Erica Lan, Ivan Lee, Avi Levy, Marina Lipshteyn, Xin Liu, Chen Liu, Guohan Lu, Yuemin Lu, Xiakun Lu, Vadim Makhervaks, Ulad Malashanka, David A. Maltz, Ilias Marinos, Rohan Mehta, Sharda Murthi, Anup Namdhari, Aaron Ogus, Jitendra Padhye, Madhav Pandya, Douglas Phillips, Adrian Power, Suraj Puri, Shachar Raindel, Jordan Rhee, Anthony Russo, Maneesh Sah, Ali Sheriff, Chris Sparacino, Ashutosh Srivastava, Weixiang Sun, Nick Swanson, Fuhou Tian, Lukasz Tomczyk, Vamsi Vadlamuri, Alec Wolman, Ying Xie, Joyce Yom, Lihua Yuan, Yanzhao Zhang, Brian Zill, "Empowering Azure Storage with RDMA", in *Proceedings of the 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Boston, Massachusetts, April 17-19, 2023 (Operational Systems Track).

Xinhao Kong, Jingrong Chen, **Wei Bai**, Yechen Xu, Mahmoud Elhaddad, Shachar Raindel, Jitendra Padhye, Alvin R. Lebeck, Danyang Zhuo, "Understanding RDMA Microarchitecture Resources for Performance Isolation", in *Proceedings of the 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Boston, Massachusetts, April 17-19, 2023.