JINGHAN HUANG

University of Illinois Urbana-Champaign
jinghan347@gmail.com ◇ jinghanhuang.github.io

RESEARCH INTEREST

- Co-design of hardware and software for network system
- System for machine learning training and inference
- Intra-host high-performance I/O

EDUCATION

University of Illinois Urbana-Champaign *August 2020 - Present*
PhD Candidate, Electrical & Computer Engineering GPA: 3.97/4.0
- *Advisor: Prof. Nam Sung Kim*

University of Illinois Urbana-Champaign *September 2016 - May 2020*
Bachelor of Science, Computer Engineering GPA: 3.96/4.0
- Graduated with Highest Honors

Zhejiang University *September 2016 - June 2020*
Bachelor of Engineering, Electronic & Computer Engineering GPA: 3.93/4.0
- The Scholarship of Zhejiang Province

WORK EXPERIENCE

IBM Research *May 2022 - August 2022*
Visiting Scholar at Yorktown Heights, New York *Mentor: Eun Kyung Lee*
- IBM Extern through UIUC Research Assistant

SKILLS AND TECHNIQUES

Programming Languages: C/C++, Python, Verilog/SystemVerilog, HLS, Shell script, C#, Go
Development Skills: Linux, DPDK, RDMA, FPGA, SmartNIC, SmartSSD, CXL, CUDA, PyTorch

PUBLICATION

- [**ISCA'25**] LIA: A Single-GPU LLM Inference Acceleration with Cooperative AMX-Enabled CPU-GPU Computation and CXL Offloading
  H. Kim, N. Wang, Q. Xia, J. Huang, A. Yazdanbakhsh, N. S. Kim
- [**ISCA'25**] UPP: Universal Predicate Pushdown to Smart Storage
  I. Jeong, J. Huang, C. Hu, D. Park, J. Kang, N. S. Kim, Y. Park
- [**MICRO'24**] Demystifying a CXL Type-2 Device: A Heterogeneous Cooperative Computing Perspective
  H. Ji, S. Vanavasam, Y. Zhou, Q. Xia, J. Huang, Y. Yuan, R. Wang, P. Gupta, B. Chitlur, I. Jeong, N. S. Kim
- [**ISCA'24**] HAL: Hardware-assisted Load Balancing for Energy-efficient SNIC-Host Cooperative Computing
  J. Huang, J. Lou, S. Vanavasam, X. Kong, H. Ji, I. Jeong, E. K. Lee, D. Zhuo, N. S. Kim

- [**ASPLOS'24**] TAROT: A CXL SmartNIC-Based Defense Against Multi-bit Errors by Row Hammer Attacks
  C. Song, M. J. Kim, T. Wang, H. Ji, J. Huang, I. Jeong, J. Park, H. Nam, M. Wi, J. H. Ahn, N. S. Kim
- [**NSDI'24**] Harmonic: Hardware-assisted RDMA Performance Isolation for Public Clouds
  J. Lou*, X. Kong*, J. Huang, W. Bai, N. S. Kim, D. Zhuo
- [**MICRO'23**] Demystifying CXL Memory with True CXL-Ready Systems and CXL Memory Devices
  Y. Sun, Y. Yuan, Z. Yu, R. Kuper, C. Song, J. Huang, H. Ji, S. Agarwal, J. Lou, I. Jeong, R. Wang, J. H. Ahn, T. Xu, N. S. Kim
- [**IISWC'23**] [**Best Paper Runner-Up Award**] Making Sense of Using a SmartNIC to Reduce Datacenter Tax from SLO and TCO Perspectives
  J. Huang, J. Lou, Y. Sun, T. Wang, E. K. Lee, N. S. Kim
- [**ATC'23**] STYX: Exploiting SmartNIC Capability to Reduce Datacenter Memory Tax
  H. Ji, M. Mansi*, Y. Sun*, Y. Yuan, J. Huang, R. Kuper, M. M. Swift, N. S. Kim
- [**HPCA'23**] RAMBDA: RDMA-driven Acceleration Framework for Memory-intensive us-scale Datacenter Applications
  Y. Yuan, J. Huang, Y. Sun, T. Wang, J. Nelson, D. Ports, Y. Wang, R. Wang, C. Tai, N. S. Kim

RESARCH EXPERIENCE

System for Machine Learning *2023 - Present*
We take advantage of SmartNICs and other ML accelerators to accelerate communication and computation in machine learning. The results have been published in **ISCA'25**.

High-performance and Energy-efficient Network System *2021 - Present*
We design hardware-assisted SmartNIC and host cooperative computing solutions, RDMA isolation for public clouds, and efficient frameworks to utilize accelerators in remote servers. The results have been published in **HPCA'23**, **IISWC'23**, **NSDI'24**, **ISCA'24**.

Intra-host Computing Offload *2022 - Present*
We take advantage of SmartNIC and SmartSSD to offload operations inside the host, such as kernel functions (KSM, zswap), memory protection against row hammer attacks, and database filtering. The results have been published in **ATC'23**, **ASPLOS'24**, **ISCA'25**.

CXL Devices Exploration *2022 - 2023*
CXL is used to connect CPUs with I/O devices, featuring memory expansion and cache coherence. We explore CXL memory devices and CXL accelerators based on real hardware. The results have been published in **MICRO'23**, **MICRO'24**.

TEACHING ASSISTANT EXPERIENCE

CS/ECE 374: Introduction to Algorithms & Models of Computation (FA20 - FA21, FA22)
ECE 385: Digital Systems Laboratory (SP20, SP25)
ECE 220: Computing Systems and Programming (FA19)
CS 225: Data Structures (SP19)

PROFESSIONAL SERVICES AND ACTIVITIES

Artifact Evaluation Committee: ASPLOS 2025, MICRO 2024, SIGCOMM 2024, ATC 2024, OSDI 2024