

# Junxue ZHANG

Research Assistant Professor, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong  
Email: [zjx@ust.hk](mailto:zjx@ust.hk) URL: <https://snowzjx.me> Office: Room 2542, Academic Building

## RESEARCH INTERESTS

My research interests are building **High-performance** (C3, C6, C12, C20 and C21; J6 and J8), **Intelligent** (C4, C13, C14 and C16; J5 and J7) and **Secure** (C5, C7 to C11, C15 and C17; J3) datacenter networking, with a particular focus on **AI-Centric Transports (AICT)** now. Please see my [research statement](#) for more details.

## ACADEMIC POSITION

School of Computer Science and Technology, University of Science and Technology of China  
*Professor<sup>a</sup>, Ph.D. Supervisor* Tentatively from 2025 Spring

Department of Computer Science and Engineering, Hong Kong University of Science and Technology  
*Research Assistant Professor* Nov. 2023 – Now

<sup>a</sup>Specially appointed as a tenured associate professor with the title of full professor.

## EDUCATION

*Ph.D., Hong Kong University of Science and Technology* Sept. 2016 – Aug. 2022

**Thesis:** Towards Efficient Transports for Datacenter Networking with High Environmental Variations

**Advisor:** Prof. Kai Chen

**Award:** Honorable Mention, HKUST CSE Best Ph.D. Dissertation Award

*MS.c., Southeast University*

Sept. 2013 – Jul. 2016

**Advisor:** Prof. Junzhou Luo

*BS.c., Southeast University*

Sept. 2009 – Jul. 2013

## GRANTS & PROJECTS

- **Optimization Technology for Datacenter Networking Transport**  
PI, affiliated with USTC **Award:** CNY ██████████  
**Funding Source:** NSFC - ██████████ 2025 - 2027
- **Software and Hardware Integrated Acceleration System for Cross-silo Federated Learning**  
PI, affiliated with HKUST **Award:** CNY 300,000  
**Funding Source:** NSFC - Young Scientist Fund 2025 - 2027
- **Adaptive Decision Optimization Technology for Operation Systems Based on Online Learning**  
PI, affiliated with HKUST **Award:** HKD 1,656,000  
**Funding Source:** HKUST-Industry Joint Lab - Cooperation Project 2024 - 2025
- **eBPF Infrastructure for High-performance Intelligent Networking**  
PI, affiliated with HKUST **Award:** HKD 330,000  
**Funding Source:** Gift Fund from Industry 2024 - 2025

## PUBLICATIONS

### Conference Publications:

- C1. [GREEN: Carbon-efficient Resource Scheduling for Machine Learning Clusters](#)  
Kaiqiang Xu, Decang Sun, Han Tian, **Junxue Zhang**, Kai Chen  
*USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2025*
- C2. [Design and Operation of Shared Machine Learning Clusters on Campus](#)  
Kaiqiang Xu, Decang Sun, Hao Wang, Zhenghang Ren, Xinchun Wan, Xudong Liao, Zilong Wang, **Junxue Zhang** and Kai Chen  
*International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2025*
- C3. [eNetSTL: Towards an In-kernel Library for High-Performance eBPF-based Network Functions](#)  
Bin Yang, Dian Shen<sup>#</sup>, **Junxue Zhang**<sup>#</sup>, Hanlin Yang, Lunqi Zhao, Beilun Wang, Guyue Liu and Kai

Chen

*The European Conference on Computer Systems (EuroSys), 2025*

- C4. [Achieving Fairness Generalizability for Learning-based Congestion Control with Jury](#)  
Han Tian, Xudong Liao, Decang Sun, Chaoliang Zeng, Yilun Jin, **Junxue Zhang**, Xinchen Wan, Zilong Wang, Yong Wang and Kai Chen  
*The European Conference on Computer Systems (EuroSys), 2025*
- C5. [Sequoia: An Accessible and Extensible Framework for Privacy-Preserving Machine Learning over Distributed Data](#)  
Kaiqiang Xu, Di Chai, **Junxue Zhang**, Fan Lai and Kai Chen  
*International Conference on Management of Data (SIGMOD), 2025*
- C6. [Fast, Scalable, and Accurate Rate Limiter for RDMA NICs](#)  
Zilong Wang, Xinchen Wan, Luyang Li, Yijun Sun, Peng Xie, Xin Wei, Qingsong Ning, **Junxue Zhang** and Kai Chen  
*ACM Special Interest Group on Data Communication (SIGCOMM), 2024*
- C7. [Efficient Decentralized Federated Singular Vector Decomposition](#)  
Di Chai, **Junxue Zhang**<sup>#</sup>, Liu Yang, Yilun Jin, Leye Wang, Kai Chen<sup>#</sup> and Qiang Yang  
*USENIX Annual Technical Conference (ATC), 2024*
- C8. [Accelerating Privacy-Preserving Machine Learning With GeniBatch](#)  
Xinyang Huang, **Junxue Zhang**<sup>#</sup>, Xiaodian Cheng, Hong Zhang, Yilun Jin, Shuihai Hu, Han Tian and Kai Chen<sup>#</sup>  
*The European Conference on Computer Systems (EuroSys), 2024*
- C9. [Accelerating Secure Collaborative Machine Learning with Protocol-Aware RDMA](#)  
Zhenghang Ren, Mingxuan Fan, Zilong Wang, **Junxue Zhang**, Chaoliang Zeng, Zhicong Huang, Cheng Hong and Kai Chen  
*USENIX Security Symposium, 2024*
- C10. [FLASH: Towards a High-performance Hardware Acceleration Architecture for Cross-silo Federated Learning](#)  
**Junxue Zhang**, Xiaodian Cheng, Wei Wang, Liu Yang, Jinbin Hu and Kai Chen  
*USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2023*
- C11. [Communication Efficient Secret Sharing with Dynamic Communication-Computation Conversion](#)  
Zhenghang Ren, Xiaodian Cheng, Mingxuan Fan, **Junxue Zhang** and Cheng Hong  
*IEEE International Conference on Computer Communications (INFOCOM), 2023*
- C12. [Enabling Load Balancing for Lossless Datacenters](#)  
Jinbin Hu, Chaoliang Zeng, Zilong Wang, **Junxue Zhang**, Kun Guo, Hong Xu, Jiawei Huang and Kai Chen  
*IEEE International Conference on Network Protocols (ICNP), 2023, Best Paper Award*
- C13. [LiteFlow: Towards High-performance Adaptive Neural Networks for Kernel Datapath](#)  
**Junxue Zhang**, Chaoliang Zeng, Hong Zhang, Shuihai Hu and Kai Chen  
*ACM Special Interest Group on Data Communication (SIGCOMM), 2022*
- C14. [Spine: An Efficient DRL-based Congestion Control with Ultra-low Overhead](#)  
Han Tian, Xudong Liao, Chaoliang Zeng, **Junxue Zhang** and Kai Chen  
*International Conference on emerging Networking EXperiments and Technologies (CoNEXT), 2022*
- C15. [Practical Lossless Federated Singular Vector Decomposition Over Billion-Scale Data](#)  
Di Chai, Leye Wang, **Junxue Zhang**, Liu Yang, Shuowei Cai, Kai Chen and Qiang Yang  
*ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2022*
- C16. [Multi-Objective Congestion Control](#)  
Yiqing Ma, Han Tian, Xudong Liao, **Junxue Zhang**, Weiyang Wang, Kai Chen and Xin Jin  
*The European Conference on Computer Systems (EuroSys), 2022*

- C17. [Sphinx: Enabling Privacy-Preserving Online Learning over the Cloud](#)  
Han Tian, Chaoliang Zeng, Zhenhang Ren, Di Chai, **Junxue Zhang**, Kai Chen and Qiang Yang  
*IEEE Symposium on Security and Privacy (Oakland), 2022*
- C18. [Towards the Full Extensibility of Multipath TCP with eMPTCP](#)  
Bin Yang, Dian Shen, **Junxue Zhang**, Fang Dong, Junzhou Luo and John. C.S. Lui  
*IEEE International Conference on Network Protocols (ICNP), 2022*
- C19. [Enabling Low Latency Edge Intelligence based on Multi-exit DNNs in the Wild](#)  
Zhaowu Huang, Fang Dong, Dian Shen, **Junxue Zhang**, Huitian Wang, Guangxing Cai and Qiang He  
*IEEE International Conference on Distributed Computing Systems (ICDCS), 2021*
- C20. [Enabling ECN for Datacenter Networks with RTT Variations](#)  
**Junxue Zhang**, Wei Bai and Kai Chen  
*International Conference on emerging Networking EXperiments and Technologies (CoNEXT), 2019*
- C21. [Resilient Datacenter Load Balancing in the Wild](#)  
Hong Zhang, **Junxue Zhang**, Wei Bai, Kai Chen and Mosharaf Chowdhury  
*ACM Special Interest Group on Data Communication (SIGCOMM), 2017*

**Journal Publications:**

- J1. [eMPTCP: A Framework to Fully Extend Multipath TCP](#)  
Dian Shen, Bin Yang, **Junxue Zhang**, Fang Dong and John. C.S. Lui  
*IEEE/ACM Transactions on Networking, 2024, Early Access*
- J2. [SoK: Fully Homomorphic Encryption Accelerators](#)  
**Junxue Zhang\***, Xiaodian Cheng\*, Liu Yang, Jinbin Hu, Ximeng Liu and Kai Chen  
*ACM Computing Survey, 2024, Volume: 26, Issue: 12*
- J3. [High-performance Hardware Acceleration Architecture for Cross-silo Federated Learning](#)  
**Junxue Zhang**, Xiaodian Cheng, Liu Yang, Jinbin Hu, Han Tian and Kai Chen  
*IEEE Transactions on Parallel and Distributed Systems, 2024, Volume: 35, Issue: 8*
- J4. [A Survey for Federated Learning Evaluations: Goals and Measures](#)  
Di Chai, Leye Wang, Liu Yang, **Junxue Zhang**, Kai Chen and Qiang Yang  
*IEEE Transactions on Knowledge and Data Engineering, 2024, Volume: 36, Issue: 10*
- J5. [LiteFlow: Towards High-performance Adaptive Neural Networks for Kernel Datapath \(extended version\)](#)  
**Junxue Zhang**, Chaoliang Zeng, Hong Zhang, Shuihai Hu and Kai Chen  
*IEEE/ACM Transactions on Networking, 2024, Volume: 32, Issue: 1*
- J6. [Load Balancing with Multi-level Signals for Lossless Datacenter Networks](#)  
Jinbin Hu, Chaoliang Zeng, Zilong Wang, **Junxue Zhang**, Kun Guo, Hong Xu, Jiawei Huang and Kai Chen  
*IEEE/ACM Transactions on Networking, 2024, Volume: 32, Issue: 3*
- J7. [Efficient DRL-based Congestion Control with Ultra-low Overhead \(extended version\)](#)  
Han Tian, Xudong Liao, Chaoliang Zeng, Decang Sun, **Junxue Zhang** and Kai Chen  
*IEEE/ACM Transactions on Networking, 2024, Volume: 32, Issue: 3*
- J8. [Enabling ECN for Datacenter Networks with RTT Variations \(extended version\)](#)  
**Junxue Zhang**, Wei Bai and Kai Chen  
*IEEE Transactions on Cloud Computing, 2023, Volume: 11, Issue: 3*
- J9. [Facilitating Application-Aware Bandwidth Allocation in the Cloud with One-Step-Ahead Traffic Information](#)  
Dian Shen, Junzhou Luo, Fang Dong, Jiahui Jin, **Junxue Zhang** and Jun Shen  
*IEEE Transactions on Services Computing, 2020, Volume: 13, Issue: 2*

**Workshop Publications:**

- W1. [Practical and Secure Federated Recommendation with Personalized Masks](#)  
Liu Yang, **Junxue Zhang**, Di Chai, Leye Wang, Kuo Guo, Kai Chen and Qiang Yang  
*International Workshop on Trustworthy Federated Learning in Conjunction with IJCAI 2022 (FL-IJCAI), 2022*
- W2. [Secure Forward Aggregation for Vertical Federated Neural Networks](#)  
Shuwei Cai, Di Chai, Liu Yang, **Junxue Zhang**, Yilun Jin, Leye Wang, Kun Guo and Kai Chen  
*International Workshop on Trustworthy Federated Learning in Conjunction with IJCAI 2022 (FL-IJCAI), 2022*
- W3. [Aegis: A Trusted, Automatic and Accurate Verification Framework for Vertical Federated Learning](#)  
Cengguang Zhang, **Junxue Zhang**, Di Chai and Kai Chen  
*International Workshop on Federated and Transfer Learning for Data Sparsity and Confidentiality in Conjunction with IJCAI 2021 (FL-IJCAI), 2021, Best Application Paper Award*
- W4. [RAT - Resilient Allreduce Tree for Distributed Machine Learning](#)  
Xinchen Wan, Hong Zhang, Hao Wang, Shuihai Hu, **Junxue Zhang** and Kai Chen  
*Asia-Pacific Workshop on Networking (APNet), 2020*
- W5. [Bridging the Edge-Cloud Barrier for Real-time Advanced Vision Analytics](#)  
Yideng Wang, Weiyan Wang, **Junxue Zhang**, Junchen Jiang and Kai Chen  
*Proceedings of the 11th USENIX Conference on Hot Topics in Cloud Computing (HotCloud), 2019*
- W6. [Rethinking Transport Layer Design for Distributed Machine Learning](#)  
Jiacheng Xia, Gaoxiong Zeng, **Junxue Zhang**, Weiyan Wang, Wei Bai, Junchen Jiang and Kai Chen  
*Asia-Pacific Workshop on Networking (APNet), 2019*
- W7. [Quantifying the Performance of Federated Transfer Learning](#)  
Qinghe Jing, Weiyan Wang, **Junxue Zhang**, Han Tian and Kai Chen  
*International Workshop on Federated Learning for User Privacy and Data Confidentiality in Conjunction with IJCAI 2019 (FL-IJCAI), 2019, Best Student Paper Award*

\* indicates equal contribution and # indicates corresponding author.

## TEACHING & TUTORING

- Instructor, HKUST COMP 3511 – Operating Systems 2024 Fall
- TA, HKUST MSBD 5014 IP – Introduction to Federated Learning and its Applications 2022
- TA, HKUST COMP4651 – Cloud Computing and Big Data Systems 2018
- TA, HKUST COMP4621 – Computer Communication Networks 2017
- TA, HKUST COMP1021 – Introduction to Computer Science 2017

## HONORS & AWARDS

- Best Paper Award @ ICNP 2023 2023
- Honorable Mention, HKUST CSE Best Ph.D. Dissertation Award 2021 – 2022
- Travel Grants, SIGCOMM 2017
- Postgraduate Fellowship, Hong Kong University of Science and Technology 2016 – 2022
- National Scholarship, China 2014
- Google Excellent Student Scholarship 2012
- National Scholarship, China 2011

## ACADEMIC TALKS

- [Accelerating Privacy-preserving Computing: A System-oriented Approach](#)  
▶ Seminar @ University of Science and Technology of China Hefei, Anhui, China, Oct. 07, 2023
- [FLASH: Towards a High-performance Hardware Acceleration Architecture for Cross-silo Federated Learning](#)  
▶ Data Session @ NSDI 2023 Boston, MA, U.S., Apr. 18, 2023  
▶ Seminar @ Southeast University Nanjing, Jiangsu, China, Feb. 23, 2023
- [LiteFlow: Towards High-performance Adaptive Neural Networks for Kernel Datapath](#)  
▶ Machine Learning Session @ SIGCOMM 2022 Amsterdam, Netherlands, Aug. 25, 2022  
▶ SIGCOMM Talks @ APNet 2022 Virtual Event, Jul. 02, 2022

► Seminar @ Southeast University

Virtual Event, Aug. 08, 2022

- [Enabling ECN for Datacenter Networks with RTT Variations](#)

► Datacenters Session @ CoNEXT 2019

Orlando, FL, U.S., Dec. 11, 2019

## INDUSTRY EXPERIENCE

Co-founder/CTO, Clustar Technology Co., Ltd.

2019 – 2023

Subordinates: 50 ~ 100

Achievements: Full stack privacy-preserving computation system including:

- Hardware accelerator: FPGA-based accelerator for cross-silo federated learning, ASIC-based accelerator for fully homomorphic encryption
- Privacy-preserving data analytics/machine learning platform: federated learning, secure multi-party computation, trusted execution environment

## ACADEMIC SERVICE

- CoNEXT 2025
- APNet 2025
- NSDI 2025
- ICNP 2024
- APNet 2024
- APNet 2023
- EuroSys 2023
- ICA3PP 2023 Workshop - Datacenter Networking and Applications
- FL-IJCAL 2022
- FL-AAAI 2022
- FL-IJCAL 2021
- FATE

TPC Member, Web Chair  
TPC Member, Finance Co-Chair  
TPC Member  
TPC Member  
TPC Member, Finance Co-Chair  
Sponsorship Co-Chair  
Shadow PC Member  
TPC Co-Chair  
TPC Member  
TPC Member  
TPC Member  
Maintainer

## JOURNAL REVIEWS

I have been reviewers for the following journals:

- IEEE/ACM Transactions on Networking
- IEEE Transactions on Cloud Computing
- IEEE Transactions on Big Data
- ACM Transactions on Sensor Networks
- ACM Transactions on Intelligent Systems and Technology
- World Wide Web Journal
- Knowledge-Based Systems
- Journal of Supercomputing