

# Jingqi Huang

[LinkedIn] [Homepage]

+1-765-237-7068  
jingqihuang124@gmail.com

## EDUCATION

---

<b>Purdue University</b> Ph.D. in Computer Science	West Lafayette, IN August 2020 - May 2025 (Expected)
<b>University of California, San Diego</b> M.S. in Electrical and Computer Engineering	La Jolla, CA September 2018 - March 2020
<b>Beijing University of Posts and Telecommunications</b> B.E. Electrical and Computer Engineering	Beijing, China August 2014 - June 2018

## SYSTEMS & PROGRAMMING SKILLS

---

- Languages: Python, Go, C/C++, Java, Bash, P4, Matlab, HTML/CSS and SQL/NoSQL
- Technologies: Kubernetes, Docker, Ansible, Helm, MongoDB, Git, gRPC, RESTful API, Prometheus, Grafana, Network System, Mobile Network, Cloud Computing, Distributed System, Mobile Computing, SDN, Mininet, Wireshark

## SELECTED AWARDS AND RECOGNITIONS

---

- **2022 Meta PhD Research Fellowship Finalist for Networking**
- 2022 Intel Intern Recognition
- Beijing University of Posts and Telecommunications Scholarship for Undergraduate Education (2015-2017)

## WORK EXPERIENCE

---

<b>Intel Labs</b> Research Scientist Intern	Hillsboro, OR May 2023 – September 2023
<ul style="list-style-type: none"><li>- Profiled and optimized a distributed, microservice-based cloud-native application, Intel Aether SD-Core, which is a cloud-native <b>5G mobile core network</b> built with <b>Go</b> and orchestrated with <b>Kubernetes</b>.</li><li>- Identified bottlenecks of the user-observed end-to-end latency of high execution time of serialization/deserialization and high read/write IO between the application and database using <b>Go pprof</b>.</li><li>- Added <b>Redis</b> support to reduce the data access latency up to <math>\sim 75\%</math> compared to previous MongoDB solution.</li><li>- Identified the bottleneck in load balancer, HTTP marshaller/unmarshaller and Go scheduler under large-scale users in the adapter between the control plane and data plane and reported to the Intel team for further optimization.</li><li>- This work is featured in the Intel Labs Intern Lightning talk.</li></ul>	
<b>Intel Labs</b> Research Scientist Intern	Hillsboro, OR May 2022 – August 2022
<ul style="list-style-type: none"><li>- Developed stateless microservices and related features in Intel Aether SD-Core using <b>Go</b> to provide high availability and scalability while ensuring resource usage efficiency and low user-observed end-to-end latency. [2, 3]</li><li>- Designed and implemented a microservice using <b>Go</b> to connect to multiple control plane microservices with HTTP interface, and the data plane with a PFCP interface.</li><li>- Developed high availability features for the stateless microservices using <b>Go</b> and <b>MongoDB</b>, including load balancer, keep-alive to ensure user requests can be handled under microservices failure.</li><li>- Designed auto-horizontal <b>Kubernetes</b> pod scaling mechanism to reduce user observed end-to-end latency by <math>\sim 50\%</math>.</li><li>- This work has been open-sourced by Intel.</li></ul>	
<b>Purdue University</b> Graduate Research Assistant	West Lafayette, IN January 2021 – Present
<ul style="list-style-type: none"><li>- Profiled and analyzed the system workload of stateful Intel Aether SD-Core. Deployment and profiling tools include <b>Python</b>, <b>Ansible</b>, <b>Kubespray</b>, <b>RKE2</b>, <b>Helm</b>, <b>Prometheus</b>, <b>Grafana</b>, <b>Go pprof</b>, and <b>Open telemetry</b>. [2, 3, 4]</li><li>- Pinpointed latency bottleneck of the Intel Aether SD-Core in high execution time of message serialization/deserialization (up to <math>\sim 60\%</math>) and high contention time for resources of 5G mobile core microservices (up to <math>\sim 70\%</math>).</li><li>- Intel Aether SD-core project has adopted and open-sourced our development and debugging effort.</li><li>- One first-author paper in submission.</li></ul>	
<b>Intel Labs</b> Research Scientist Intern	Hillsboro, OR May 2021 – August 2021
<ul style="list-style-type: none"><li>- Implemented in-network ML aggregation and Map/Reduce operations on programmable data planes to reduce training and inference latency in distributed machine learning using <b>P4</b>, <b>Python</b>, <b>Bash</b> and <b>BMV2</b>.</li><li>- This work reduced the network latency by <math>\sim 50\%</math>, and saved the network bandwidth by <math>\sim 75\%</math>.</li></ul>	

## University of California, San Diego

Research Assistant

La Jolla, CA  
Sep 2018 – Mar 2020

- Implement and evaluate vehicle-to-everything (V2X) communication over millimeter-wave network use case for 5G New Radio use case using mmWave access points **Airfide Sparow+**, **Matlab**, **Python** and **Bash**. Profiling results guide the mmWave beamforming management mechanisms and interference cancellation. [5]
- Designed and developed X-Array for prototyping and evaluating omnidirectional millimeter-wave (mmWave) network using mmWave access points **Airfide Sparow+**, **Matlab**, **Python**, and **Bash**. [6]

## Beijing University of Posts and Telecommunications

Research Assistant

Beijing, China  
Jan 2017 – Jun 2018

- Designed and developed KPad [9], a system to increase channel utilization in Wi-Fi MU-MIMO.
- Developed Romil [1, 8] for robust indoor mmWave communication.

## PUBLICATIONS

---

### Journal Articles

1. Anfu Zhou, Shaoqing Xu, Song Wang, **Jingqi Huang**, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. Robotic Millimeter-Wave Wireless Networks. **IEEE/ACM Transactions on Networking**, 28(4):1534–1549, 2020

### Conference and Workshop Papers

2. **Jingqi Huang**, Bilal Saleem, Jiayi Meng, Iftekharul Alam, Christian Maciocco, Y. Charlie Hu, and Muhammad Shahbaz. Towards a Performant and Scalable Cloud-Native 5G Mobile Core Architecture. In **SRC TECHCON**, 2023
3. Jiayi Meng, **Jingqi Huang**, Y Charlie Hu, Yaron Koral, Xiaojun Lin, Muhammad Shahbaz, and Abhigyan Sharma. Characterizing and Modeling Control-Plane Traffic for Mobile Core Network. In **ACM Internet Measurement Conference (IMC)**, 2023
4. **Jingqi Huang**, Jiayi Meng, Iftekharul Alam, Christian Maciocco, Y. Charlie Hu, and Muhammad Shahbaz. Accelerating 5G (Mobile Core) Control Plane using P4. In **P4 Workshop**, 2022
5. Haotian Deng, Qianru Li, **Jingqi Huang**, and Chunyi Peng. iCellSpeed: Increasing Cellular Data Speed with Device-Assisted Cell Selection. In **ACM MobiCom**, 2020
6. Song Wang\*, **Jingqi Huang**\*, and Xinyu Zhang. Demystifying Millimeter-Wave V2X: Towards Robust and Efficient Directional Connectivity Under High Mobility. In **ACM MobiCom**, 2020. (\*Equal contribution)
7. Song Wang\*, **Jingqi Huang**\*, Xinyu Zhang, Hyoil Kim, and Sujit Dey. X-array: Approximating Omnidirectional Millimeter-Wave Coverage Using an Array of Phased Arrays. In **ACM MobiCom**, 2020. (\*Equal contribution)
8. Anfu Zhou, Shaoqing Xu, Song Wang, **Jingqi Huang**, Shaoyuan Yang, Teng Wei, Xinyu Zhang, and Huadong Ma. Robot Navigation in Radio Beam Space: Leveraging Robotic Intelligence for Seamless mmWave Network Coverage. In **ACM MobiHoc**, 2019
9. Song Wang\*, **Jingqi Huang**\*, and Anfu Zhou. KPad: Maximizing Channel Utilization for MU-MIMO Systems Using Knapsack Padding. In **IEEE ICC**, 2018. (\*Equal contribution)

## CORE CURRICULUM

---

- **Ph.D.:** Datacenter and Cloud Networks, Data Communication and Computer Networks, Database Systems, Distributed Database Systems, Algorithm Design Analysis & Implementation, Data Mining, Statistical Machine Learning
- **Master:** Software Foundations I, Linear Algebra and Application, Multi User Communication system, Digital Communications, Principles of Wireless Networks, Probabilistic Coding, Special Topic in Signal & Image/Robotic, Big Network Data
- **Undergraduate:** Operating System, Cloud Computing, Databases, Software Engineering, Data Structure, Introductory Java Programming, Signals and Systems, Product Development, Network and Protocols, Security and Authentication, Middleware