

Bojie Li

Last update on July 11, 2023

bojieli@gmail.com • +86.15011272877 • Haidian, Beijing, China • <https://01.me>

Work Experience

Associate Chief Expert (副首席专家) & Assistant Scientist (助理科学家) & Technical Expert A (技术专家A)	BEIJING, CHINA
Computer Networking and Protocol Lab (计算机网络与协议实验室), Huawei Manager: Dr. Kun Tan	Mar '22 – present
Huawei TopMinds Program (华为天才少年)	BEIJING, CHINA
Computer Networking and Protocol Lab (计算机网络与协议实验室), Huawei Manager: Dr. Kun Tan	June '19 – Mar '22

Education

University of Science and Technology of China (中国科学技术大学) Ph.D. in Computer Science	HEFEI, ANHUI, CHINA Sept. '14 – June '19
Joint Ph.D. program with Microsoft Research Asia. Advisor: Prof. Enhong Chen and Dr. Lintao Zhang	
University of Science and Technology of China (中国科学技术大学) B.S. in Computer Science (School of Gifted Young, 少年班学院)	HEFEI, ANHUI, CHINA Sept. '10 – July '14
No.2 Middle School, Shijiazhuang (石家庄二中) Honor Class in STEM (理科实验班)	SHIJIAZHUANG, HEBEI, CHINA Sept. '04 – July '10
Zhongshan Primary School, Shijiazhuang (石家庄市中山路小学)	SHIJIAZHUANG, HEBEI, CHINA Sept. '04 – July '10

Internship Experience

Microsoft Research Asia Systems Research Group Advisor: Dr. Lintao Zhang	BEIJING, CHINA Oct. '16 – May '19
Microsoft Research Asia Wireless and Networking Research Group Advisor: Dr. Kun Tan	BEIJING, CHINA July '13 – Sept. '14 and July '15 – Oct. '16

Selected Projects in Huawei

Unified Bus (UB, 灵衢): Next Generation Data Center Interconnect Ongoing project, instructed by Dr. Kun Tan and Dr. Heng Liao	HUAWEI Apr. '20 – July '23
<ul style="list-style-type: none">Unified Bus (UB, 灵衢) is Huawei's full-stack interconnect fabric to compete with NVIDIA NVLink, Intel CXL, and Infiniband.I am one of the key architects of Unified Bus, directly managing 15 software developers in Hangzhou and Shanghai and indirectly managing an innovation team of 5 local researchers in Isarel and 5 researchers in Beijing.By co-designing the interconnect with Kunpeng (鲲鹏) ARM CPU and Ascend (昇腾) AI NPU, UB achieves 3.2 Tbps bandwidth and sub-μs latency with 10K+ CPUs and NPUs in water-cooled Tiangong (天工) supercomputer for AI and HPC. In comparison, NVLink (7.2 Tbps) only scales to 512 GPUs and Infiniband for large-scale clusters only have 0.4 Tbps bandwidth and 2 μs latency.UB enables direct communication among GPUs in data center scale, bypassing the CPU and thus improving communication performance in LLM training to 5x with 10K NPUs.By overlapping communication and computation, UB improves AllReduce performance in CPU-based ML inference by 20% in search, recommendation, and advertisement systems.Design an instruction set for SmartNIC offloading, improving end-to-end performance of distributed transactions, distributed in-memory database and enterprise storage by 25%~114%.Enable remote direct memory access without pinning memory, increasing effective memory size to 5x with tiered memory and accelerating Spark-based big data application initialization from 2 minutes to 6 seconds.Design a fast process wake-up mechanism in the Linux kernel for NIC interrupts, thus reducing interrupt-mode RDMA latency by 59%~80% on x86 and 52%~77% on ARM.Design a hardware-based Memory-Level Cache that utilizes local DDR as a cache of remote memory, improving Redis performance by 21%~45% and making Redis on remote memory only 2.5% slower than Redis on local memory.	

AKG: Automatic Kernel Compiler for AI Chips using Polyhedral Transformations HUAWEI

Published in PLDI '21, second author, instructed by Dr. Peng Di *June '19 – Apr. '20*

- AKG is the equivalent of CUDNN in the Ascend (昇腾) ecosystem. The major challenge of Ascend NPUs is its TPU-like architecture which requires developers to explicitly manage the memory hierarchy, data movement and synchronization among matrix and vector computation units in the NPU.
- AKG automatically compiles mathematical formulas to low-level assembly of an AI NPU. AKG uses polyhedral compilation to fully automate tiling, scheduling and memory management in a complicated memory hierarchy of an NPU.
- My role: architect and technical leader of the polyhedral compilation framework and dynamic shape team of 15 developers.
- My technical contributions: Propose and implement novel solutions to open problems in polyhedral compilation (non-affine access and parametric tiling); Refine polyhedral schedule constraints to fit NPU architecture; Design a combined manual and automatic scheduling mechanism; Design the architecture and key algorithms to support dynamic shape, e.g., algebraic simplification of multivariate polynomials.

uRPC: A High Performance Unified RPC for Cloud & Client HUAWEI

Ongoing project, instructed by Dr. Kun Tan and Dr. Heng Liao *Apr. '22 – July '23*

- Design a high-performance RPC protocol for Cloud & Client, achieving 5x~10x performance of Google's gRPC, the de-facto standard of RPC.
- Design a cross-language and cross-platform serialization library with easy-to-use APIs similar to Google's Protocol Buffers, but storing unified data structure in off-heap memory, achieving 5x serialization and deserialization performance of Protocol Buffers.
- Design a highly efficient cross-language and cross-platform RPC framework designed for low latency and high throughput. We replace HTTP headers with simple binary headers. Using RDMA hardware, uRPC achieves $< 5 \mu s$ latency, which is 50x lower than gRPC.
- As a unified RPC for Cloud & Client, uRPC integrates optimizations in wireless LAN and client-to-cloud WAN, improving performance of an AI-based photo beautifier by 4x.

Selected Projects during Ph.D. Study in USTC & MSRA

KV-Direct: High-Performance Key-Value Store with Programmable NIC MICROSOFT RESEARCH ASIA

Published in SOSP '17, first author, instructed by Dr. Lintao Zhang *May '16 – Apr '17*

- World's fastest key-value store for AI parameter server, graph computation and web caching.
- Extends RDMA primitives to key-value operations, enabling remote direct key-value access to the main host memory. Also supports vector operations and user-defined functions.
- Achieves up to 180 M key-value operations per second per programmable NIC, equivalent to the throughput of 20 – 30 CPU cores. Built an 1.2 billion KV ops server with 10 FPGA NICs.
- Leverages PCIe bandwidth efficiently and hide PCIe latency with optimized hash table, slab allocator, out-of-order execution, load dispatch and client-side batching.

ClickNP: Highly Flexible Network Processing with FPGA MICROSOFT RESEARCH ASIA

Published in SIGCOMM '16, first author, instructed by Dr. Kun Tan *July '15 – Jan. '16*

- Designed for FPGA-based acceleration of Azure network virtualization platform.
- The first FPGA-accelerated platform for general network functions, written completely in high-level language and achieving 40 Gbps line rate as well as $< 2 \mu s$ latency at any packet size.
- Support high throughput (25Gbps) and low latency ($1 \mu s$) joint CPU-FPGA processing.
- Implement the ClickNP tool-chain, which can integrate with various commercial HLS tools.
- **Work with two undergraduates** to design and implement 100 elements and 5 network functions.

FPGA-based Bing Ranking DNN and HTTPS Accelerator MICROSOFT RESEARCH ASIA

Global 2nd place in Cloud and Enterprise, Microsoft Hackathon 2016; Most Impact Award in Beijing

Venue, Microsoft Hackathon 2017 *Apr. '16 – Oct. '17*

- FPGA-based accelerator for deep neural networks in Bing Search Ranking. Leverage pipeline parallelism to store all parameters of dense layers in the SRAMs of distributed FPGA cluster to avoid swapping out to DRAM. Optimize DRAM and PCIe accesses for sparse layers. Resulting in 5x performance for Bing Ranking.
- FPGA-based accelerator for RSA crypto in HTTPS, serving 10K+ new connections per second with a single server, which is 20x performance of Nginx web server.

- SocksDirect: Datacenter Sockets can be Fast and Compatible MICROSOFT RESEARCH ASIA
Published in SIGCOMM '19, first author; Global 1st place in IT Pros category, Microsoft Hackathon;
instructed by Dr. Lintao Zhang July 17 – Jan 19
- OS kernel consume ~80% CPU time in network communication for Web applications in a server.
 - We design and implement SocksDirect, a user-space high performance socket system fully compatible with existing socket applications.
 - SocksDirect leverages RDMA and shared memory (SHM) for interhost and intra-host communication, respectively. SocksDirect achieves isolation by employing a trusted monitor daemon to handle control plane operations such as connection establishment and access control. The data plane is peer-to-peer between processes, in which we remove multi-thread synchronization, buffer management, large payload copy and process wakeup overheads in common cases.
 - SocksDirect achieves 7~20x better message throughput and 17~35x better latency than Linux socket, and reduces Nginx HTTP latency to 1/5.5.

- Fast and Incremental Web Crawler USTC
Personal Project Jan. 17 – Oct. 17
- A fast web crawler framework based on Scrapy with browser mimicking and rotating proxies.
 - Scrapes 100M+ webpages from 40+ websites within one month for a commercial project, resulting in a database of 1TB+ structured data.
 - Incremental updating mechanism to dynamically probe and update information from the websites.

Publications (selected)

- 1Pipe: Scalable Total Order Communication in Data Center Networks**
Bojie Li, Gefei Zuo, Wei Bai and Lintao Zhang
Proceedings of the 2021 ACM SIGCOMM Conference (SIGCOMM'21)
- AKG: Automatic Kernel Generation for Neural Processing Units using Polyhedral Transformations**
Jie Zhao, **Bojie Li**, Wang Nie, Zhen Geng, Renwei Zhang, Xiong Gao, Bin Cheng, Chen Wu, Yun Cheng, Zheng Li, Peng Di, Kun Zhang and Xuefeng Jin
Proceedings of the 42nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'21)
- SocksDirect: Datacenter Sockets Can be Fast and Compatible**
Bojie Li, Tianyi Cui, Zibo Wang, Wei Bai and Lintao Zhang
Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'19)
- KV-Direct: High-Performance In-Memory Key-Value Store with Programmable NIC**
Bojie Li, Zhenyuan Ruan, Wencong Xiao, Yuanwei Lu, Yongqiang Xiong, Andrew Putnam, Enhong Chen and Lintao Zhang
Proceedings of the 26th ACM Symposium on Operating Systems Principles (SOSP'17)
- ClickNP: Highly Flexible and High-performance Network Processing with Reconfigurable Hardware**
Bojie Li, Kun Tan, Layong (Larry) Luo, Yanqing Peng, Renqian Luo, Ningyi Xu, Yongqiang Xiong, Peng Cheng and Enhong Chen
Proceedings of the 2016 ACM conference on SIGCOMM (SIGCOMM'16)
- Multi-Path Transport for RDMA in Datacenters**
Yuanwei Lu, Guo Chen, **Bojie Li**, Kun Tan, Yongqiang Xiong, Peng Cheng, Jiansong Zhang, Enhong Chen and Thomas Moscibroda
Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI'18)
- Fast and Cautious: Leveraging Multi-path Diversity for Transport Loss Recovery in Data Centers**
Guo Chen, Yuanwei Lu, Yuan Meng, **Bojie Li**, Kun Tan, Dan Pei, Peng Cheng, Layong (Larry) Luo, Yongqiang Xiong, Xiaoliang Wang and Youjian Zhao
Proceedings of the 2016 USENIX Annual Technical Conference (ATC'16)

Professional Activities

- Member, ACM ChinaSys Yong Committee (ACM ChinaSys青年委员会委员) Dec. 21 – now
- Reviewer, ACM Transactions on Computer Systems (TOCS) June 20 – now
- Teaching Assistant, Advanced Software Engineering course, MSRA Sep. 15 – Jan. 16
- Co-founder, USTC icourse.club (USTC评课社区联合创始人) (9K+ users, 23K+ reviews) Feb. 15 – now
- Founder, USTC Freeshell (USTC容器云平台创始人) (2K+ users, 4K+ VMs) Feb. 13 – Sep. 14

Founder, USTC Blog (USTC博客托管平台创始人) (3K+ blogs)	Nov. '12 – Sep. '14
Maintainer, USTC Open Source Mirrors (USTC开源软件镜像维护者) (1TB+ traffic/month)	May '12 – May '13
Member, Gewu Digital Startup (科大格物物联创业团队成员, 因选择读博而退出, 公司后发展成为国仪量子)	Jan. '13 – July '13
President, Linux User Group, USTC (科大Linux用户协会会长)	May '12 – May '13
Founder and President, Technology Division, Student Union, Class of Gifted Young, USTC (科大少年班学院学生会技术部创始人和首任部长)	May '11 – May '12
President, Technology Division, Office of Publicity, Youth League, USTC (科大校团委网络工作办公室技术部首任部长)	May '11 – Jan. '12
Vice President, Gewu Zhizhi Student Club of Physics, USTC (科大格物致知社(物理学术社团)副会长)	May '11 – May '12

Awards in Huawei

Innovation Spark Award (任总颁发的“难题揭榜”火花奖) (100/110000)	Feb. '23
Presidential Team Award of 2012 Labs (2012实验室总裁团队奖) (1/40)	Dec. '22
Huawei Central Software Institute Innovation Contest 2nd prize (10/1300)	Nov. '22
20 Minutes Report to Executive Management Team (每月2名基层员工向公司最高管理团队EMT汇报, 我是第42期) (84/200000)	Aug. '22
Huawei Golden Individual Award (公司金牌个人奖) (1/100)	Dec. '21
Huawei Golden Team Award (公司金牌团队奖) (1/40)	Dec. '21
Huawei Central Software Institute Innovation Contest 1st prize (2/650)	Nov. '21
Huawei Central Software Institute Hackathon 1st place (1/38)	Oct. '21
Outstanding New Employee in 2012 Labs (1%)	Aug. '20

Awards in USTC & MSRA

ACM China Doctoral Dissertation Award (total 2 awardees) (ACM中国优秀博士学位论文奖, 全国共2名)	Dec. '19
Chinese Academy of Sciences Presidential Scholarship (中国科学院院长奖, 全国共400名)	June '19
Outstanding Graduate in Anhui Province (安徽省品学兼优毕业生) (196/4233 = 4.6% in USTC)	Apr. '19
Outstanding Graduate in USTC (中国科学技术大学优秀毕业生)	Apr. '19
Microsoft Research Asia Fellowship Award (微软学者奖学金) (total 10 awardees in Asia)	Oct. '17
China National Ph.D. Scholarship (博士生国家奖学金)	Oct. '17
Global 1st place in IT Pros Category, Microsoft Hackathon 2017 (1/200+)	Aug. '17
Most Impact Award in Beijing Venue, Microsoft Hackathon 2017 (1/200+)	Aug. '17
Best Presentation Award, MSRA Student Techfest 2016 (1/50+)	Oct. '16
Global 2nd place in Cloud and Enterprise, Microsoft Hackathon 2016 (2/1000+)	Aug. '16
Microsoft Research Asia Young Fellowship Award (微软小学者奖学金) (total 10 awardees)	Aug. '13
Outstanding Student Leader in USTC (中科大优秀学生干部)	May '13
Outstanding Lecturer, Web Development Workshop, USTC (“网站开发讨论班”优秀讲师)	Jan. '12
Best Innovation Award, USTC RoboGame Contest (中科大机器人大赛最佳创意奖) (1/52)	Oct. '11

Awards before Undergraduate Study

Bronze Medal, National Olympiad in Informatics (NOI) 2009 (高中全国信息学竞赛全国铜牌)	Jan. '09
First Prize, National Olympiad in Informatics in Hebei Province (NOIP) 2008 (高中全国信息学竞赛河北省一等奖) (rank 2 nd)	Nov. '08
First Prize, National Olympiad in Mathematics in Hebei Province 2009 (高中全国数学竞赛河北省一等奖) (top 10)	Oct. '09
Top 20 in High School Entrance Exam, Shijiazhuang (石家庄市中考前20名)	June '07

First Prize (full score), National Olympiad in Informatics for Junior High School (初中全国信息学竞赛一等奖, 满分) *Oct. '06*

Top 10 Students in Shijiazhuang Primary Schools (石家庄市十佳金童) *Aug. '03*

Merit Student in Shijiazhuang Primary Schools (石家庄市小学三好学生) *Aug. '03*

First Prize, The Hope Cup Mathematical Contest National Finals ("希望杯" 少年数学邀请赛全国总决赛一等奖) *Aug. '03*

Gold Medal, Hua Luogeng Cup Mathematical Contest National Finals ("华罗庚金杯" 少年数学邀请赛全国总决赛金牌) (National top 10) *May '03*
