

HAIYANG SHI

🏠 Redmond, WA 98052

✉ haiyangshi.876@gmail.com

☎ 614-607-4016

SUMMARY OF QUALIFICATIONS

- 10+ years of research and work experience in designing and developing high-performance and scalable systems
- In-depth knowledge of high-performance interconnects (e.g., Remote Direct Memory Access (RDMA)) and protocols, Erasure Coding (EC), In-Network Computing, and distributed systems
- Years of experience in using message queues, databases, and Big Data processing frameworks
- Strong problem-solving and communication skills
- Self-motivated and efficient in cross-group and cross-culture collaboration

EDUCATION

The Ohio State University

Columbus, OH

Ph.D. in Computer Science and Engineering. Advisor: Prof. Xiaoyi Lu

Aug 2016 – Dec 2020

Dissertation: Designing High-Performance Erasure Coding Schemes for Next-Generation Storage Systems

Tianjin University

Tianjin, China

B.E. in Computer Science and Technology

Aug 2008 – Jul 2012

EXPERIENCE

Research Scientist

Infrastructure System Lab @ ByteDance, Seattle, WA.

Feb 2021 – Present

- Research on software/hardware co-design solutions to enhance the performance of company infrastructure systems, such as databases and networking, etc.
- Design and implement caching systems that utilize emerging storage technologies (e.g., persistent memory) to cater to the diverse demands of in-house database systems
- Develop hardware-accelerated transport protocols tailored for data center networks, significantly improving the efficiency and speed of data communication within the company's network infrastructure
- Conduct performance studies of infrastructure systems
- Understand industrial trends and be able to learn and work on cutting-edge technologies such as RDMA, persistent memory, CXL, DPU, etc.

Graduate Research Assistant

The Ohio State University, Columbus, OH

Aug 2016 - Dec 2020

- **High-Performance Erasure Coding Schemes for Next-Generation Storage Systems:** Research project; Multiple papers get published in top conferences
 - Led the research and implementation of involved EC libraries and schemes
 - Designed and developed a multi-rail EC library leveraging heterogeneous hardware, a tripartite graph based EC scheme taking advantage of EC NIC (Network Interface Card) offloading, and a set of in-network EC primitives to accelerate EC computation in distributed storage systems
 - Invented a set of micro-benchmarks to evaluate the performance of various EC schemes
- **Fast and Scalable Remote Procedure Call (RPC) Frameworks:** Accelerated Apache Thrift and brpc for modern data centers
 - Led the design and development of RDMA-oriented RPC frameworks

- Characterized and applied networking optimizations for fast message delivery
- Designed and implemented efficient flow control and communication protocols
- Participated in benchmarking and testing the frameworks in several data centers
- RPC frameworks have been exploited in an enterprise database system and demonstrate up to 20% speedup on 1TB TPC-H benchmark
- **Big Data Processing Frameworks over High-Performance Networks:** High-performance derivative of Apache Hadoop accelerated by RDMA technologies
 - Redesigned and implemented RDMA-accelerated I/O pipelines for both replication and erasure coding storage schemes in HDFS 3.x
 - Involved in benchmarking and testing related software stacks
 - Benchmark results show that the enhanced Hadoop improves overall throughput by up to 3.1×
 - Open-sourced as a part of OSU HiBD project (<http://hibd.cse.ohio-state.edu>)

Research Intern

Alibaba, Seattle, WA. Mentor: Liwei Peng

May 2019 – Aug 2019

- Benchmarked distributed training of Transformer and BERT models with varied configurations to study the impact of RDMA interconnects and all-reduce algorithms
- Developed tools to analyze the bottleneck of distributed machine learning

Software Engineer

MiningLamp, Beijing, China

May 2015 – Jul 2016

- **Scopa:** A property graph based analysis system presenting data as meaningfully defined objects and relationships: people, places, events, and the connections between them.
 - One of the major engineers in designing and developing the system
 - Developed a graph traversal engine for the underlying graph database - TitanDB (or JanusGraph)

Weibo, Beijing, China

Jul 2012 – Apr 2014

- Participated in designing and developing several subsystems for managing profiles and posts of hundreds of millions of users

SELECTED PUBLICATIONS

[VLDB'23] KRYPTON: Real-time Serving and Analytical SQL Engine at ByteDance, Jianjun Chen, Rui Shi, Heng Chen, Li Zhang, Ruidong Li, Wei Ding, Liya Fan, Hao Wang, Mu Xiong, Yuxiang Chen, Benchao Dong, Kuankuan Guo, Yuanjin Lin, Xiao Liu, Haiyang Shi, Peipei Wang, Zikang Wang, Yemeng Yang, Junda Zhao, Dongyan Zhou, Zhikai Zuo, Yuming Liang, In Proceedings of the 49th International Conference on Very Large Data Bases (VLDB), 2023.

[ICDE'23] Accelerating Cloud-Native Databases with Distributed PMem Stores, Jason Sun, Haoxiang Ma, Li Zhang, Huicong Liu, Haiyang Shi, Shangyu Luo, Kai Wu, Kevin Bruhwiler, Cheng Zhu, Yuanyuan Nie, Jianjun Chen, Lei Zhang, Yuming Liang, In Proceedings of the 39th IEEE International Conference on Data Engineering (ICDE), 2023.

[SC'21] HatRPC: Hint-Accelerated Thrift RPC over RDMA, T. Li*, H. Shi*, and X. Lu, In Proceedings of the 34th International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2021. * T. Li and H. Shi contributed equally to this work.

[SC'20] INEC: Fast and Coherent In-Network Erasure Coding, H. Shi, and X. Lu, In Proceedings of the 33rd International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2020.

[SC'19] **TriEC: Tripartite Graph Based Erasure Coding NIC Offload**, *H. Shi, and X. Lu*, In Proceedings of the 32nd International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2019. **Best Student Paper Finalist**.

[HPDC'19] **UMR-EC: A Unified and Multi-Rail Erasure Coding Library for High-Performance Distributed Storage Systems**, *H. Shi, X. Lu, D. Shankar, and D. K. Panda*, In Proceedings of the 28th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2019.

[Bench'18] **EC-Bench: Benchmarking Onload and Offload Erasure Coders on Modern Hardware Architectures**, *H. Shi, X. Lu, and D. K. Panda*, In Proceedings of International Symposium on Benchmarking, Measuring, and Optimizing (Bench), 2018. **Best Paper**.

AWARDS

- SC'19 Best Student Paper Finalist
- SC'19 ACM's Student Research Competition Travel Award
- HPDC'19 Student Travel Award
- Bench'18 Best Paper Award
- ACNN'17 Student Travel Award

PROFESSIONAL INVOLVEMENT

- Publicity chair of HPBDC'20
- Reviewer and external/sub reviewer of TOCS'23, HPCC'21, Cluster'21, SC'21, TPDS'21, CCGrid'21, TPDS Special Section on AI/ML/DL 2020, HiPC'20, Bench'20, UCC'20, TPDS'19