

Zheqi Shen

Summary

Recent graduate with a Ph.D. in computer science seeking an engineer or researcher position, believing that a good solution comes from both algorithmic and systemic efforts.

My research focuses on designing efficient parallel algorithms for large-scale problems, with good performance in both theory and practice.

Additionally, I am experienced in programming and have hands-on knowledge of filesystem and OS development. I run a homelab with variant servers, customizing cooling schemes and the management panel, to develop personal projects.

Education

- July, 2025 (expected) **PhD, Computer Science & Engineering**
University of California, Riverside (UCR), CA | Advised by Yan Gu
 - Research on the topic of efficient parallel algorithms and scalable systems
- 2016–2020 **Bachelor of Engineering, Computer Science**
ShanghaiTech University (ShanghaiTech), China
 - Practices in high-performance I/O and distributed program optimization

Skills and Knowledge

- Languages **C++ (proficient, 16-year coding experience since middle school)**, Python, C, x86_64 Assembly, Matlab, JavaScript, Rust, SQL, Verilog (for FPGA)
 - Developed parallel linker, toy distributed OS kernel, network stack, userspace filesystem, etc
- Parallel Design, analyze, and implement **parallel algorithms and data structure**
Algorithms
 - The first work-optimal parallel LIS algorithm and the efficient parallel vEB tree
 - Research on lock-free parallel ANN algorithms, scaling to 96 cores with 58× self-speedup
 - Experience in debugging and profiling parallel programs, optimizing cache behavior and memory footprint
- Systems **Operating system, network, storage**
 - Hands-on practices in OS kernel programming, RDMA config, IPMI monitoring, and tiered storage
 - Built up a homelab with virtualized servers connected via an L3 switch and IB FDR network

Experience

UCR | Parallel Algorithm Lab

- Sept. 2020 – Present **Research Assistant**, with Professor Yan Gu
 - Developed high-performance approximate nearest neighbor search (ANNS) systems, improving build scalability in multi-threading by 1.6x over SOTA, achieving 3.5x faster filtered search throughput. The new algorithms were wrapped into a C++ template library with a fully modularized design.
 - Designed and implemented efficient parallel algorithms at large scales and novel data structures, reducing cache misses by up to 73.5% and saving memory bandwidth by 66.4%.
 - Designed new methods to parallelize iterative algorithms and metrics to analyze parallel programs.

University of Maryland | ParAlg Lab

- Summer 2022 **Visiting Student Research**, with Professor Laxman Dhulipala
 - Developed functional updates and snapshots for graph-based ANN algorithms and implemented the adaption to the graph container built on the cache-efficient tree embedding.
 - Used vTune to locate the performance bottleneck and optimized the uarch front-end dataflow.

- Sept. 2018 – **Undergraduate Student Researcher**, with Professor *Shu Yin*
Jun. 2020 ○ Designed a middleware between applications and the storage to customize the cache/prefetch strategy with awareness of the data arrangement in the sideband, reducing the I/O load and the tail latency.

Achievements and Awards

- 2025 **Dissertation Completion Fellowship Awards** (four nominations in the dept.), *UCR*
2020 **Dean's Distinguished Fellowship Award**, *UCR*
2019 **Outstanding Student Awards**, *ShanghaiTech*
2018 **Silver Prize**, *ASC Student Supercomputer Challenge*
 - Accelerated the distributed computing software, RELION, with $528\times$ speedup by porting the hotspot to GPUs, reducing the communication overhead, and parallelizing critical for-loops.
 - Tuned the code of CFL3D based on the profiling, achieving 25.6% performance improvement
 - Customized the servers' frequency and fan strategy to achieve the best performance under the constrained power budget; developed an agent to monitor the sensor data and dynamically distribute the power according to the running application and the characters of the workload.

2018 **Fourth Prize, Fan Favorite Prize**, *ISC Student Cluster Competition*
 - Profiled and optimized Nektar++, rewriting the main loops for better parallelism and cache behaviors
 - Tuned the MPI configurations and the packet size to improve the communication efficiency

2018 **Third Prize**, *Robomaster Robotics Competition*
 - Programmed on SoC and developed the electric control of the continuous dodging function

2017 **Outstanding Student Scholarship**, *ShanghaiTech*

Publications and Presentations

Manuscripts

- 2025 **ANNlib: A Development Framework for Efficiently Building ANN Systems**
Zheqi Shen, Jingbo Su, Yan Gu, and Yihan Sun. (under final preparation)

In Conference Proceedings

- 2025 **Pkd-tree: Parallel kd-tree with Batch Updates**
Ziyang Men, Zheqi Shen, Yan Gu, and Yihan Sun. (SIGMOD'25)
- 2024 **BYO: A Unified Framework for Benchmarking Large-Scale Graph Containers**
Brian Wheatman, Xiaojun Dong, Zheqi Shen, Laxman Dhulipala, Jakub Łacki, Prashant Pandey, and Helen Xu. (VLDB'24)
- 2024 **ParANN: Scalable and Deterministic Parallel Graph-Based Algorithms for Approximate Nearest Neighbor Search**
Magdalen Dobson, Zheqi Shen, Guy E. Blelloch, Laxman Dhulipala, Yan Gu, Harsha Vardhan Simhadri, and Yihan Sun. (PPoPP'24)
- 2023 **Parallel longest increasing subsequence and van Emde boas trees**
Yan Gu, Ziyang Men, Zheqi Shen, Yihan Sun, and Zijin Wan. (SPAA'23)
- 2022 **Many sequential iterative algorithms can be parallel and (nearly) work-efficient**
Zheqi Shen, Zijin Wan, Yan Gu, and Yihan Sun. (SPAA'22)

Talks

- 2024 Techniques and Challenges Towards Better Approximate Nearest Neighbor Search
2024 Approximate Nearest Neighbor Search (ANNS): Techniques and Challenges
2023 Techniques and Challenges Towards Better Approximate Nearest Neighbor Search
2023 Approximate Nearest Neighbor Search (ANNS): Techniques and Open Problems
2022 Many sequential iterative algorithms can be parallel and (nearly) work-efficient

Community Involvement

- 2022–2025 **Conference & Journal Reviewer**, *in topics of algorithm, data-management, and HPC*
- SPAA(Junior PC), PPOPP, SIGMOD, HiPC, JPDC, ICPP, ALENEX, TOPC, ESA, Euro-Par
- 2021–2024 **Organizer**, *UCR Programming Contest (UCRPC)*
- Set up the online-judge website and organize the on-site activities
 - Designed and tested contest problems with the auto-generated test cases
- 2019 **Team Leader**, *GeekPie_HPC Association at ShanghaiTech*
- Held lectures biweekly with 20+ participants on average
 - Advised the new teammates and trained for the contests
- 2016–2018 **Student Assistant**, *Teaching Affairs Office at ShanghaiTech*
- Developed and maintained the orientation web system (deployed and used for three years)
 - Participated in developing the IT asset management system