

Jiazhen Hu

Software Engineer · M.S. Computer Science (Georgia Tech, 2028)

Atlanta, GA

+1 678 618 2817 | jayzwho1121@gmail.com | https://hujiazhen.com | JayzWho

Education

Georgia Institute of Technology

M.S. in Computer Science

Atlanta, USA

Expected Aug 2028

Southeast University

B.Eng. in Computer Science and Technology · Major GPA: 3.93 / 4.0

Nanjing, China

Sep 2021 – Jul 2025

Experience

X Asset Management (Hedge Fund, \$1.5B+ AUM)

Shanghai, China

Quantitative Developer Intern · C++ Low-Latency Systems

May 2026 – Aug 2026

- Rebuilt real-time limit-order-book reconstruction for SSE Level-2 tick data, fixing out-of-order Trade/Order delivery with a credit-based mechanism; raised 10-level accuracy 66% → 91.9%, removing ghost orders and crossed quotes
- Built an automated accuracy-scoring framework against exchange snapshots, replacing manual inspection with 5-minute fault localization (27% → 66% → 91.9% hit rate)
- Optimized two order-book containers (14.9M-op trace): value-based storage cut insert cost 25% (14% faster than `std::unordered_map`); a bitmap top-K scan became 14% faster than `std::map`

Projects

Exchange Matching Engine Core

Changzhou, China

C++ High-Performance Systems Development

Feb 2026 – Apr 2026

- Built a limit order matching engine (FIFO, dual-thread, lock-free queues) with a cache-aligned 64B `Order` struct to avoid cache contention, and integer pricing for numerical stability
- Profiled bottlenecks with `perf` FlameGraph and added a cache-friendly hash table + $O(1)$ memory pool; achieved 3.15M ops/s ($2.65\times$) and 2.49ms P99 latency ($5.1\times$ reduction) vs. mutex

Government Data Sharing Platform

Changzhou, China

C++ Backend Development

Nov 2025 – Feb 2026

- Built a high-concurrency `epoll`-based C++ backend with a multi-thread pipeline decoupling I/O from business logic; delivered an end-to-end system with external HTTP services and multi-process daemon support
- Implemented data-driven interface configuration for code-free API extension, plus an OCI-based Oracle extraction/sync layer with checkpoint resume and auto-generated upsert logic

Zynq-based Embedded Collaborative AI Acceleration

Nanjing, China

C / C++ Embedded Systems Development

Dec 2024 – Jun 2025

- Implemented LeNet-5 inference on a Zynq-7000 SoC (CPU+FPGA co-execution) with Int8 quantization ($8\times$ compression) and FPGA-accelerated convolution; kept 97.19% accuracy at $3.17\times$ speedup via mixed partitioning

Skills

Systems/Prog. Concurrency

C / C++; STL; Linux systems programming; TCP/IP Socket; `epoll`
Lock-free programming (SPSC); C++ memory model; performance analysis
(`perf`, `rdtsc`)

Database Toolchain

Oracle; SQL; OCI C library development
Git; CMake; Makefile; gdb; Google Benchmark & Test

Embedded/AI Languages

Xilinx Vivado; Zynq SoC; PyTorch (model training & deployment)
Chinese (Native); English (Fluent, TOEFL 117/120)