

Pusen Yi

pusen.yi@yahoo.com | linkedin.com/in/pusen-yi | github.com/kwfcfc

Education

Northeastern University, MS in Computer Science – Boston, MA, USA

Sept 2024 – present

- GPA: 3.9/4.0
- Relevant coursework: Object Oriented Design, Software Engineering, Data Structures, Introduction to AI, Algorithms

Peking University, BA in Foreign Language and Literature – Beijing, China

Sept 2020 – June 2024

- Relevant coursework: Software Foundations, Algorithms and Data Structures, Linux and Open Source Software

Projects

Self-Study: xv6 Operating System

Dec 2025 – Jan 2026

Self-studied and extended the xv6 Unix-like teaching operating system from MIT 6.1810 on RISC-V.

- Implemented virtual memory mechanisms including page tables, copy-on-write (COW), mmap, and trap handling.
- Extended and debugged core kernel subsystems including process scheduling, system calls, file system, networking, and synchronization primitives (locks), completing 9 labs with ~3,300 lines of kernel-level C code.
- Worked extensively in C within a kernel environment, performing low-level debugging with gdb.

Minnow TCP Implementation

July 2025 – Aug 2025

A simplified TCP/IP stack implementation in C++ based on Stanford's CS144 Computer Networks course.

- Implemented a simplified TCP/IP stack including TCP sender/receiver, stream reassembly, byte stream abstraction, ARP cache, and IPv4 router.
- Benchmarked and optimized ByteStream and TCP reassembly components, redesigning buffering and segment merging logic to reduce memory copies and improve cache locality.
- Achieved **20–40 Gbit/s** ByteStream throughput and **10–30 Gbit/s** reassembly throughput—around **10x faster** than the reference implementation used in the course (~1–3 Gbit/s).

JPMorgan Chase Software Engineering Job Simulation on Forage

Aug 2025

Completed a job simulation project (Midas system) focused on building a transaction processing and balance management platform using Spring Boot, Kafka, and H2 Database.

- Integrated Kafka to consume real-time transaction data streams.
- Developed RESTful APIs to enable external services and client applications to query account balances.
- Implemented database integration with H2 for transaction persistence and balance reconciliation.

Tech credit RAG System

May 2025 – June 2025

A Retrieval-Augmented-Generation (RAG) system to assist in evaluating possible "tech credits" (coding practices or design patterns supporting iterative development).

- Integrated LLMs with in-memory and Chroma vector database to retrieve relevant academic and code examples.
- Implemented document parsing and embedding pipelines that identify code snippets in target repository.
- Built a modular RAG pipeline using LangChain, LangGraph, and OpenAI API with reproducible experimentation in Jupyter.

Skills

Programming Languages: C/C++, Rust, Go, Python, Java, JavaScript, OCaml, Racket, Rocq

Full-stack Development: React, Svelte, Hono, RESTful API, HTML, CSS

DevOps and Cloud Platforms: Docker, Containerization, Nginx, Ansible, Cloudflare Workers, CI/CD pipelines

Messaging and Data Infrastructure: Kafka, PostgreSQL, MySQL, Cassandra

Development Tools: Git, Emacs, LaTeX, Devcontainer, Bash, Jupyter Notebook