# Sheng Xu

412-537-6090 | shengx@andrew.cmu.edu | https://www.linkedin.com/in/sheng-xu666/
Personal website: https://www.sheng-x.com

## EDUCATION

Carnegie Mellon University, Pittsburgh, PA

M.S in Computer Science — *GPA: 4.00 / 4.00* *Dec 2020*

B.S. in Computer Science, — *GPA: 3.54/4.00 (graduated with university honors)* *May 2019*

## PROJECTS (see www.sheng-x.com for more details)

Compiler for C0 (Scoreboard rank #1 project for CMU 15411 Compilers) *OCaml*

- Defeated gcc -O1 on 70% of benchmarks, matched gcc -O1 performance on the other 30%
- Wrote optimizations including sparse conditional constant, partial redundancy elimination propagation, iterative register coalescing, and deadcode elimination
- Implemented features of C0 including pointers, structs, arrays, functions, and exception handling

OS Microkernel (Course project for CMU 15410 Operating Systems) *C / x86 assembly code*

- Implemented virtual memory, scheduler, context switching, exception handlers, and device drivers
- Used locking and synchronization to support a preemptive and thread-safe kernel
- Turned kernel into hypervisor layer capable of hosting multiple guest kernels using paravirtualization

DBMS Concurrent Schema Change (Course project for CMU 15721 Advanced Database) *C++ / SQL*

- Added schema change functionality and ALTER TABLE command to CMU's yet-to-be-named database
- Implemented schema change in a non-blocking manner and supported lazy migration

Flash Translation Layer (Scoreboard rank #1 project for CMU 15746 Storage Systems) *C++*

- Adapted wear leveling techniques from papers to implement an FTL that maximizes SSD lifetime
- Did a followup project that proves the potential of applying machine learning to FTL algorithms

Cloud File System (Scoreboard rank #2 project for CMU 15746 Storage Systems) *C / Linux FS / Amazon S3*

- Built a cloud/local hybrid linux file system supporting 20 syscalls, storing large files on AWS S3
- Supported taking snapshots of the system and mounting/unmounting
- Minimized cloud cost by write-back caching and deduplication with rabin fingerprinting

## EXPERIENCE

Infrastructure for CMU Compilers Course (Quora internship cancelled due to COVID-19) *May-Aug 2020*

- Designed 3 new projects and wrote the infrastructure for the rehauled version of 15411 Compilers
- Managed an EC2-based auto-scaling autograder, developing on a fork of course-management tool Autolab

Blend, Software Engineering Intern *May-Aug 2018*

- Decoupled Blend's largest module, income, from the main repo, and built it as a service
- Improved backend unit testing speed by 10 times, saving backend engineers 30 minutes per day
- Assisted in integrating Docusign to Blend, tailoring to needs of 100+ banks including Wells Fargo and US bank

## RELEVANT COURSES

Systems: Compilers, Operating System, Storage System, Distributed System, Computer Security, Advanced DB, Advanced OS
Algorithms and PL: Foundations of Programming Languages, Advanced Algorithms, TCS Toolkit, Competitive Programming
Math and ML: Probability and Computing, Graph theory, Real analysis, Continuous-time finance, Deep learning, Machine learning

## TEACHING AND AWARDS

Head Teaching Assistant, 15411 Compilers *Fall 2020*

Teaching Assistant, 15451 Algorithms and 15440 Distributed Systems *Spring 2019 - Spring 2020*

Research with Mor-Harchol Balter (wrote 1 book chapter for CMU 15359) *Fall 2018*

ACM-ICPC East Central North America Regionals (Team CMU, 14$^{th}$ place / 124) *Fall 2016*

## SKILLS

Programming Languages: *C*, Python, Ocaml, C++, Java, Golang, Bash, Javascript, Node.js, SQL, Halide, Swift

Tools: Git, x86 assembly, AWS, Linux, LaTeX, Simics, Docker, Pytorch, Tensorflow