

# Shi Liu

☎ +1 (310) 210-0167 | ✉ shiliu@g.ucla.edu | 🌐 shiliu-ucla

## Education

### University of California, Los Angeles

PH.D. IN COMPUTER SCIENCE

- Advised by Prof. Harry Xu. Research interest: resource disaggregation, operating systems, cloud computing.

*Los Angeles, US*

*Sep. 2020 - June 2025 (est.)*

### University of Science and Technology of China

B.E. IN COMPUTER SCIENCE

- GPA **4.05 / 4.30**. Ranked **1 / 202**. Highest honor receiver.

*Hefei, China*

*Sep. 2016 - June 2020*

## Experiences

### A Two-Way Data Plane for High Throughput, Low Latency Far Memory

GRADUATE STUDENT RESEARCHER AT UCLA

- Designed a kernel-runtime co-designed, hybrid data path that can transfer remote data in variable granularities.
- Developed functionalities and optimized paging in Linux kernel memory management sub-system.
- Implemented a C++ runtime for object-granularity data fetches with RDMA.
- Ported real-world applications and optimized their performance, e.g., DataFrame, MapReduce, and Memcached.
- Achieved up to **3x** higher throughput and a **99%** reduction of tail latency, compared against state-of-the-art remote memory systems.

*Los Angeles, US*

*June 2022 - June 2024*

### Java Virtual Machine Redesigned for Memory-disaggregated Data Centers

GRADUATE STUDENT RESEARCHER AT UCLA

- Participated in three projects to improve the efficiency of Java virtual machine running in memory-disaggregated settings.
- Developed two distributed garbage collectors based on production garbage collectors such as G1 GC and Shenandoah GC.
- Designed and implemented a mechanism to line up garbage collector tracing memory access and application access.
- Optimized performance of the JVM systems on real-world applications, e.g., Spark, Cassandra, and Neo4j.
- Achieved up to **4x** higher application throughput and **6x** higher GC throughput.
- Published three papers in top system conferences (OSDI and PLDI), including one winning the OSDI best paper award.

*Los Angeles, US*

*Sep. 2020 - June 2022*

### Distributed Shared Memory Redesigned with Rust Ownership for Ultra Efficiency

RESEARCH INTERN AT ALIBABA GROUP (US)

- Worked on a novel distributed shared memory architecture that minimizes synchronization overheads with ownership.
- Developed the underlying RDMA communication library between servers.
- Implemented language runtime and applications in Rust.
- Achieved up to **30x** higher throughput compared against state-of-the-art distributed shared memory systems.

*Bellevue, US*

*June 2023 - June 2024*

### Automatic Parallelization of DNN models

RESEARCH INTERN AT MICROSOFT RESEARCH ASIA

- Focused on automatic generation of distributed training plans of DNN models.
- Proposed a search algorithm to discover the best partitioning scheme to train large models with hybrid parallelism methods.

*Beijing, China*

*June 2021 - Aug. 2021*

## Publications

### A Tale of Two Paths: Toward a Hybrid Data Plane for Efficient Far-Memory Applications (OSDI '24)

**DRust: Language-Guided Distributed Shared Memory with Fine Granularity, Full Transparency, and Ultra Efficiency (OSDI '24)**

**Canvas: Isolated and Adaptive Swapping for Multi-Applications on Remote Memory (NSDI '23)**

**MemLiner: Lining up Tracing and Application for a Far-Memory-Friendly Runtime (OSDI '22) Jay Lepreau Best Paper**

**Mako: A Low-Pause, High-Throughput Evacuating Collector for Memory-Disaggregated Datacenters (PLDI '22)**

**Semeru: A Memory-Disaggregated Managed Runtime (OSDI '20)**

## Skills

**Programming Languages** C++, C, Python, Rust

**Systems** Linux kernel, Java virtual machine, garbage collectors, RDMA, distributed systems

**Tools** Linux, Bash, Git, Docker, Perf, GDB, public cloud