Jun Kong

Ph.D. Student, Arizona State University, Tempe
jkong123@asu.edu *kongjun18.github.io*

My research spans across the domains of Storage Systems, Data Infrastructures, AI-driven database systems, and AI/ML with Storage Systems.

Internship Availability: May 15th, 2026 – August 18th, 2026

RESEARCH INTERESTS

- **Data infrastructure: Distributed KV stores (FoundationDB); Graph databases (NebulaGraph); Index design for large-scale storage systems; AI-driven storage and database systems.**
- **LLM/AI systems infrastructure: Storage for efficient LLM serving and training, including KV-cache management/offloading, checkpoint/dataset I/O efficiency.**

EDUCATION

Arizona State University **Arizona, US**
Ph.D in Computer Science, advised by Prof. Zhichao Cao *July 2025 - ONGOING*

Lanzhou University **Lanzhou, China**
Bachelor in Computer Science *SEP 2019 - JUN 2023*

- *Outstanding Graduation Thesis Award (Top 10%)*
- *First-Class Scholarship (Top 5%)*

PROJECTS IN PROGRESS

Low Latency SSD-based Indexing for Large-Scale KV Cache Storage in LLM Inference

- *Motivation: When the KV-cache index outgrows DRAM, it must be offloaded to SSD; SSD latency then dominate index lookups, driving up P99 latency and capping LLM serving throughput.*
- *What I'm doing: Designing SSD-friendly indexing for massive KV caches to boost locality and cut random I/O, while keeping an in-memory sparse fingerprint directory that fits DRAM and guides SSD lookups.*

MVCC Turbo: CPU-Efficient Multi-Version Indexing for FoundationDB Storage Servers

- *Motivation: FoundationDB servers execute MVCC via PTreap multi-version indexing; pointer-heavy traversals and key comparisions waste CPU cycles, limiting TPS and inflating latency.*
- *What I'm doing: Profiling storage servers to pinpoint CPU hotspots, then redesigning CPU-friendly multi-version index structures with better range-query performance to reduce per-transaction compute cost.*

INDUSTRY EXPERIENCE

Megvii Technology *AI Infrastructure engineer* *JUN 2023 – JUL 2025*

I worked on scalable storage at MEGVII (A leading Chinese AI startup), focusing on the intersection of AI and storage spanning from object storage systems to training-optimized AI dataset storage systems.

Overlay: An Self-Developed Object Storage System (OSS) with 100+ PiB data in production.

- *Diagnosed and fixed concurrency/performance bugs (e.g. data races, deadlocks, broken heartbeats).*
- *Developed an asynchronous file I/O layer with io_uring to curb Go runtime OS thread overspawning under high concurrency workloads, improving system reliability.*
- *Migrated the storage stack to Kubernetes by containerizing services and authoring Helm charts; enabled safe rolling upgrades and observability, simplifying deployments and day-2 operations for maintenance.*

Nori: Trillion-image–scale dataset loading accelerator for ML training; sustained 600+ GB/s.

- *Designed priority-based scheduling, reducing waiting time for most tasks from hours to seconds.*
- *Redesigned database access patterns and improves maximum system capacity from 800K tasks to 3M tasks.*

UNDERGRUADUATE PROJECTS

LZU OS: A 64-bit education-oriented OS kernel from scratch *LINK* *NOV 2020 - APR 2022*

Lanzhou University, Advised by Associate Prof. Li Liu

- *Built core subsystems from scratch (device tree, bootloader, virtual memory, process magement) and re-implemented Linux 2.6-style buddy + SLAB allocators.*
- *Authored step-by-step tutorials and labs that walk readers from boost to cheduling and memory management, turning the kernel into a reusable teaching platform for OS fundamental.*
- *Won Excellence Award, China National College Student Computer System Ability Contest 2021.*

Simplified Ibex RISC-V Core Re-implementation *LINK* *JUN 2021 - MAR 2022*

Lanzhou University, Advised by Associate Prof. Anping He

- *Engineered a simplified RISC-V (RV32I) processor core from scratch using the Chisel Hardware Description Language, leveraging the open-source Ibex CPU architecture as a reference model.*
- *Verified functionality through simulation and FPGA synthesis.*

INTERNSHIP EXPERIENCE

Megvii Technology *Software Development Internship* *JUN 2022 - OCT 2022*

I worked on a Kubernetes orchestration and package management system similar to Helm.

- *Added support for OCI(Open Container Initiative) standard and transformed packages to OCI format.*

AWARDS AND HONORS

- *Lanzhou University Second-Class Scholarship (Top 10%)*
- *Excellence Award, China National College Student Computer System Ability Contest 2021.*

SKILLS

- *Languages: Proficient in C/C++, Go, Python and Bash.*
- *Tools & More: Experienced in Kubernetes, Docker, AWS, MongoDB, RocksDB and other platforms.*