

# Omkar Desai

[omkarbdesai@gmail.com](mailto:omkarbdesai@gmail.com) | +1 (201) 442-9613 | [linkedin.com/in/omkar-desai/](https://www.linkedin.com/in/omkar-desai/) | [github.com/swiftomkar](https://github.com/swiftomkar)

## EDUCATION

---

Ph.D. in Computer Science (Systems Research), *Syracuse University* Expected May 2025  
Master of Engineering, *Stevens Institute of Technology* Feb 2019  
Bachelor of Engineering, *University of Mumbai* Aug 2017

## EXPERIENCE

---

**Syracuse University**, Syracuse, NY Aug 2020 – Present

*Systems research lab, Research Assistant*

High-impact research in computer systems and storage for transformer-based AI model training with publications in top-tier conferences

- Optimizing HPC systems for AI model training, addressing the gap between GPU and CPU compute capacity [Paper under review]
- Build caching systems for mitigating data pre-processing stalls in large DNN model training using next-gen memories [Paper under review]
- Design of memory-efficient indexing techniques for storage devices that expose a key-value interface [HPDC '23]
- Characterize block I/O workloads and predict I/O performance of a device for a given workload [HotStorage '22]

**Samsung Semiconductor, Inc.**, San Jose, CA May 2022 – Aug 2022

*Memory Solutions Lab, PhD Research Intern*

Conducted simulation & modeling of large-scale storage systems, focusing on reliability and performance for RAID, distributed and parallel file systems, and NFS.

- Developed an all flash array (AFA) simulator for NVMe SSDs based on MQSim+QEMU to simulate application I/O accurately on SSD models
- Demonstrated the scalability of the simulator, accurately simulating an AFA RAID of up-to 256 NVMe SSDs
- Identified and mitigated long tail latencies in distributed AFA storage systems. Patented AI model to predict GC and divert I/Os away, reducing tail latencies by a magnitude [U.S. Patent pending]

**Samsung Semiconductor, Inc.**, San Jose, CA May 2021 – Aug 2021

*Memory Solutions Lab, PhD Research Intern*

Research on leveraging computational storage devices to offload database query processing from host CPU to embedded accelerators. This approach is aimed at improving query processing performance while reducing energy consumption.

- Developed FTL components for memory-efficient indexing in a key-value SSD based on computational storage
- Implemented a machine learning-based system, creating a linear regression-inspired recursive model index for key-value indexing [US patent 11954345]

**Practo**, Bangalore, IN Jun 2019 – Sep 2020

*DevOps group, Software Engineer*

Worked with the Kubernetes, DevOps, and data engineering teams to design & develop infrastructure, deployment, and data strategies for the largest healthcare platform in India with 20 million active users

- Developed a scalable user behavior and events pipeline to ingest 100K events/min from the website to customize recommendations for users
- Led development of ETL tools for data pipelines. Used them for pre-processing data for training DLRM models

**Delos Living**, New York, NY May 2018 – Dec 2018

*Delos Labs, Software Engineer Intern*

Worked with clinicians and building scientists to build an IoT-powered indoor wellness ecosystem using smart sensors and appliances for R&D

- Designed and developed a scalable cloud-native platform on AWS, capable of supporting up to 10,000 homes to facilitate large-scale testing and evaluation efforts at the R&D lab.

## PUBLICATIONS & PATENTS

---

- **Omkar Desai**, Daniel Carlson, Janki Bhimani, and Bryan Kim. 2024. Preparation Meets Opportunity: Enhancing the Data Pipeline for DNN Training with Seneca [Under review]
- **Omkar Desai**, Daniel Carlson, Janki Bhimani, and Bryan Kim. 2024. [A caching system for concurrent DNN model training](#) (ECS Research Day 2024) **(Won first place for poster presentation)**
- Manoj Saha, **Omkar Desai**, Bryan Kim, and Janki Bhimani. 2023. [Leveraging Keys In Key-Value SSD for Production Workloads](#). (HPDC '23) <https://dl.acm.org/doi/abs/10.1145/3588195.3595949>
- **Omkar Desai**, Seungmin Shin, Eunji Lee, and Bryan S. Kim. 2022. [A principled approach for selecting block I/O traces](#). (HotStorage '22) <https://doi.org/10.1145/3538643.3539754> **(Won best presentation award)**
- **Omkar Desai**, Changho Choi, and Yangwook Kang. [Two level indexing for key-value persistent storage device](#). [US patent: 11954345]
- **Omkar Desai**, Shuyi Pei, Jing Yang, and Rekha Pitchumani. [AI model informed garbage collection scheme for AFA RAID systems](#) [Patent pending]
- **Omkar Desai** and Parth Bhatt. 2017. [Home Automation Using Wi-Fi and Sensor Networks](#). *Technofocus*, ISSN:2321-0532

## SKILLS

---

**Domains:** [Distributed Systems](#), [AI/ML systems](#), [Databases](#), [Storage systems](#), Operating Systems, Data engineering, [Flash storage](#)

**Open-Source software contributions:** PyTorch, Apache Airflow, Linux strace, Berkeley Packet Filter (BPF)

**Systems:** File Systems, RAID, Ceph, Lustre, ZFS, LevelDB, RocksDB, Kubernetes, Apache Airflow, Kafka, Hadoop

**Languages & Frameworks:** C++, C, Python, PyTorch, CUDA, Flask, SQLAlchemy, Java, Spring, Redis

**Tools:** Berkeley Packet Filter, Blktrace, strace, AWS, Git, Docker, etc.

## **TEACHING & MENTORING**

---

**Syracuse University, Teaching Assistant**

Aug 2023 – May 2024

- Lead teaching assistant for undergraduate and graduate Operating Systems class
- Designed and delivered high-quality material, was recognized as the best TA

**Syracuse University, Graduate Mentor**

May 2023 – Aug 2023

- Closely mentored undergraduate summer research intern on a research project - training systems for transformer-based LLMs
- Mentored several Graduate (M.S.) students with research and career guidance