

DEVASHISH R. PURANDARE

 [devashishp](#) |  [devashishp](#) |  [sincerely.dev](#) |  devashish.purandare@gmail.com

WORK EXPERIENCE

SK HYNIX AMERICA INC., San Jose, CA | Architecture Enabling Intern 2019

Analyzed cold-storage workload patterns and implemented a cold workload generator to test QLC SSDs. Introduced compression to reduce Ceph's storage bandwidth for recovery. Awarded Most Valued Professional for September 2019.

Technologies: C, C++, Python, NVMe, and Ceph.

RIVERBED INC., Sunnyvale, CA | Engineering Intern 2017

Worked with the Steelhead network optimization team on modernizing the network stats collection system. Replaced legacy Python code with efficient, multi-threaded Go allowing faster aggregation and processing of statistics.

Technologies: Golang, Python, data analysis.

CENTER FOR RESEARCH IN STORAGE AND SYSTEMS, UC Santa Cruz | Graduate Student Researcher 2017 – present

Research in SSDs and systems. Maintenance of research clusters: setting up servers, networking, sysadmin, and sharing. Mentored undergraduates, presented research, and provided feedback to other researchers. Teaching assistant for various courses including programming, data structures, database management, and operating systems.

PROJECTS

Shimmer: Hint Generation and Data Placement on Modern SSDs 2022 – 2023

Shimmer is a dynamic library that interposes on `libc` functions to modify them to inject lifetime and placement hints, and group data ZNS and FDP SSDs. Shimmer improves write throughput for RocksDB and MongoDB by 30-60%, and reduces tail latency by 14× over filesystems (F2FS) and application-specific solutions (ZenFS) with no changes to the applications or the operating system.

Technologies: Rust, C, Python, libc, VFS layer.

Publications: [[CIDR'22](#)] [[Under Submission](#)].

Persimmon: A filesystem for Zoned Namespaces 2021 – 2023

Persimmon is a fork of the F2FS filesystem specialized for Zoned Namespace SSDs that uses append-only data structures for metadata. Persimmon improves utilization, reduces tail latency, garbage collection overhead, and write amplification.

Technologies: C, filesystems, and kernel programming.

Publications: [[ICCD'23](#)].

Analysis of peta-scale scientific archives 2016 – 2019

Analysis of CERN's file system traces over a year to identify trends in large-scale storage. Involves an analysis of over 2.4 billion unique actions spanning 300 million files, with suggestions to improve throughput and reduce duplication.

Technologies: Spark, Python data analysis stack: pandas, numpy, and arrow.

Publications: [[CHEOPS'22](#)] [[SIGOPS operating systems review #56](#)].

EDUCATION

University of California, Santa Cruz, Baskin School of Engineering 2016 – present

Ph.D. COMPUTER SCIENCE March 2024 (expected)

Dissertation: “Enabling host-device collaboration on modern flash for improved performance and lifetime”

MS COMPUTER SCIENCE 2022

Savitribai Phule Pune University, Pune, India. 2012 – 2016

BE COMPUTER ENGINEERING – First Class with Distinction

SKILLS

Programming: Rust, C, C++, Python, Java, Go, SQL, NVMe.

Tools: Linux storage stack, filesystems, `io_uring`, shell, Spark, Docker, Kubernetes, \LaTeX , qemu, strace.

SERVICE

!!Con West | Organizer 2019, 2020

Conference on the joy, surprise, and excitement of computing.

Mastodon: [discuss.systems](#) | Administrator 2022 – present

A Mastodon instance for systems researchers run by Dan Ports, Irene Zhang, and me.