



Department of Computer Science

Designing Exascale Distributed Systems

Saurabh Kadekodi
Google

Hosted by Iulian Neamtiu

DATE: Friday, March 10, 2023

TIME: 11:00 AM – 12:00 PM

LOCATION: GITC 4402

WebEx Link: <https://njit.webex.com/njit/j.php?MTID=mb76545fe49ed4976841525934e10d1bb>

<https://cs.njit.edu/seminars>

Abstract:

Fundamental physical limitations have slowed down hardware scaling, thus ending the “free” scaling benefits of processing power and storage capacity. At the same time, data is growing at an unprecedented rate. This data juggernaut is highly disruptive. It morphs benign assumptions into critical bottlenecks, and forces radical system (re-)designs. My work replaces design decisions of distributed systems that are disrupted by scale with new, data-driven solutions that are efficient, scalable, nimble, and robust. As an example, I will describe disk-adaptive redundancy (DARE): a novel redesign of data reliability in exascale storage clusters driven by insights gleaned from studying over 5.3 million disks from production environments of Google, NetApp and Backblaze. I will also describe three new DARE systems that reduce conservative over-protection of data by up to 20% amounting to millions of dollars of cost savings along with a significant carbon footprint reduction, while always meeting desired data reliability targets. Additionally, I will briefly describe some past and current research efforts to improve the availability and performance of local and distributed storage systems including new erasure codes that reduce observed unavailability events at Google by up to 33%, a novel aging framework that can systematically age local file systems to look over 20 years old in less than 6 hours, and an efficient packing and indexing layer in public cloud infrastructures that boosts the throughput of accessing tiny objects by over 60x while simultaneously reducing the cost of accessing them by over 25000x. Finally, I will touch upon the open challenges in designing exascale distributed systems and highlight promising future directions.

Bio:

Saurabh Kadekodi obtained his PhD in the Computer Science Department at Carnegie Mellon University (CMU) in 2020 as part of the Parallel Data Laboratory (PDL) under the guidance of Prof. Gregory Ganger and Prof. Rashmi Vinayak. After graduation Saurabh joined Google as a Visiting Faculty Researcher, and is currently a Research Scientist in the Storage Analytics team. Saurabh is broadly interested in designing distributed systems with special focus on the performance and reliability of storage systems. At Google, Saurabh is working towards implementing his PhD thesis on disk-adaptive redundancy and other exciting research ideas in some of the largest systems in the world.