Saurabh Kadekodi

http://www.cs.cmu.edu/~saukad *saukad@cs.cmu.edu* *847-912-2652*

Education

Carnegie Mellon University *2014 - present*
Ph.D. in Computer Science
Advisors: Prof. Gregory Ganger, Prof. Rashmi Vinayak and Prof. Garth Gibson
Relevant Coursework *- Storage Systems, Analytical Performance Modeling, Practical Information and Coding Theory for Computer Systems, Advanced and Distributed Operating Systems, Computer Architecture, Advanced Database Systems.*
Teaching *- Served as a TA for Storage Systems (15-746) and Distributed Systems (15-446 / 640)*

Northwestern University *2012 - 2013*
Masters in Computer Science
Advisor: Prof. Peter Dinda
Relevant Coursework *- Operating Systems, Distributed Systems, Resource Virtualization, Independent Study (Compression in Checkpointing and Fault Tolerance Systems), Art of Multi-Core Concurrent Programming, Design and Analysis of Algorithms (Honors), Probability and Stochastic Processes, Information Theory, Graph Theory.*

Pune Institute of Computer Technology (University of Pune) *2005 - 2009*
Bachelors in Computer Engineering
Relevant Coursework *- Computer Organization, Operating Systems, Advanced Computer Architecture, Distributed Systems, Data Structures.*

Research

Carnegie Mellon University

Exploiting heterogeneity in cluster storage systems *(distributed storage systems)*
Exploiting heterogeneity in cluster storage systems can achieve lower storage cost, reduced data access latency and improved device utilization. In particular, decisions regarding the degree of replication in cluster storage systems, when informed by the storage device reliability heterogeneity—even from the same tier—can lead to lower storage cost, and, placement of data when informed by the capacity and data heat heterogeneity can achieve lower data access latency and higher device utilization.

***Geriatrix*: Aging what you see and what you don't see** *(file systems, benchmarking)*
Designed and developed a file system aging tool for artificially aging a file system to encourage realistic and fair benchmarking. The tool takes as input the file age and file size distributions from already aged file system images and performs a randomized sequence of creations and deletions to age the intended file system to mimic the characteristics of the reference file system image. (technical report).

SMRfs: A file system for Shingled Magnetic Recording *(file systems)*
This research aims at building one of the first shingled-aware file systems. Specifics of my work included ways to minimize the size of unshingled partitions (typically used for frequently updated data viz. metadata, small files, etc.) on shingled disks along with the analysis and implementation of inline and opportunistic cleaning algorithms. My work also involved the realization of Caveat-Scriptor; a write-anywhere semantics on SMR disks resulting in substantially low garbage collection and significantly reduced tail latencies. (code / wiki).

Burst-Buffers in HPC Systems *(distributed file systems)*
Ongoing research on exploring efficient distributed file system architectures for supercomputers with burst-buffer installations. Burst-buffers are a layer of SSDs intended to gracefully accept a checkpoint dump from *O(million)* cores and eventually drain to underlying parallel file system.

University of Wisconsin - Madison - *Graduate Research Intern* *Summer 2014*
Advisor: Prof. Remzi Arpaci-Dusseau

SSD over Infiniband *(file systems)*
This study compared the performance between a locally connected SSD and remotely connected SSD (over infiniband). Using the lightweght SCSI RDMA protocol (SRP) for communication, we analyzed the costs in accessing remote SSDs and gained insights into enhancing software architectures of next-gen data centers from the storage perspective.

Price of Ext4 *(file systems)*
Measured the software overhead of the Ext4 file system with the advent of storage devices with microsecond latencies. We threw light on the shifting of bottlenecks in the various submodules of Ext4 and suggested optimizations to make it future-proof.

Publications

Geriatrix: Aging what you see and what you don't see *(file systems)* *ATC 2018*
Kadekodi Saurabh, Nagarajan Vaishnavh, Ganger Gregory R. and Gibson Garth A.

A case for packing and indexing in cloud file systems *(file systems)* *HotCloud 2018*
Kadekodi Saurabh, Fan Bin, Madan Adit, Gibson Garth A. and Ganger Gregory R.

Software-Defined Storage for Fast Trajectory Queries using a DeltaFS Indexed Massive Directory *(file systems)* *PDSW 2017*
Zheng Qing, Amvrosiadis George, Kadekodi Saurabh, Kuchnik Michael Cranor Chuck, Gibson Garth A., Settlemyer Brad, Grider Gary and Guo Fan (slides)

Caveat-Scriptor: Write Anywhere Shingled Disks *(file systems)* *HotStorage 2015*
Kadekodi Saurabh, Pimpale Swapnil and Gibson Garth A.

Space Maps in Ext4 *(file systems)* *OLS 2010*
Kadekodi Saurabh and Jain Shweta

Projects

Space Maps in Ext4 *(file systems)*
Designed and developed an extent-based free-space management technique for the Ext4 filesystem, called Space Maps, along with an allocator that uses Space Maps for disk-space allocation. Consisting of a red black tree and a log, Space Maps enhanced the allocation speed by 30% and deallocation speed by 80% and aided in reducing file and free space fragmentation.

myFTL *(SSD, firmware)*
Designed and implemented a flash translation layer (FTL) with block-mapping, garbage collection (with four policies) and wear-leveling in FlashSim (an FTL simulation software). This project was enhanced and released as a course project for a 70+ student graduate level storage systems course (15-746) at CMU.

Compression in Checkpointing and Fault Tolerant Systems *(fault tolerance, compression)*
Studied the hazy nature of compression algorithms used in checkpoint / restore systems, and went on to suggest possible enhancements and future directions in library-level checkpoint compression for faster and more efficient checkpointing with reduced disk footprint (technical survey).

Database Garbage Collection *(databases, garbage collection)*
Designed, developed and evaluated a co-operative (i.e. not stop-the-world) multi-threaded, lock-free, epoch based garbage collection mechanism for Peloton, a hybrid in-memory database system. Explored tradeoffs between optimizing for average latency versus tail latency due to absence of dedicated garbage collection thread (code).

VM Co-Migration *(virtualization)*
Designed and developed a UDP based VM migration module in Palacios - an OS independent embeddable VM monitor. It supported multiple-source multiple-destination migrations specifically aimed at distributed applications in HPC environments (viz. supercomputers) to exploit page-sharing among participating nodes giving increased parallelism for migration (technical report).

DNA Compression *(compression)*
Explored a run-length based preprocessing scheme exploiting the power-law behavior of genomic data to reveal possibilities of Markovian compression and variable length encoding algorithms for higher compression ratio than provided by existing dictionary based compression algorithms (technical report).

NIC of Time *(networking)*
Designed and developed a tool for exploring the state space of all possible combinations of offloaded functionalities on the NIC vs their presence in the kernel. The tool performs extensive analysis of throughput and CPU utilization to suggest one or a group of features that should be offloaded to the NIC depending on the particular workload under consideration (code).

Active Databases *(distributed systems)*
Implemented a proof-of-concept of decentralized active databases on top of Kademlia - a distributed hash table on a decentralized peer-to-peer network. Active Databases essentially mean event-driven databases following event-condition-action (ECA) rules (technical report).

Industry

Alluxio - *PhD Intern* *Summer 2017*
Designed and developed a packing and indexing layer in the writeback cache of Alluxio (a distributed user-level file system) to transform arbitrary user workloads to write-patterns that are better suited for the cloud. My optimization resulted in >60× increase in throughput and >25000× reduction in cost required to ingest data into Amazon S3.

Google - *PhD Intern* *Summer 2016*
Worked at the intersection of the Storage Analytics and Colossus (distributed file system) teams on analyzing access patterns from Google services and designing a load balancing strategy to avoid disk hotspots in data centers with over hundred thousand disks with heterogeneous capacities and speeds.

Apple - *PhD Intern* *Summer 2015*
Spent a summer as a part of the file systems team in the CoreOS divison.

Soft Corner - *Software Developer* *Oct 2010 - Sep 2012*
PatientScribe *(healthcare web application on ipad, fault tolerance)*
Worked on data reliability, fault tolerance and audit trails at the clinical level. Developed global data reliability using data analytics. Also developed a distributed data-storage algorithm for seamless remote backup and restore.

Project-X *(social network for academia)*
Designed and developed a custom stackable context-based framework for the social network. Developed a suggester module based on a graph-based traversal of the user connection web to help generate meaningful suggestions based on connected entities. Extended WWW SQL Designer by adding collaboration (using Node.js), versioning and a chat module to build a real-time block diagram and discussion tool.
Patent - Kadekodi Saurabh, Narayanan Shrikanth, Ranade Sanyogita, Patil Bharat. 2012. *Project-X*. India Patent Application 2024/MUM/2012, filed July 12, 2012 *(pending)*.

Spring Computing Pvt. Ltd. - *Software Engineer* *Jun 2009 - Oct 2010*
Emdebian Research *(file systems)*
Created a 9 MB JFFS2 filesystem image using Emdebian Crush (an online repository that helps create BusyBox based root filesystem images) and ported the resulting filesystem image on an ARM board.
Kernel Porting on Embedded Devices *(embedded systems)*
Participated in full cycle from porting Linux kernel 2.6.30 onto 3 customized ARM boards, testing and tweaking peripherals to adapt to the new kernel and making public releases of the boards.

Leadership

Product Owner / Architect of Project-X - a social networking solution to bridge the gap between students, universities and the industry.

President of the PICT Art Circle - a group of 50 theater enthusiasts. *Jun 2007 - May 2009*

Mentorship of Masters Students - in 4 independent studies performed in shingled disks, one capstone project about packing small writes in HPC applications, one capstone project about garbage collection of packed blobs in cloud file systems. *Fall 2014, Fall 2017, Fall 2018*

Mentoring Teaching Assistants - for future iterations of the storage systems course. *Spring 2016*

Board member of CMU's Indian Graduate Student Association (IGSA). *Dec 2016 - present*

Skills

Programming Languages - Proficient in C, C++, Python and JavaScript with significant experience in PHP, Java and Shell Scripting.

Systems - Proficient in Linux kernel / userspace development, file system architecture and in web application development.