

## EDUCATION

<b>Syracuse University, NY</b>	Aug 2020-present
<b>Doctor of Philosophy (Ph.D.) in Computer Science</b> – Storage Systems Researcher	
<b>Stevens Institute of Technology, NJ, USA</b>	Feb 2019
Master of Engineering in Computer Engineering	
<b>University of Mumbai, India</b>	Aug 2017
Bachelor of Engineering in Biomedical Engineering	
<b>Shri. Bhagubhai Mafatlal Polytechnic, India</b>	Jun 2014
Diploma in Computer Science	

## SKILLS

<b>CORE:</b>	<u>Systems Engineering, Operating Systems, Databases, Distributed Systems, Data engineering, Flash based storage, IoT</u>
<b>Systems:</b>	Unix File Systems, RAID, NVMe, LevelDB, RocksDB, Kubernetes, Nginx, Apache Airflow, Kafka, Hadoop
<b>Languages &amp; Frameworks:</b>	C++, C, Python, Flask, SQLAlchemy, Java, Spring, Redis, Mvn, Javascript, Swift, HTML5, Node.js
<b>Database:</b>	Transactional & analytical databases based on SQL & NoSQL (MySQL, DynamoDB, PostgreSQL, MongoDB)
<b>Tools:</b>	AWS, Serverless, Git, raspberry pi, Docker, Kubernetes, Tracing (BPF, strace)

## WORK EXPERIENCE

<b>Syracuse University, <u>Research Assistant</u>, Syracuse, NY</b>	Aug 2020 - Present
<ul style="list-style-type: none"><li>Storage Systems Research. Currently working on <u>NSF grant – CPR for Flash-Based Storage Systems</u> under <u>Prof. Bryan Kim</u>.</li><li>Studying big data Key Value databases like LevelDB, RocksDB, Redis in depth in order to advance the database field.</li><li>Researching ways to make distributed storage more resilient and intelligent to understand heterogeneity and ageing of SSDs.</li><li>Developing an <u>open source tool for remote monitoring of storage devices</u> like SSD's to study their Performance and Reliability.</li></ul>	
<b>Practo, <u>Software Engineer / Data Engineer</u>, Bangalore, IND</b>	Jun 2019- Sep 2020
<ul style="list-style-type: none"><li>Developed and maintained the user behavior and events ingestion pipeline that forms the backbone of Data Analytics, Machine Learning and powers the Recommendations infrastructure at practo.com.</li><li>Worked on bringing in depth visibility on the production Kubernetes cluster using Prometheus, Grafana</li><li>Deployed and maintained Apache Airflow on Kubernetes in a novel configuration; Wrote in house libraries &amp; deployment automations.</li><li>Developed ETL workflows for massive data cleaning operations; trained the business analytics teams on ETL (Extract Transform Load)</li><li>Managed a multi-terabyte, production, database cluster for efficient data analysis, built data visualization tools for analysts.</li></ul>	
<b>Delos Living, <u>Software Engineer (Intern)</u>, New York, NY</b>	May 2018- Dec 2018
<ul style="list-style-type: none"><li>Designed and implemented a cloud native system for a cloud connected IoT platform.</li><li>Working with an R&amp;D team at Delos Labs, my work included designing and building cloud-based infra and IoT tools to build an indoor wellness ecosystem.</li><li>Developed a data pipeline to record data from a heterogenous set of wearables and Environmental sensors (1000 data pts/hr). This forms the backbone of the POC and research evaluation efforts at the R&amp;D lab to create advanced models of human stress and sleep. This data also powers the Machine learning efforts in home automation.</li></ul>	

## PROJECTS & ACHIEVEMENTS

<b><u>Capacity Performance Reliability (CPR) for Flash-Based Storage Systems</u></b>	Oct 2020 – Sep 2023
<i>Summary: exploiting tradeoffs among CPR and designing a capacity-variant interface that allows the SSD to maintain performance while gracefully reducing the capacity. link: <a href="https://www.nsf.gov/awardsearch/showAward?AWD_ID=2008453">nsf.gov/awardsearch/showAward?AWD_ID=2008453</a></i>	
<ul style="list-style-type: none"><li>Quantify the error-induced performance degradation by building an SSD aging framework.</li><li>Build a capacity-variant system and demonstrate the effectiveness of a capacity variant SSD.</li><li>Develop new filesystems and RAID systems to study how capacity-variance can be extended to a heterogenous set of SSDs</li></ul>	
<b><u>Creating Synergies between Memory, Disk and Log in Log Structured Key-Value Stores</u></b>	Aug 2020 – Dec 2020
<i>Summary: Improving the background I/O performance on LevelDB – an open sourced key value store by Google</i>	
<ul style="list-style-type: none"><li>Improved the background process efficiency, made performance improvements and ensured consistency</li></ul>	
<b><u>Parallel programming on distributed clusters</u></b>	Jan 2018- May 2018
<i>Summary: parallel computing, high-performance algorithms, and complexity analysis</i>	
<ul style="list-style-type: none"><li>Developed high-speed sorting and numerical computation algorithms for using C++ and MPI, achieved 5x improvement in performance</li></ul>	

Research Papers:

- Desai, Omkar. 2017. "Home Automation Using Wi-Fi and Sensor Networks." *Technofocus*, Mar 2016 ISSN:2321-0532 Mar 2016