Research Statement

Janki Bhimani, Assistant Professor,

Knight Foundation School of Computing and Information Science, Florida International University jbhimani@fiu.edu, bhimanijanki@gmail.com · Linkedin · Website · Google Scholar · Mobile: +1(857)991-9868

(Please visit my Website for the most updated information.)

The global datasphere is projected to reach 163 zettabytes by 2025, indicating a profound shift in our daily lives driven by data. As datacenters adapt to meet escalating computational and storage demands, the challenge of efficiently managing this abundance of information becomes critical for system architects. My decade-and-a-half-long research focuses on **pioneering simple yet disruptive and efficient solutions for memory management and storage systems to address this challenge**. The goal is to establish robust, parallel, and performant computing systems capable of thriving in our data-intensive landscape. This work not only aligns with current federal funding priorities but is imperative for shaping the future, where innovation and efficiency converge to redefine the boundaries of large-scale data-intensive workloads.

[Publication Decisions.] My work as a faculty member primarily focuses on storage systems, system design and architecture, modeling and simulation, and cloud computing domains. Consequently, I have directed my efforts towards the highest-ranked journals in these disciplines, including ACM TOS, IEEE TC, ACM TOMACS, IEEE TCC, and IEEE TVT. For instance, the IEEE Journal of Transactions in Cloud Computing has an impact factor of 11.1, and most of these journals publish less than 10% of submitted manuscripts. Due to the often lengthy time-to-press for journal publications and the time-intensive nature of qualitative research, I frequently target preliminary findings for highly-selective peer-reviewed conferences and workshops such as HPCA, DAC, DATE, HotStorage, CLOUD, and HPDC. This approach enables me to disseminate preliminary findings quickly to target audiences while also building a narrative for subsequent journal publications. Currently, I have 45 peer-reviewed publications, with 20 published after joining FIU in the last five years. I am also the lead inventor of 10 patents. My work has received multiple Best Paper Awards and nominations for the prestigious IEEE MICRO Top Picks Award and In Company of Women Award in the Science and Technology category, showcasing its impact. Additionally, my research is highly cited, with about **1000 citations**, an h-index of 16, and an i10-index of 22, reflecting broad influence. Recognizing the excellent quality of my research, I am a recipient of the NSF CAREER Award, FIU Top Scholar Award in the category of- Research, Creative Activities, and Award-Winning Publications, and the Outstanding **Research Award** by the KFSCIS. I have also been invited to speak on my research by various organizations, including the Entrepreneurs' Organization of Miami and industry research labs such as IBM and Samsung Research.

[Research Funding.] In my research, I have achieved significant milestones by securing a total of thirteen grants valued at over \$6 million, with ten of them valued at over \$3 million as the lead/sole Principal Investigator from federal, state, and industry sources such as NSF, Cyber Florida, and Samsung Semiconductors Notably, I am proud of my recent highly competitive NSF CAREER Award towards designing efficient in-storage indexing techniques, and NSF CISE Core small award as the sole PI, focusing on designing new memory management techniques for in-memory analytic frameworks and databases leveraging machine learning (ML). These funding's underscores the significance of my research work in the national discourse. Currently, I am awaiting decisions on five proposals, two of which I lead. Looking forward, my plans involve pursuing large grants from NSF and other opportunities such as DOE CAREER and AFOSR grants to further advance my research.

[Research Collaborations and Advising.] My commitment to advising and mentoring is reflected in the successful graduation of one Ph.D. student who is a tenure-track Assistant Professor at Missouri State University. Over the past five years, I have advised a diverse group, including eight Hispanics, one Asian American, and three Women. Currently, I am supervising a dynamic team comprising six graduate and three undergraduate students and I am in process of hiring two more graduate students starting Fall 2024. Establishing fruitful collaborations is integral to my research approach. I have established partnerships with esteemed institutions such as the University of Maryland, University of Chicago, Argonne National Lab, and Syracuse University, as well as industry leaders like Samsung and IBM Research, contributing to the advancement of our collective research goals.

[Research Experience and Successes.] My research philosophy is centered on advancing "end-to-end system design" to maximize the utilization of available resources and effectively address societal challenges. I am deeply passionate about optimization and emerging technologies. Engaging in interdisciplinary research is a particular interest of mine, and I consider benchmarking and modeling as essential tools for conducting thorough investigations and driving disruptive improvements. At a high level, my research has concentrated on improving two critical components of computer systems: Memory and Storage. Over the last five years as a faculty member, my research contributions can be categorized primarily into the following three directions.

I/O Behavior Modeling, Performance Prediction and Optimization: In the era of "Big Data," where multiple data processing applications coexist in data centers, I/O activities exhibit significant variations. The configuration of existing storage systems, typically done during installation and then permanently maintained, is becoming insufficient. Modern data processing systems present three key challenges. First, simultaneous operations from multiple applications create interference, impacting the performance of solid-state drives (SSDs), unlike hard disk drives (HDDs) with limited bandwidth. Second, while HDD reliability relies on internal mechanical components, SSDs are sensitive to user and

operating system I/O workloads, necessitating optimal configuration for persistent storage. Third, diverse SSD types, including multi-stream SSDs and Key-Value (KV) SSDs, come with distinct internal algorithms and parameters. Selecting and tuning these algorithms based on I/O activities is crucial for optimal performance and flash endurance. Therefore, the research objective is to model the complex I/O activities of diverse applications, dynamically tuning the internal algorithm parameters of flash-based SSDs for optimal performance and reliability.

Learning and Management in Tiered Memory Systems: Further, in the past decade, ML has undergone astounding growth, permeating various industries, including storage systems. To tackle the challenges presented by vast amounts of data and optimize memory accesses, tiered memory systems are gaining popularity. These systems employ high-speed memory like DRAM for frequently accessed upper-tier data and slower but larger memories like NVMe NAND flash, 3D-Xpoint, and CXL memories for lower tiers. These systems prove crucial for efficient data management, contributing to improved performance and efficiency, reducing data access times, and lowering overall computing costs. Thus, we design novel tiered memory management techniques, leveraging ML's power and addressing its limits and overheads as a versatile solution to enhance various aspects, including parameter tuning, task scheduling, scanning, migration, and allocation, ultimately optimizing performance, Quality of Service (QoS), and resource utilization in in-memory databases and analytic frameworks.

Towards Efficient In-storage Indexing and Device Endurance: In the domain of in-storage indexing, a concept originating in the 1990s, the practical implementation has encountered obstacles, as exemplified by Seagate's 2014 endeavor with Kinetic HDDs, constrained by HDD I/O limitations. However, the prospect of efficient in-storage indexing has recently gained momentum with flash-based SSDs. Early efforts, including our preliminary research, demonstrate that the combination of a fast and lightweight KV database or POSIX-compliant file system with a keyvalue SSD (KVSSD) performing in-storage indexing outperforms traditional block SSDs. The development of Efficient In-Storage Indexing Devices (ISIDs) poses challenges that necessitate attention to ensure optimal performance and functionality. Firstly, the role of storage device models is pivotal in computer systems research, addressing research gaps related to performance analysis, algorithm development, system evaluation, resource management, and realistic simulations. The lack of a low-cost open-source research platform hampers rapid adoption. Secondly, the design of ISIDs for diverse workloads demands meticulous consideration of indexing techniques, query optimization, data access patterns, and data distribution within the constraints of limited device resources. Therefore, we **develop ISID models that capture internal feature dependencies and support dynamic model calibration. This aims to develop new index management techniques efficiently utilizing limited on-device resources while considering flash-specific constraints to optimize endurance and latency for a multi-tenant environment.**

In conjunction with the above, my team has embarked on pursuing various other cutting-edge projects. First, we challenge the assumption that flash-based SSDs are less susceptible to diverse environmental conditions such as vibrations, temperature, and humidity than HDDs. Extensive testing indicates that, even within specified datasheet limits, short-term exposures exhibit lingering effects, and long-term exposure results in over **30%** performance degradation. This raises concerns in datacenter performance, affecting tail latency and Service Level Agreements (SLAs) but has even more crucial implications on automotive operating in harsh conditions, with its transition to use more and more complex electronic control systems (ECS) with integrated flash. Leveraging these insights, with a particular focus on the impact of temperature on the reliability of electronic control systems (ECS), we study automotive applications with multi-core processing architectures, taking into account temperature considerations and system-level reliability. We **optimize system-wide reliability through a mathematical programming model and a genetic algorithm**, **accurately predicting system-wide mean-time-to-failure (MTTF) with substantial speed-ups, thereby enhancing reliability analysis**.

Second, another initiative involves critically examining wear leveling in SSDs, addressing challenges, assessing effectiveness, and advocating for capacity variance. Third, in the space of Systems-for-ML, our team is designing new models to enhance data storage and ingestion pipelines for ML workloads. We are also investigating novel approaches to optimize Deep Neural Network (DNN) checkpointing and versioning beyond traditional file I/O methods and designing models. Forth, in the space of ML-for-Systems, we are leveraging disciplined data science to revolutionize decision-making in storage I/O. Fifth, we are working towards optimizing data spilling in distributed query engine. Lastly, my commitment extends to advocating for positive change in educational systems through the NSF HSI project VOCES - Voices for Organizing Change in Educational Systems, leading efforts in the direction of social media impacts and improving course curriculum. Thus, to conclude, my diverse undertakings showcase my comprehensive engagement in advancing knowledge and addressing pivotal challenges across multiple facets of data storage and computing as well as the education system.

[Future Plan.] In my envisioned trajectory, I am committed to further solidifying my standing as an expert in memory management and storage systems, with a dedicated focus on emerging technologies. My work will continue to contribute to and influence the national discourse around efficient, reliable, and enduring data management techniques. Each successive phase of my research is strategically crafted to make practical and measurable strides towards optimizing data management. My dedication to solving important research problems and consistently publishing in high-quality venues remains unwavering. I plan to enhance my recruitment strategies, aiming to intentionally diversify my doctoral student cohort. Looking forward, I aspire to leverage the national network I have cultivated to foster collaborations on a broader scale. Initiating endeavors for large-scale, multi-institutional funding is a pivotal goal, intended to amplify the impact of my research endeavors and revolutionize data management practices.