

Mihir Nanavati

Software Engineer, IOP Systems

www.mihirnanavati.com
jobs@mihirnanavati.com

Industrial Experience

- 2023 – **Software Engineer, IOP Systems, New York City, NY.**
Building an automated performance testing platform to measure and optimize the performance of programs and help increase the efficiency and utilization of the underlying hardware.
- 2021 – 2023 **Staff Software Engineer, Twitter, New York City, NY.**
Infrastructure Efficiency: Increased density for co-located workloads running on shared on-prem hardware resulting in savings of \$xxM in 2022 and projected savings of \$xxxM in 2023+ by providing better cache, socket, and network interface isolation between workloads.
Compute/Orchestration: Scaled Twitter’s compute clusters to support millions of cores by developing a federation layer that transparently aggregated them into a single monocluster. Built a write-aggregating and connection-pooling proxy for cluster services to announce themselves to a ZooKeeper-based service discovery system.
- 2018 – 2020 **Postdoctoral Researcher, Microsoft Research New York, New York City, NY.**
Applied machine learning to optimize the performance of data store indexes (in a research context), and the efficiency of Azure’s virtual machine provisioning (as part of the *Real World Reinforcement Learning* team), which resulted in a 15-25% reduction in downtime [SoCC’21].
- 2014 – 2016 **Senior Software Engineer, Coho Data, Vancouver, BC.**
Lead development of Decibel [NSDI’17]: a “next-generation” network-attached block storage device which presented virtualized and isolated remote storage on top of shared devices. By bypassing the kernel for I/O, it achieved millions of IOPS at microsecond latencies (comparable to local devices) on a 1U server—over an order of magnitude improvement over the existing system.
- 2006 – 2009 **Security Software Engineer, Miel e-Security Systems, Mumbai, India.**
Developed Helios to detect rootkits and other hidden malware on Windows using cross-view detection. Helios was later integrated into the Intel open-source RPIER forensic toolkit.

Education

- 2019 **Ph.D. in Computer Science, University of British Columbia, Vancouver.**
Thesis: Operator, Number Please: Mediating Access to Shared Resources for Efficiency and Isolation
Advisors: Andy Warfield and Bill Aiello
- 2011 **M.Sc. in Computer Science, University of British Columbia, Vancouver.**
Thesis: Breaking Up is Hard to Do: Security and Functionality in a Commodity Hypervisor
Advisors: Andy Warfield and Bill Aiello
- 2007 **B.E. in Computer Engineering, Gujarat University (LD Engineering), Ahmedabad, India.**

Academic Research Highlights

- [Nature’20] **Reproducible, Distributed Genomics, University of British Columbia.**
Reduced the runtime of whole genome sequencing analysis pipelines from weeks and months to hours and days by distributing compute jobs across private clusters and public clouds using a custom-built serverless container framework. Enabled reproducibility by tracking data provenance and reduced wasted compute by caching intermediate results for frequently executed analyses.

- [Eurosys'13] **Fine-Grained Memory Remapping**, *University of British Columbia*.
Transparently eliminated false sharing in running applications with a combination of binary instrumentation-based, sub-page memory remapping and low-overhead cache line contention detection. Moving falsely shared regions of memory onto independent, non-contending cache lines improved performance in certain multi-core benchmarks by 3–6x.
- [SOSP'11] **Hypervisor Disaggregation**, *University of British Columbia and Citrix Systems*.
Reduced the overall trusted code (TCB) in the Xen hypervisor by 90% by decomposing a monolithic privileged domain into a set of independent domains, each with reduced privilege. Also reduced the temporal scope of exploits by periodically rebooting privileged domains to known-good states.
- [HotSec'11] **Authorship Attribution**, *University of British Columbia*.
Demonstrated the fragility of double-blind peer-review in academic conferences by building a stylometric classifier that uses tells in labelled reviews to identify the authors of anonymous ones.

Selected Publications

- [SoCC'21] **Sayer: Using Implicit Feedback to Optimize System Policies**
Mathias Lécuyer, Sang Hoon Kim, Mihir Nanavati, Junchen Jiang, Siddhartha Sen, Amit Sharma, and Aleksandrs Slivkins
ACM Symposium on Cloud Computing (SoCC), 2021
- [Nature'20] **Massive Haplotypes Underlie Ecotypic Differentiation in Sunflowers**
Marco Todesco et al. (Rieseberg Lab including Mihir Nanavati)
Nature, Plant Genomics, 2020
- [HotCloud'20] **Disaggregation and the Application**
Sebastian Angel, Mihir Nanavati, and Siddhartha Sen
USENIX Workshop on Hot Topics in Cloud Computing (HotCloud), 2020
- [NSDI'17] **Decibel: Isolation and Sharing in Disaggregated Rack-Scale Storage**
Mihir Nanavati, Jake Wires, and Andrew Warfield
USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2017
- [Eurosys'13] **Whose Cache Line is it Anyway? Operating System Support for Live Detection and Repair of False Sharing**
Mihir Nanavati, Mark Spear, Nathan Taylor, Shriram Rajagopalan, Dutch T. Meyer, William Aiello, and Andrew Warfield
ACM European Conference on Computer Systems (Eurosys), 2013
- [SOSP'11] **Breaking Up is Hard to Do: Security and Functionality in a Commodity Hypervisor**
Patrick Colp, Mihir Nanavati, Jun Zhu, William Aiello, George Coker, Tim Deegan, Pete Loscocco, and Andrew Warfield
ACM Symposium on Operating Systems Principles (SOSP), 2011
- [HotSec'11] **Herbert West – Deanonymizer**
Mihir Nanavati, Nathan Taylor, William Aiello, and Andrew Warfield
USENIX Workshop on Hot Topics in Security (HotSec), 2011