

Vinay S Banakar

#2289, Dept. of Computer Science
1210 W Dayton St, Madison, WI 53706

vin@cs.wisc.edu
www.vinaybanakar.com

EDUCATION

University of Wisconsin-Madison
Ph.D in Computer Science
M.S. in Computer Science

Advisors: Andrea Arpaci-Dusseau & Remzi Arpaci-Dusseau
2020 (ongoing)
2020 - 2023

PES University
B.E in Computer Science and Engineering

2013-2017
CGPA: 8.26/10

AREAS OF INTEREST

Operating systems, distributed storage systems, and systems for ML

PUBLICATIONS

- [1] **WiscSort: External Sorting for Byte-Addressable Storage** VLDB'23
Vinay Banakar, Kan Wu, Yuvraj Patel, Kimberly Keeton, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau
- [2] **Understanding and Benchmarking the Impact of GDPR on Database Systems** VLDB'20
Supreeth Shastri, Vinay Banakar, Melissa Wasserman, Arun Kumar, and Vijay Chidambaram
- [3] **CIED - Rapid Composability of Rack Scale Resources Using Capability Inference Engine Across Datacenters** IEEE Infra'20
Vinay Banakar, Pavan Upadhyaya, and Maneesh Keshavan
- [4] **Analyzing the Impact of GDPR on Storage Systems** ACM HotStorage'19
Vinay Banakar, Aashaka Shah, Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram

PATENTS

- [1] **Intent driven hardware placement using rack capability inference engine across datacenters, 2019** US20210014998A1
Vinay Banakar, Pavan Upadhyaya, and Maneesh Keshavan
- [2] **Intelligent orchestration of disaggregated applications based on class of service, 2019** US20200249999A1
Tom Golway, Vinay Banakar, and Sandeep Panda
- [3] **Preemptive compatibility failure detection using graph structure learning in datacenters, 2018** US10938623B2
Vinay Banakar, Pavan Upadhyaya, and Maneesh Keshavan
- [4] **Topology based root cause triangulation of hardware issues** US10831587B2
Pavan Upadhyaya, Maneesh Keshavan, Naveena Kedlaya, and Vinay Banakar

ONGOING WORK

UW Madison Google	Building new data structure paradigms for modern memory systems Data structures that dynamically organizes data based on workload and memory characteristics in tiered environments to improve memory efficiency.
----------------------	---

RECENT PROJECTS

- *OCC hierarchical validation*: Designed a new hierarchical validation scheme for databases that use optimistic concurrency control [Fall'22].
- *Disaggregated-PM aware datastructures*: Explored the performance v/s functionality trade-offs and skewed read/write performance of a disaggregated-PM architecture by implementing a B+-tree and distributed exter-

nal sort over InfiniBand RDMA [Fall'21].

- *RecoverKV*: Strongly consistent, partition tolerant Key-Value store built in Go using quorum protocol [Spring'21].
- *DIAS*: Distributed machine learning Inference as A Service [Spring'21].
- *AutoTune-IO*: Optimizing Linux IO scheduler parameters using Bayesian Optimization [Fall'20]

TEACHING EXPERIENCE

- *CS739*: Distributed Systems TA (UW Madison, Spring 2022)
- *CS537*: Operating Systems TA (UW Madison, Fall 2021)
- *ECE252*: Computer Organization TA (UW Madison, Spring 2021)
- *CS220*: Data programming TA (UW Madison, Fall 2020)

RESEARCH & INDUSTRY EXPERIENCE

Google Visiting Researcher <i>Systems Research Group</i> Summer 2023 - Present	Advisor: Dr. Kimberly Keeton Built optimization techniques to repackage objects to pages for better memory efficiency in Spanner. Currently building next generation tiering aware data structure library for KV and vector databases.
UW Madison Research Assistant Summer 2021 - Present	Advisors: Prof. Andrea Arpaci-Dusseau and Prof. Remzi Arpaci-Dusseau Designed and built <i>WiscSort</i> , a high-performance concurrent sorting system for byte-addressable storage (BAS) that is 7x faster than state-of-the-art. Introduced a new storage model (BRAID) for efficient BAS/CXL device utilization.
Microsoft Research Assistant Spring 2023	Advisor: Dr. Jesús Camacho Rodríguez Focused on Database Redesign for non-Uniform Memory & Storage (DRUMS). Investigated new HTAP database designs for CXL memory and storage devices.
HPE RnD Labs Senior Systems Engineer 2017-2020	Advisors: Dr. Kimberly Keeton and Dr. Sharad Singhal - Developed applications for disaggregated persistent Fabric Attached Memory (<i>openFAM</i>) to evaluate it against traditional cluster-based HPC programming models (<i>openSHMEM</i> and <i>MPI</i>). - Built features for large scale datacenter management software (OneView) and developed a Redfish compliant server hardware simulator.
UT Austin Research Fellow 2018-2019	Advisor: Prof. Vijay Chidambaram - Modified Redis, Postgres and OracleDB to strictly comply with <i>GDPR</i> requirements and found up to 20x drop in throughput. This illustrated how retro fitting existing storage designs to work efficiently with new privacy policies is inadequate. Also built <i>GDPRBench</i> , a <i>GDPR</i> benchmark that let users to assess compliance level of a storage system and helps evaluate compliance-performance tradeoff.
HPE RnD Labs Research intern 2017	Implemented a virtual host simulation platform that mimics ESXi instances as hosts in a VMware vCenter cluster. Scaled up to 1000+ simulated hosts that were leveraged for performance evaluation in OneView.

AWARDS AND ACCOMPLISHMENTS

- *Silver award* for innovation at Hewlett Packard Enterprise, 2020
- First prize in TechnoBiz track at *8th IEEE conference on Cloud Computing in Emerging Markets*, 2019.
- Second place at MyWired Open hack 2016 and SimpliHack 2015.

SERVICE

- Poster presentation: VLDB 2023, Database Affiliates 2022, UW Madison.
- Program committee: *HPE TechCon'20*, *IEEE HiPC'24*, *EuroSys'25 (shadow)*, *IPDPS'25*
- Artifact evaluation review Committee: *ASPLOS 2020*, *SOSP 2019*
- Book Contribution: *Effective Cybersecurity - Understanding and Using Standards and Best Practices*.
- Open source contributions: *Apache Ratis*, *Postgresql* and *YCSB*
- Invited Talks: Virtue insight Blockchain 2019 conference, HPE Technical Symposium'18 and 19, Microsoft Gray Systems Lab 2023, and Google 2023.