

RASHMI VINAYAK

www.cs.cmu.edu/~rvinayak
rvinayak@cs.cmu.edu

ACADEMIC POSITIONS

- **Carnegie Mellon University, Computer Science Department**, Associate Professor, 2023-present.
- **Carnegie Mellon University, Computer Science Department**, Assistant Professor, 2017 - 2023.
- **UC Berkeley**, Post-doctoral researcher in AMPLab & BLISSLab, 2016-17.

EDUCATION

- **University of California at Berkeley, 2011-16**
PhD, Electrical Engineering & Computer Science
- **Indian Institute of Science, Bangalore, India, 2008-10**
Master of Engineering
- **National Institute of Technology Karnataka, Surathkal, India, 2003-07**
Bachelor of Technology

AWARDS AND HONORS

- **USENIX NSDI'24 Community Award** (awarded to 2 papers out of the program of 112 papers).
- **Sloan Research Fellowship 2023**.
- **IEEE Information Theory Society Goldsmith Lecturer 2023** (awarded to one recipient per year).
- **VMware Systems Research Award 2021** (awarded to one recipient per year).
- **USENIX NSDI'21 Community Award** (one of the two best paper awards of the conference).
- Our paper recognized as **one of the best storage-related papers at USENIX OSDI 2020** and **invited to ACM Transactions on Storage**.
- **NSF CAREER Award 2020-25**.
- **Tata Institute of Fundamental Research Memorial Lecture Award 2019** (awarded to one recipient per year).
- PhD thesis awarded **UC Berkeley Eli Jury Award 2016**, given for outstanding achievement in the area of Systems, Communications, Control, or Signal Processing.
- **Rising Stars in EECS 2016**.
- **Google Anita Borg Memorial Scholarship 2015-16**, for impact on diversity, demonstrated leadership and academic background.
- **Microsoft Research PhD Fellowship 2013-15**.
- **Facebook PhD Fellowship 2012-13**.
- **IEEE Data Storage Best Student Paper Award** and **Best Paper Award** for the years 2011/2012.

PUBLICATIONS

Google Scholar profile: [Rashmi Vinayak's Google Scholar profile](#).

CONFERENCE PAPERS

- C1. S. Chopra, F. Maturana and K. V. Rashmi, "On Low Field Size Constructions of Access-Optimal Convertible Codes", in *IEEE International Symposium on Information Theory (ISIT)*, 2024.
- C2. Y. Zhang, J. Yang, Y. Yue, Y. Vigfusson, K. V. Rashmi, "SIEVE: Simple and Efficient Eviction Policy for Turn-key Web Cache Replacement", in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2024.
- C3. A. Park, T. Leong, F. Maturana, W. Zheng, K. V. Rashmi, "Communication-efficient, Fault Tolerant PIR over Erasure Coded Storage", in *IEEE Symposium on Security and Privacy (IEEE S&P)*, 2024.
- C4. J. Yang, Y. Zhang, Z. Qiu, Y. Yue, K. V. Rashmi, "FIFO queues are all you need for cache eviction", in *ACM Symposium on Operating Systems Principles (SOSP) 2023*.
- C5. T. Zhang, K. Liu, J. Kosaian, J. Yang, K. V. Rashmi, "Efficient Fault Tolerance for Recommendation Model Training via Erasure Coding", in *International Conference on Very Large Databases (VLDB)*, 2023.
- C6. M. Rudow, F. Yan, A. Kumar, G. Ananthanarayanan, M. Ellis, K. V. Rashmi, "Tambur: Efficient loss recovery for videoconferencing via streaming codes", in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2023.
- C7. F. Maturana and K. V. Rashmi, "Locally Repairable Convertible Codes: Erasure Codes for Efficient Repair and Conversion", in *IEEE International Symposium on Information Theory (ISIT)*, 2023.
- C8. M. Rudow and K. V. Rashmi, "Learning-Augmented Streaming Codes for Variable-Size Messages Under Partial Burst Losses", in *IEEE International Symposium on Information Theory (ISIT)*, 2023.
- C9. M. Rudow, V. Guruswami, and K. V. Rashmi, "On expanding the toolkit of locality-based coded computation to the coordinates of inputs", in *IEEE International Symposium on Information Theory (ISIT)*, 2023.
- C10. M. Rudow, N. Charalambides, A. O. Hero III, and K. V. Rashmi, "Compression-Informed Coded Computing", in *IEEE International Symposium on Information Theory (ISIT)*, 2023.
- C11. J. Yang, Z. Mao, Y. Yue, and K. V. Rashmi, "GL-Cache: Group-level learning for efficient and high-performance caching", in *USENIX USENIX Conference on File and Storage Technologies (FAST)*, 2023.
- C12. S. Kadekodi, F. Maturana, S. Athlur, A. Merchant, K. V. Rashmi, and G. Ganger, "Tiger: disk-adaptive redundancy without placement restrictions", in *USENIX Operating Systems Design and Implementation (OSDI)*, 2022.
- C13. J. Yang, A. Sabnis, D. S. Berger, K. V. Rashmi, R. K. Sitaraman, "C2DN: How to Harness Erasure Codes at the Edge for Efficient Content Delivery", in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2022.
- C14. M. Rudow and K. V. Rashmi, "Learning-Based Streaming Codes are Approximately Optimal for Variable-Size Messages", in *IEEE International Symposium on Information Theory (ISIT)*, 2022.
- C15. F. Maturana and K. V. Rashmi, "Bandwidth Cost of Code Conversions in the Split Regime", in *IEEE International Symposium on Information Theory (ISIT)*, 2022.
- C16. J. Kosaian, and K. V. Rashmi, "Arithmetic-Intensity-Guided Fault Tolerance for Neural Network Inference on GPUs", in *International conference on high performance computing, networking, storage and analysis (SC)* 2021.

- C17. J. Kosaian, A. Phanishayee, D. Dey, M. Philipose, and K. V. Rashmi, “Boosting the Throughput and Accelerator Utilization of Specialized CNN Inference Beyond Increasing Batch Size”, in *International Conference on Machine Learning (ICML)* 2021.
- C18. J. Yang, Y. Yue, K. V. Rashmi, “Segcache: memory-efficient and high-throughput DRAM cache for small objects”, in *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2021.
- C19. F. Maturana and K. V. Rashmi, “Irregular Array Codes with Arbitrary Access Sets for Geo-Distributed Storage”, in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- C20. M. Rudow, K. V. Rashmi and V. Guruswami, “A locality-based lens for coded computation”, in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- C21. F. Maturana and K. V. Rashmi, “Bandwidth Cost of Code Conversions in Distributed Storage: Fundamental Limits and Optimal Constructions”, in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- C22. S. Kadekodi, F. Maturana, S. J. Subramanya, J. Yang, K. V. Rashmi, and G. Ganger, “Pacemaker: Avoiding HeART attacks in storage clusters with disk-adaptive redundancy”, in *USENIX Operating Systems Design and Implementation (OSDI)*, 2020.
- C23. J. Yang, Y. Yue, and K. V. Rashmi, “A large scale analysis of hundreds of in-memory cache clusters at Twitter”, in *USENIX Operating Systems Design and Implementation (OSDI)*, 2020.
- C24. F. Maturana and K. V. Rashmi, “Convertible Codes: New Class of Codes for Efficient Conversion of Coded Data”, in *Innovations in Theoretical Computer Science (ITCS)* 2020.
- C25. M. Rudow and K. V. Rashmi, “Online Versus Offline Rate in Streaming Codes for Variable-Size Messages”, in *IEEE International Symposium on Information Theory (ISIT)*, 2020.
- C26. F. Maturana, C. Mukka, and K. V. Rashmi, “Access-optimal Linear MDS Convertible Codes for All Parameters”, in *IEEE International Symposium on Information Theory (ISIT)*, 2020.
- C27. J. Kosaian, K. V. Rashmi, and S. Venkataraman, “Parity Models: Erasure-Coded Resilience for Prediction Serving Systems”, in *ACM Symposium on Operating Systems Principles (SOSP)* 2019.
- C28. D. Ray, J. Kosaian, K. V. Rashmi, and S. Seshan, “Optimizing video upload for time-shifted viewing of social live streams”, in *ACM SIGCOMM*, 2019.
- C29. S. Kadekodi, K. V. Rashmi, and G. Ganger, “Cluster storage systems gotta have HeART: improving storage efficiency by exploiting disk-reliability heterogeneity”, in *USENIX Conference on File and Storage Technologies (FAST)*, 2019.
- C30. M. Rudow and K. V. Rashmi, “Streaming Codes for Variable-Size Arrivals”, in *Proceedings of Allerton Conference on Communication, Control, and Computing*, 2018.
- C31. K. V. Rashmi, M. Chowdhury, J. Kosaian, I. Stoica and K. Ramchandran, “EC-Cache: Load-Balanced, Low-Latency Cluster Caching with Online Erasure Coding,” in *USENIX Operating Systems Design and Implementation (OSDI)*, 2016.
- C32. P. Nakkiran, K. V. Rashmi, and K. Ramchandran, “Optimal Systematic Distributed Storage Codes with Fast Encoding,” in *IEEE International Symposium on Information Theory (ISIT)*, 2016.
- C33. K. V. Rashmi, P. Nakkiran, J. Wang, N. Shah, K. Ramchandran, “Having Your Cake and Eating It Too: Jointly Optimal Codes for I/O, Storage and Network-bandwidth In Distributed Storage Systems,” in *USENIX Conference on File And Storage Technologies (FAST)*, 2015.
- C34. K. V. Rashmi, and R. Gilad-Bachrach, “DART: Dropouts meet Multiple Additive Regression Trees,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- C35. K. V. Rashmi, N. Shah, D. Gu, H. Kuang, D. Borthakur and K. Ramchandran, “A ”Hitchhiker’s” Guide to Fast and Efficient Data Reconstruction in Erasure-coded Data Centers,” *ACM SIGCOMM*, 2014.
- C36. N. Shah, K. V. Rashmi, K. Ramchandran, “One Extra Bit of Download Ensures Perfectly Private Information Retrieval,” in *IEEE International Symposium on Information Theory (ISIT)*, 2014.

- C37. P. Nakkiran, N. Shah, K. V. Rashmi, "Fundamental Limits on Communication for Oblivious Updates in Storage Networks", in *IEEE Global Communications Conference (GLOBECOM)*, 2014.
- C38. K. V. Rashmi, N. Shah and K. Ramchandran, "A Piggybacking Design Framework for Read-and Download-efficient Distributed Storage Codes," in *IEEE International Symposium on Information Theory (ISIT)*, 2013.
- C39. N. Shah, K. V. Rashmi, and K. Ramchandran, "Efficient and Distributed Secret Sharing in General Network," in *IEEE International Symposium on Information Theory (ISIT)*, 2013.
- C40. K. V. Rashmi, N. Shah, K. Ramchandran and P. Kumar, "Regenerating Codes for Errors and Erasures in Distributed Storage," in *IEEE International Symposium on Information Theory (ISIT)*, 2012.
- C41. K. V. Rashmi, N. Shah and P. Kumar, "Enabling Node Repair in Any Erasure Code for Distributed Storage," in *IEEE International Symposium on Information Theory (ISIT)*, 2011.
- C42. N. Shah, K. V. Rashmi, and P. Kumar, "Information-theoretically Secure Regenerating Codes for Distributed Storage," in *IEEE Global Communications Conference (GLOBECOM)*, 2011.
- C43. K. V. Rashmi, N. Shah, P. Kumar and K. Ramchandran, "Explicit and Optimal Exact-Regenerating Codes for the Minimum-Bandwidth Point in Distributed Storage," in *IEEE International Symposium on Information Theory (ISIT)*, 2010.
- C44. N. Shah, K. V. Rashmi, and P. Kumar, "A Flexible Class of Regenerating Codes for Distributed Storage," in *IEEE International Symposium on Information Theory (ISIT)*, 2010.
- C45. K. V. Rashmi, N. Shah, P. Kumar and K. Ramchandran, "Explicit construction of optimal exact regenerating codes for distributed storage," in *Allerton Conference on Control, Computing and Communication*, 2009.

JOURNAL PAPERS

- J1. Francisco Maturana and K. V. Rashmi, "Bandwidth Cost of Code Conversions in Distributed Storage: Fundamental Limits and Optimal Constructions", in *IEEE Transactions on Information Theory*, 2023.
- J2. Michael Rudow and K. V. Rashmi, "Online Versus Offline Rate in Streaming Codes for Variable-Size Messages", in *IEEE Transactions on Information Theory*, 2023.
- J3. Francisco Maturana and K. V. Rashmi, "Convertible Codes: Enabling Efficient Conversion of Coded Data in Distributed Storage", in *IEEE Transactions on Information Theory*, 2022.
- J4. Michael Rudow and K. V. Rashmi, "Streaming Codes for Variable-Size Messages", *IEEE Transactions on Information Theory*, 2022.
- J5. Juncheng Yang, Yao Yue, and K. V. Rashmi, "A large scale of analysis of hundreds of in-memory cache clusters at Twitter", in *ACM Transactions on Storage (TOS)*, 2021.
- J6. Jack Kosaian, K. V. Rashmi, and Shivaram Venkataraman, "Learning-Based Coded-Computation," *IEEE Journal on Selected Areas in Information Theory*, March 2020.
- J7. K. V. Rashmi, N. Shah, K. Ramchandran, and P. Kumar, "Information-theoretically Secure Erasure Codes for Distributed Storage," *IEEE Transactions on Information Theory*, Vol. 64, no. 3, pp. 1621 - 1646, Mar. 2018.
- J8. K. V. Rashmi, N. Shah and K. Ramchandran, "A Piggybacking Design Framework for Read-and Download-efficient Distributed Storage Codes," *IEEE Transactions on Information Theory*, vol. 63, no. 9, pp. 5802–5820, Sept. 2017.
- J9. N. B. Shah, K. V. Rashmi and K. Ramchandran, "Distributed Secret Dissemination Across a Network," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 7, pp. 1206-1216, Oct. 2015.
- J10. N. B. Shah, K. V. Rashmi, P. V. Kumar and K. Ramchandran, "Distributed Storage Codes with Repair-by-Transfer and Non-achievability of Interior Points on the Storage-Bandwidth Tradeoff," *IEEE Transactions on Information Theory*, vol. 58, no. 3, 1837 - 1852, Mar. 2012.

- J11. N. B. Shah, K. V. Rashmi, P. V. Kumar and K. Ramchandran, "Interference Alignment in Regenerating Codes for Distributed Storage: Necessity and Code Constructions," *IEEE Transactions on Information Theory*, Apr. 2012.
- J12. K. V. Rashmi, N. B. Shah and P. V. Kumar, "Optimal Exact-Regenerating Codes for the MSR and MBR Points via a Product-Matrix Construction," *IEEE Transactions on Information Theory*, vol. 57, no. 8, pp. 5227 - 5239, Aug. 2011.

WORKSHOP PAPERS

- W1. J. Hu, J. Kosaian, K. V. Rashmi, "Rethinking Erasure-Coding Libraries in the Age of Optimized Machine Learning", in *USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2024.
- W2. J. Yang , Z. Qiu, Y. Zhang, Y. Yue, K. V. Rashmi, "FIFO Can be Better than LRU: the Power of Lazy Promotion and Quick Demotion", in *Workshop on Hot Topics in Operating Systems (HotOS)*, 2023.
- W3. K. Liu, J. Kosaian, K. V. Rashmi, "Erasure-Coding-Based Fault Tolerance for Recommendation Model Training", in *International Symposium on Checkpointing for Supercomputing (SuperCheck-SC 21)*, workshop held in conjunction with ACM SC, 2021.
- W4. K. Liu, J. Kosaian, and K. V. Rashmi, "Erasure-Coding-Based Fault Tolerance for Recommendation Model Training," in *Personalized Recommendation Systems and Algorithms Workshop (PeRSONAI)*, workshop held in conjunction with MLSys 2021.
- W5. M. Rudow, K. V. Rashmi, and V. Guruswami, "Locality driven coded computation", in *Coding Theory and Machine Learning (CodML) workshop at ICML*, 2019.
- W6. K. V. Rashmi, N. Shah, D. Gu, H. Kuang, D. Borthakur and K. Ramchandran, "A Solution to the Network Challenges of Data Recovery in Erasure-coded Distributed Storage Systems: A Study on the Facebook Warehouse Cluster," in *USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2013.

PROFESSIONAL SERVICE

ORGANIZING COMMITTEES

- ACM Symposium on Operating Systems Principles (SOSP) 2023, Student Scholarship Chair.

PROGRAM COMMITTEES

- ACM ACM Symposium on Operating Systems Principles (SOSP) 2024
- USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2024
- IEEE International Symposium on Information Theory (ISIT) 2024
- USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2022
- Conference on Machine Learning and Systems (MLSys) 2022
- IEEE International Symposium on Information Theory (ISIT) 2022
- USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2021
- IEEE International Symposium on Information Theory (ISIT) 2021
- Conference on Machine Learning and Systems (MLSys) 2021
- USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2021 (External PC)
- USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2020
- USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2020
- International Conference on Machine Learning (ICML) 2019 CodML Workshop 2019

- USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2018
- SysML 2018

DIVERSITY AND INCLUSION ACTIVITIES

- **Co-chair, Women in IEEE Information Theory Society (WITHITS), 2022-24.**
- **Panelist, ITALT Professional Development Panel, ITALT (organized by WiML-T and Let-All), 2024.**
- **Speaker, Pittsburgh Women in Mathematics and Computing Symposium (WMCS), 2023**
- **Mentor, ACM SOSP 2023 mentoring program, 2023.**
- **Women in IEEE Information Theory Society (WITHITS), co-host for the coding theory round table at IEEE ISIT 2021.**
- **N2Women Panel, USENIX NSDI, 2021.**
- **Mentor, ACM SOSP 2021 mentoring program, 2021.**
- **Mentor, USENIX OSDI/ATC 2021 mentoring program, 2021.**
- **Committee member on a new course on Diversity, Equity, and Inclusion, Computer Science Department, CMU, 2021.**
- **SCS4All, School of Computer Science, CMU, 2017 - present.**
- **Women@SCS, School of Computer Science, CMU, 2017 - present.**
- **Women in IEEE Information Theory Society (WITHITS), panelist in “women in academia” at Bombay Information Theory Seminar 2020.**
- **Women in Academia, Alpha Chi Omega, CMU, 2020.**
- **Diversity Committee, Electrical Engineering & Computer Science (EECS) Department, UC Berkeley, 2014-15.**
- **Rising Stars in EECS organizer - An academic career workshop for women, UC Berkeley, Volunteer Organizer, 2014.**
- **Undergraduate Mentoring Program, Women in Computer Science and Engineering (WICSE), UC Berkeley, Mentor, 2014-16.**
- **Women in Computer Science and Engineering (WICSE), UC Berkeley, Officer, 2013-14.**
- **Big Sister Mentoring Program, Women in Computer Science and Engineering (WICSE), UC Berkeley, Mentor, 2012-2016.**

OTHER SERVICE

- **NSF proposal review panel, 2023.**
- **NSF proposal review panel, 2021.**
- **Journal reviewing:** ACM Transactions on Computer Systems, ACM Transactions on Storage, IEEE Transactions on Information Theory, IEEE Transactions on Computers, IEEE Communication Letters, IEEE Signal Processing Magazine.

TALKS

- **Invited talks session at International Conference on Signal Processing and Communications (SPCOM), 2024**
“Convertible Codes with Local Repair”
- **PUC Chile Institute of Mathematical and Computational Engineering Seminar, 2024**
“Unlocking the Code: How Math Protects Your Data at Scale”
- **UC Berkeley Skylab Seminar, 2024**
“Co-designing Systems and Algorithms for Efficient and Performant Data Systems”
- **Simons Institute Workshop on Application-Driven Coding Theory at UC Berkeley , 2024**
“Convertible codes: Adaptive coding for large-scale data storage”
- **Information Theory and Applications (ITA) Workshop, 2024**
“Adaptive coding for large-scale data storage: what, why, and how?”
- **UC Santa Cruz ECE Department Seminar, 2023**
“Improving Resource Efficiency and Reliability in Large-scale Storage Systems”
- **IEEE European Summer School on Information Theory, 2023**
“Coding theory for distributed systems”
- **JTG/IEEE Indian Summer School on Information Theory, 2023**
“Coding theory for distributed systems”
- **EURECOM, France, 2023**
“Disk-adaptive Coding for Distributed Storage: Theory and Systems”
- **Information Theory and Applications (ITA) Workshop, 2023**
“Disk-adaptive Coding for Distributed Storage: Theory and Systems”
- **Microsoft Research India, 2023**
“DARE: Disk-Adaptive Redundancy for Improving Efficiency in Storage Systems”
- **VMware Company-Wide R&D Innovation Offsite (RADIO) Seminar, 2022**
“Efficient, Performant, and Resilient Data Storage and Caching Systems”
- **VMware Data Storage Team Seminar, 2022**
“Efficient, Performant, and Resilient Data Storage and Caching Systems”
- **Stanford Information Theory Forum, 2021**
“Convertible Codes: Enabling Redundancy Tuning in Large-scale Storage Systems”
- **Rutgers Signal and Information Processing Seminar Series, 2021**
“Convertible Codes: Enabling Redundancy Tuning in Large-scale Storage Systems”
- **American Mathematical Society Fall Central Sectional Meeting, 2021**
“Convertible Codes: Efficient Conversion of Coded Data for Large-scale Storage Systems”
- **Information Theory Workshop (ITW) Special Session on Frontiers of Coding Theory and Practice, 2021**
“Convertible Codes: Enabling Efficient Conversion of Coded Data in Distributed Storage”
- **UW Madison Systems Information Learning and Optimization (SILO) Seminar, 2021**
“Convertible Codes: Efficient Conversion of Coded Data in Large-scale Storage Systems”
- **Stanford Compression Workshop, 2021**
“Learning-Based Coded-Computation”
- **Facebook Distributed Systems Faculty Summit, 2020**
“Resource-efficient cluster storage by exploiting disk-reliability heterogeneity”
- **IEEE International Symposium on Information Theory (ISIT) 2020 Live Panel Session on “Machine-learning based approaches to coding”, 2020**
“Learning-based approaches to coded computation”

- **Shannon Channel Seminar**, 2020
“Convertible Codes: A New Class of Codes for Efficient Conversion of Coded Data in Distributed Storage”
- **Information Theory and Applications (ITA) Workshop**, 2020
“A locality based approach for coded computation”
- **BITS(Bombay Information Theory Seminar) Workshop Tutorial**, 2020
“Resilient and Efficient Distributed Storage and Computation via Coding Theoretic Tools”
- **Tata Institute of Fundamental Research**, 2020
“Convertible Codes: A New Class of Codes for Efficient Conversion of Coded Data in Distributed Storage”
- **AI Systems Workshop at ACM SOSP**, 2019
“Learning based coded-computation: A novel approach for resilient computation in ML inference systems”
- **Microsoft Research Redmond**, 2019
“Resource-efficient redundancy for large-scale data processing and storage systems”
- **Facebook Communications and Networking Faculty Summit**, 2019
“Vantage: Optimizing Video Upload for Time-shifted Viewing of Social Livestreams”
- **ICML Workshop on Coding Theory for Large-Scale Machine Learning**, 2019
“Coded-Computation for ML Inference: Learning-based approach”
- **Information Theory and Applications (ITA) Workshop**, 2019
- **Indian Institute of Science**, 2019
“Smart Redundancy for Big data Systems: Theory and Practice”
- **Microsoft Research**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **Princeton University**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **Carnegie Mellon University**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **Cornell University**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **Massachusetts Institute of Technology**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **University of Illinois Urbana Champaign**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **University of Pennsylvania**, 2017
“Smart Redundancy for Big data Systems: Theory and Practice”
- **University of Southern California**, 2017
“Smart redundancy for big-data systems: Theory and Practice”
- **Stanford Information Theory Forum**, Oct. 2016
“Erasure coding for big-data systems: Theory and Practice”
- **Alluxio Inc.**, Sept. 2016
“EC-Cache: Load-Balanced, Low-Latency Cluster Caching with Online Erasure Coding”
- **Cisco**, July 2016
“Erasure coding for next-generation distributed storage systems”
- **AMPLab Berkeley Retreat**, June 2016
“EC-Cache: Load-balanced, Low-latency Cluster Caching with Online Erasure Coding”

- **Information Theory and Applications (ITA) workshop**, Feb. 2016
“A Hitchhiker’s Guide to Resource-Efficient Fault Tolerance in Data Centers: Theory & Practice”
- **Allerton conference on Communication, Control, and Computing, Special session on coding theory**, Oct. 2015
“Piggybacking for Fast and Efficient Data Reconstruction in Erasure-Coded Data Centers”
- **Google**, June 2015
“A Hitchhiker’s Guide to Fast and Efficient Data Reconstruction in Erasure-coded Data Centers”
- **AMPLab Retreat**, Jan. 2015
“Hitchhiker: Efficient Erasure Coding for Data Centers”
- **NetApp**, Oct. 2014
“Piggybacking and Hitchhiker: Retaining the Angels but not the Demons of Reed-Solomon”
- **Facebook**, Feb. 2012
“Erasure coding for distributed storage systems”

TEACHING

CARNEGIE MELLON UNIVERSITY

- **Distributed Systems**, Spring 2024, Undergraduate course
- **Algorithms in the Real World**, Fall 2023, Graduate course
- **Algorithms in the Real World**, Spring 2022, Graduate course
- **Distributed Systems**, Fall 2021, Undergraduate course
- **Graduate Algorithms**, Spring 2021, Graduate course
- **Distributed Systems**, Fall 2020, Undergraduate course
- **Distributed Systems**, Spring 2020, Undergraduate course
- **Algorithms in the Real World**, Fall 2019, Graduate course
- **Practical information and coding theory for computer systems**, Fall 2018, Graduate course
- **Probability and computing**, Spring 2018, Undergraduate course

UC BERKELEY (Graduate Student Instructor)

- **Random Processes in Systems**, Fall 2015, Graduate course
- **Coding Theory for Communication and Beyond**, Fall 2013, Undergraduate course

STUDENTS

CURRENT PHD STUDENTS

- Juncheng Yang
- Sanjith Athlur (co-advised with Greg Ganger)
- Timothy Kim (co-advised with Greg Ganger)
- Yixuan Mei

GRADUATED PHD STUDENTS

- Saurabh Kadekodi, co-advised with Prof. Greg Ganger, 2020 (First employment: Postdoctoral researcher at Google Research)
- Jack Kosaian, 2022 (First employment: Nvidia)

- Michael Rudow, 2023 (First employment: McKinsey & Company)
- Francisco Maturana, 2023 (First employment: Amazon Web Services)

MASTERS STUDENTS

- Saransh Chopra, 5th year Master's program in CS, 2023-present
- Justin Zhang, 5th year Master's program in CS, 2023-present
- Jiyu Hu, Independent Study, 2023
- Timothy Kim, Independent Study, 2022
- Abhishek Kumar, Independent Study, 2022
- Shaobo Guan, Master's Capstone project, 2021
- Kaige Liu , 5th year Master's program in CS, 2020
- Arvind Sai Krishnan, Master's Capstone project, 2020
- Vilas Bhat, Master's Capstone project, 2020
- Jiaan Dai, Master's Capstone project, 2019
- Jiaqi Zuo, Master's Capstone project, 2019
- Jiongtao Ye, Master's Capstone project, 2019
- Sai Kiriti Badam, Master's Capstone project, 2019
- Xuren Zhou, Master's Capstone project, 2019

UNDERGRADUATE STUDENTS

- Helen Wang (CMU), 2024-present
- Emily Zhang (CMU), Summer 2023
- Qinghan Chen (CMU), 2022-present
- Zhuofan Chen (CMU), 2022-present
- Jonathan Chiu (CMU), 2021-2022
- Justin Zhang (CMU), 2022-23
- Tianyu Zhang (CMU), 2022-23
- Ziming Mao (Yale), 2022
- Ian Chiu (CMU), 2019-2020
- Chaitanya Mukka (IISc), Fall 2019
- Sanya Agarwal (CMU), Summer 2019
- Weizhong Zhang, Summer 2018
- Eliot Robson (CMU), Spring 2018
- Jingyan Wang (UC Berkeley), 2014-2015
- Preetum Nakkiran (UC Berkeley), 2014-2016

MEDIA COVERAGE

- SIEVE (USENIX NSDI 2024) featured on:
 - The Hacker News
 - Several other news outlets: TL;DR, Digital Daze, LiveScience, Tech Explorist, BNN Speaking
 - Multiple popular blogposts including Marc Brooker's blog post: "Why Aren't We SIEVE-ing?"
- S3-FIFO (ACM SOSP 2023) featured on:
 - Pointer newsletter, a popular engineering newsletter.
 - Multiple other newsletters around the globe in various languages such as Japanese, Korean, and Russian
- "SCS Faculty Earn 2023 Sloan Research Fellowships", CMU SCS News, Feb. 2023.
- "Vinayak receives Meta Research Award", CMU SCS News, Sept. 2022.
- "Vinayak Receives VMware Systems Research Award", CMU SCS News, Jan 10 2022.
- "Research by CMU, Twitter Could Improve Cache Efficiency by 60%. Team Wins Top Paper Award at USENIX NSDI Conference." CMU SCS News, May 11 2021.
- "Segcache: a memory-efficient, scalable cache for small objects with TTL. In collaboration with Carnegie Mellon University, Twitter is building the next generation of storage backend, Segcache, into Pelikan. Segcache enables high memory efficiency, high throughput, and excellent scalability for Twitter's cache workloads." Twitter's Caching Blog Pelikan Cache, April 13 2021.
- "Rashmi Vinayak Receives Facebook Distributed Systems Research Award. One Of Just Eight Researchers To Receive An Award." CMU SCS News, April 2 2020.
- "CSD's Vinayak Wins NSF CAREER Award. Research Aims To Improve Efficiency of Large-Scale Data Centers." CMU SCS News March 4 2020.
- HeART (USENIX FAST 2019), a space-efficient redundancy management solution for large-scale storage systems, featured as a top pick from FAST 2019 on Digital Preservation blog.
- USENIX FAST 2015 paper on IO, bandwidth, and storage optimal codes and the associated systems solution, picked as the best paper of FAST 2015 by StorageMojo, a popular storage systems blog.
- Hitchhiker (ACM SIGCOMM 2014) featured on:
 - The Morning Paper, a popular blog on CS literature.
 - StorageMojo, a popular storage systems blog.