RASHMI VINAYAK

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

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

PROFESSIONAL EXPERIENCE

- **Carnegie Mellon University, Computer Science Department**, Assistant Professor, 2017-present.
- **UC Berkeley, AMPLab** & **BLISSLab, Berkeley**, Post-doctoral researcher, 2016-17.
- **Microsoft Research, Redmond**, Research Intern, May-August 2014.
- **Microsoft Research, Redmond**, Research Intern, May 2013 - August 2013.
- **Facebook Inc**, Menlo Park, Software Engineering Intern, May 2012 - August 2012.
- **Indian Institute of Science**, Bangalore, Research Associate, July 2010 - June 2011.
- **Nvidia Graphics**, Bangalore, ASIC Designer, July 2007 - July 2008.
- **Goldman Sachs**, Bangalore, Summer Intern, May-June 2006.

AWARDS AND HONORS

- **USENIX NSDI'21 Community (Best Paper) Award**, one of the two best paper awards of the conference.
- Our paper recognized by the USENIX OSDI 2020 PC chairs as **one of the best storage-related papers at OSDI 2020** and **invited to ACM Transactions on Storage**.
- **NSF CAREER Award** 2020.
- **Tata Institute of Fundamental Research Memorial Lecture Award** 2020.
- **Facebook Distributed Systems Research Award 2019** (1 out of the 8 awardees).
- **Google Research Award** 2019.
- **NSF CISE Research Initiation Initiative (CRII) Award** 2019.
- **Facebook Communications and Networking Research Award 2018** (1 out of the 7 awardees).
- PhD thesis awarded **UC Berkeley Eli Jury Award 2016**, for outstanding achievement in the area of Systems, Communications, Control, or Signal Processing.
- **Information Theory and Applications (ITA) Graduation Day speaker,** 2016.
- **Rising Stars in EECS**, 2016.
- **Google Anita Borg Memorial Scholarship** 2015-16.
- **Microsoft Research PhD Fellowship** 2013-15.
- **Facebook 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.

REFEREED CONFERENCE/WORKSHOP PAPERS

- Jack 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.
- Jack Kosaian, Amar Phanishayee, Debadeepta Dey, Matthai Philipose, and K. V. Rashmi, "Boosting the Throughput and Accelerator Utilization of Specialized CNNInference Beyond Increasing Batch Size", in *International Conference on Machine Learning (ICML)* 2021.
- Juncheng Yang, Yao 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.
- Francisco 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.
- Michael Rudow, K. V. Rashmi and Venkatesan Guruswami, "A locality-based lens for coded computation", in *IEEE International Symposium on Information Theory (ISIT)*, 2021.
- Francisco 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.
- Saurabh Kadekodi, Francisco Maturana, Suhas Jayaram Subramanya, Juncheng Yang, K. V. Rashmi, and Greg Ganger, "Pacemaker: Avoiding HeART attacks in storage clusters with disk-adaptive redundancy", in *USENIX Operating Systems Design and Implementation (OSDI)*, 2020.
- Juncheng Yang, Yao Yue, and K. V. Rashmi, "A large scale of analysis of hundreds of in-memory cache clusters at Twitter", in *USENIX Operating Systems Design and Implementation (OSDI)*, 2020.
- Michael 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.
- Francisco Maturana, Chaitanya Mukka, and K. V. Rashmi, "Access-optimal Linear MDS Convertible Codes for All Parameters", in *IEEE International Symposium on Information Theory (ISIT)*, 2020.
- Francisco 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.*
- Jack 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.
- Devdeep Ray, Jack Kosaian, K. V. Rashmi, and Srini Seshan, "Optimizing video upload for time-shifted viewing of social live streams", in *ACM SIGCOMM*, 2019.
- Saurabh Kadekodi, K. V. Rashmi, and Greg 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.
- Michael Rudow and K. V. Rashmi, "Streaming Codes for Variable-Size Arrivals", in *P*roceedings of Allerton Conference on Communication, Control, and Computing, 2018
- 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.
- 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.

- 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.
- K. V. Rashmi, and R. Gilad-Bachrach, "DART: Dropouts meet Multiple Additive Regression Trees," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- 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.
- 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.
- P. Nakkiran, N. Shah, K. V. Rashmi, "Fundamental Limits on Communication for Oblivious Updates in Storage Networks", in *IEEE Global Communications Conference (GLOBECOM)*, 2014.
- 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.
- 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.
- 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.
- 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.
- 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.
- N. Shah, K. V. Rashmi, and P. Kumar, "Information-theoretically Secure Regenerating Codes for Distributed Storage," in *IEEE Global Communications Conference (GLOBECOM)*, 2011.
- 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.
- 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.
- N. Shah, K. V. Rashmi, P. Kumar and K. Ramchandran,"Explicit codes minimizing repair bandwidth for distributed storage," in *IEEE Information Theory Workshop (ITW)*, 2010.
- 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.

REFEREED JOURNAL PAPERS

- 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 (to appear).
- Jack Kosaian, K. V. Rashmi, and Shivaram Venkataraman, "Learning-Based Coded-Computation," *IEEE Journal on Selected Areas in Information Theory, March 2020.*
- 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.*

- 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.*
- 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.
- 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.
- 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.
- 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.

PROFESSIONAL SERVICE

PROGRAM COMMITTEES

- **USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2021**, Program Committee member.
- **IEEE International Symposium on Information Theory (ISIT) 2021**, Program Committee member.
- **Conference on Machine Learning and Systems (MLSys) 2021**, Program Committee member.
- **USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2021**, External Program Committee member.
- **USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2020**, Program Committee member.
- **USENIX Symposium on Networked Systems Design and Implementation (NSDI) 2020**, Program Committee member.
- **International Conference on Machine Learning (ICML) 2019 CodML Workshop 2019**, Program Committee member.
- **USENIX Symposium on Operating Systems Design and Implementation (OSDI) 2018**, Program Committee member.
- **SysML 2018**, Program Committee member.

DIVERSITY, EQUITY, & INCLUSION, AND MENTORING

- **Women in IEEE Information Theory Society (WITHITS)**, 2017 - present.
- 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 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**, 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 Maganize.

MEMBERSHIPS IN PROFESSIONAL SOCIETIES

- ACM, USENIX, IEEE Information Theory Society

INVITED TALKS

- **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"
- **AISystems 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 approac"
- **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 Pennslyvania**, 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

CMU

- **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

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