

# RASHMI K. VINAYAK

rvinayak@cs.cmu.edu

## Academic Positions

---

- **Carnegie Mellon University, Aug. 2017 - present**  
Assistant Professor, Computer Science Department
- **University of California at Berkeley, Sep. 2016 - July 2017**  
Postdoctoral Researcher, AMPLab and BLISS, UC Berkeley  
Advisors: Prof. Ion Stoica and Prof. Kannan Ramchandran

## Education

---

- **University of California at Berkeley, Aug. 2011 - Sep. 2016**  
PhD, Electrical Engineering & Computer Science  
Dissertation title: Erasure Coding for Big-data Systems: Theory and Practice  
Advisor: Prof. Kannan Ramchandran
- **Indian Institute of Science, Bangalore, India, 2008-10**  
Master of Engineering  
Advisor: Prof. P. Vijay Kumar
- **National Institute of Technology Karnataka, Surathkal, India, 2003-07**  
Bachelor of Technology

## Awards and Honors

---

- **Eli Jury Award** for outstanding achievement in the area of systems, communications, control, or signal processing, EECS department, UC Berkeley 2015-16.
- **Google Anita Borg Memorial Scholarship** 2015-16.
- **Microsoft Research PhD Fellowship** 2013-15.
- **Facebook Fellowship** 2012-13.
- **IEEE Data Storage Best Paper Award** and **Best Student Paper Award** for the years 2011/2012.

## Publications

---

### Peer-reviewed Journal Publications

- **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*, June 2017.
- **K. V. Rashmi**, N. Shah, K. Ramchandran, and P. Kumar, "Information-theoretically Secure Erasure Codes for Distributed Storage," *IEEE Transactions on Information Theory* (accepted with minor revisions).
- **K. V. Rashmi**, N. Shah and P. Kumar, "Optimal Exact-Regenerating Codes for the MSR and MBR Points via a Product-Matrix Construction," *IEEE Transactions on Information Theory*, 2011.  
**IEEE Data Storage Best Paper Award and Best Student Paper Award for the years 2011/12.**

- N. Shah, **K. V. Rashmi**, and K. Ramchandran, "Distributed Secret Dissemination Across a Network," *IEEE Journal of Selected Topics in Signal Processing*, 2015.
- N. Shah, **K. V. Rashmi**, P. 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*, 2012.
- N. Shah, **K. V. Rashmi**, P. Kumar and K. Ramchandran, "Interference Alignment in Regenerating Codes for Distributed Storage: Necessity and Code Constructions," *IEEE Transactions on Information Theory*, 2012.

## Peer-reviewed Conference Publications

- **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 *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 *ISIT*, 2013.
- N. Shah, **K. V. Rashmi**, and K. Ramchandran, "Efficient and Distributed Secret Sharing in General Network," in *ISIT*, 2013.
- **K. V. Rashmi**, N. Shah, K. Ramchandran and P. Kumar, "Regenerating Codes for Errors and Erasures in Distributed Storage," in *ISIT*, 2012.
- **K. V. Rashmi**, N. Shah and P. Kumar, "Enabling Node Repair in Any Erasure Code for Distributed Storage," in *ISIT*, 2011.
- N. Shah, **K. V. Rashmi**, and P. Kumar, "Information-theoretically Secure Regenerating Codes for Distributed Storage," in *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 *ISIT*, 2010.
- N. Shah, **K. V. Rashmi**, and P. Kumar, "A Flexible Class of Regenerating Codes for Distributed Storage," in *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*, 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.

## Invited Talks

---

- **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 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**, 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”
- **Paradise lab, Caltech**, May 2012  
“Codes for Distributed Storage”
- **MURI meeting, UC Berkeley**, Feb. 2012  
“Codes for Distributed Storage”
- **Facebook**, Feb. 2012  
“Erasure coding for distributed storage systems”

## Mentorship

---

- **Jack Kosian** (University of Michigan Ann Arbor undergraduate → PhD student at CMU)  
Systems deploying erasure codes
- **Preetum Nakkiran** (UC Berkeley undergraduate → PhD student at Harvard)  
Sparse and resource-efficient storage codes.
- **Jingyan Wang** (UC Berkeley undergraduate → PhD student at CMU)  
I/O-efficient distributed storage codes and systems.
- **Zihan Liu** (PhD student at UC Berkeley)  
Reducing latency of computations in big-data systems using structured redundancy.
- **Avishek Ghosh** (PhD student at UC Berkeley)  
Coding in the caching layer of big-data systems.

## Professional Service

---

- **Reviewer** for IEEE Transactions on Information Theory, IEEE Transactions on Parallel and Distributed Systems, IEEE Journal of Selected Topics in Signal Processing, IEEE Transactions on Communications, IEEE Transactions on Computers, IEEE Signal Processing Magazine, IEEE Communication Letters, International Symposium on Information Theory, Globecom, Information Theory Workshop.
- **Mentor** for female graduate and undergraduate students, **Women in Computer Science and Electrical Engineering** (WICSE), UC Berkeley, 2012-16.

## Industry

---

- Research Intern, **Microsoft Research, Redmond**, May-Aug 2014
  - Improving accuracy of the multiple additive regression trees (MART) algorithm for regression, classification, and ranking.
- Research Intern, **Microsoft Research, Redmond**, May-Aug 2013
  - Ultra-large-scale K-means clustering on GPU.
- Software Engineering Intern, **Facebook Inc., Menlo Park**, May-Aug 2012
  - Erasure coding in Facebook's HDFS (Hadoop Distributed File System) clusters.
- Research Associate, **Indian Institute of Science, Bangalore**, July 2010 - June 2011
  - Reliable and efficient media distribution and retrieval in content distribution networks.
- ASIC Designer, **Nvidia Graphics Pvt. Ltd., Bangalore**, July 2007 - July 2008
  - Design of high performance computer chip-sets for gaming.

## Teaching

---

- **Graduate Student Instructor, Random Processes in Systems, Fall 2015**
  - Graduate course taught by Prof. Thomas Courtade.
- **Graduate Student Instructor, Coding Theory for Communication and Beyond, Fall 2013**
  - Undergraduate course taught by Prof. Anant Sahai.

## References

---

**Kannan Ramchandran**

Professor  
EECS Department  
UC Berkeley

**P. Vijay Kumar**

Professor  
ECE Department  
Indian Institute of Science, and  
Adjunct Professor  
University of Southern California

**Ran Gilad-Bachrach**

Researcher  
Machine Learning Group  
Microsoft Research Redmond and Israel

**Ion Stoica**

Professor  
Computer Science Division  
UC Berkeley

**Randy Katz**

The UMC Distinguished Professor  
EECS Department  
UC Berkeley