

Mor Harchol-Balter

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
harchol@cs.cmu.edu; (412) 268-7893
www.cs.cmu.edu/~harchol

RESEARCH AREAS

Design and performance analysis of computer systems including both theory and implementation:

Algorithmic Work: Designing and analyzing algorithms for: resource allocation, task/job scheduling, load sharing, routing, cycle stealing, replication, power-management, multi-class/multi-server scheduling, multi-core parallel scheduling, fairness. Correlated arrival processes, and analysis under high-variability workloads. “*All-Can-Win Theorem*”, demonstrating that scheduling policies which are biased towards favoring small jobs can also be preferable to large jobs.

Stochastic Analysis & Queueing Techniques: Developing new methods for stochastic analysis. Examples include: (i) *Recursive Dimensionality Reduction (RDR)*, a technique to reduce a multi-dimensional Markov chain to a 1-dimensional chain, by using busy period transitions; (ii) *Recursive Renewal Reward (RRR)* and *Clearing Analysis on Phases (CAP)*, techniques that allow one to obtain closed-form solutions for many one-dimensional infinite repeating Markov chains, including the M/M/k with setup chain; (iii) *Exact analysis of Replication Systems*: the first exact solution (in product-form) for replication systems, involving any number of servers, and number of classes, and any degree of replication; (iv) *SOAP Analysis of Scheduling Policies*, the first exact response time analysis of a huge class of scheduling policies with no prior analysis, including the Gittins Index; (v) *Optimal Scheduling for Multi-server systems*, the first algorithms and analysis for optimal scheduling in the M/G/k, both with known and unknown job sizes; (vi) *First queueing analysis of Multiserver-Jobs*; (vii) *First queueing analysis of Parallel Speedup Jobs*.

Systems implementation: Resource management/scheduling in data centers and supercomputing centers. Implementation of QoS tail latency guarantees for flows in networked storage systems (PriorityMeister and SNC-Meister). Online scheduling of jobs with heterogeneous preferences (TetriSched). Dynamic power management in multi-tier data centers (AutoScale), adopted by Facebook. Scheduling of Memory Controllers (ATLAS). Pricing and queueing optimizations for cloud services. Kernel-level SRPT connection scheduling in Web servers (SyNC). Scheduling to cope with transient overload in Web servers.

Modeling and Workload characterization: Known for discovery of Pareto heavy-tailed distribution of UNIX process CPU lifetimes. Statistical characterization of workloads including UNIX processes, web, OLTP, supercomputing, memory workloads, datacenter workloads, and parallel multi-core jobs.

EMPLOYMENT

2020 - pres Bruce J. Nelson Professor in Computer Science. Computer Science Department. Carnegie Mellon University.

2014 - pres Full Professor. Computer Science Department. Carnegie Mellon University.

2012 - 2014 Associate Professor. Computer Science Department. Carnegie Mellon University.

2008 - 2011 Associate Professor and Associate Department Head.
Computer Science Department, Carnegie Mellon University.

2005 - 2008 Associate Professor without tenure. Computer Science Department. Carnegie Mellon University.

- 1999-2005 McCandless Assistant Professor. Computer Science Department. Carnegie Mellon University.
- 1996-1999 NSF Postdoctoral Fellow. Massachusetts Institute of Technology.
Funded by: NSF Postdoctoral Fellowship in the Mathematical Sciences.
Mentor: Prof. Tom Leighton.

EDUCATION

- 1990-1996 University of California, Berkeley.
Ph.D. in Computer Science, August 1996.
Thesis: *Network Analysis without Exponentiality Assumptions*.
Advisor: Prof. Manuel Blum. Committee: Manuel Blum, Sheldon Ross, Venkat Anantharam.
- 1984-1988 Brandeis University, Waltham, Massachusetts.
B.A. in Computer Science and Mathematics, May 1988. Summa Cum Laude.
Thesis: *Evaluation of Data Encryption Standard Using Walsh-Hadamard Matrix Factorization*.
Advisor: Prof. Martin Cohn.

HONORS

- 2023 Keynote Speaker: Queueing Theory for Emerging Classical & Quantum Systems (QTECQS 2023)
- 2023 APS Tutorial Speaker at INFORMS Annual Conference.
- 2023 Elected SIG Chair for ACM SIGMETRICS (4 year position).
- 2023 SNAPP Seminar Speaker.
- 2022 INFORMS 2022 George Nicholson Best Student-led Paper in Operations Research this year.
- 2022 Keynote Speaker: International Teletraffic Congress (ITC 2022)
- 2022 Keynote Speaker: Matrix Analytic Methods Conference (MAM 2022)
- 2021 Keynote Speaker: Ubiquitous Computing, Electronics & Mobile Communications (UEMCON 2021)
- 2021 Distinguished Lecture Series Speaker: Penn State, Computer Science Department.
- 2021 Keynote Speaker: Data-Driven Queueing Challenges (DDQC 2021)
- 2021 ACM SIGMETRICS 2021 Best Paper Award.
- 2021 Second Inaugural Speaker: Scheduling Seminar
- 2021 Plenary Speaker: STOC Theory Fest Short Plenary Pick, STOC 2021.
- 2020 Google Faculty Award.
- 2020 ACM SIGMETRICS 2020 Best Video Award.
- 2020 Bruce J. Nelson Endowed Chair in Computer Science
- 2019 ACM SIGMETRICS 2019 Best Student Paper Award.
- 2019 Keynote Speaker: Queueing Theory and Network Applications (QTNA 2019).
- 2019 Ruth and Joel Spira Outstanding Teaching Award.
- 2019 Fellow of IEEE.
- 2018 IFIP PERFORMANCE 2018 Best Student Paper Award.
- 2018 Keynote Speaker: Young European Queueing Theory Conference (YEQT 2018).
- 2018 APS Best Student Paper Award Finalist.
- 2018 Microsoft Faculty Award.
- 2018 Distinguished Lecture Series Speaker: UCSD, Computer Science Department.
- 2018 Distinguished Speaker: Information Modeling and Control of Complex Systems Workshop (IMaCCS 2018)
- 2018 Plenary speaker at Dutch Queueing Colloquium 40th Anniversary.
- 2018 Fellow of ACM.
- 2017 Keynote Speaker: MIT LIDS Student Conference.

2016 Keynote Speaker: CanQueue 2016.
 2016 Keynote Speaker: ACM SIGMETRICS 2016.
 2016 Distinguished Speaker: Information Modeling and Control of Complex Systems Workshop (IMaCCS 2016)
 2016 EuroSys Best Student Paper.
 2011-16 Intel Faculty Gift every year.
 2015 Google Faculty Award.
 2015 Facebook Faculty Award.
 2015 Keynote Speaker: IEEE ICDCS 2015.
 2014 Keynote Speaker: GreenMetrics 2014.
 2014 CMU Mudge House Dinner with the Deans Honorary Event for Influential Teachers.
 2014 Apple Pie with Alpha Chi Honorary Event for CMU Faculty with Impact on Students.
 2013 Distinguished Speaker Series, Northwestern University.
 2013 General Chair ACM SIGMETRICS 2013.
 2012 Microsoft Faculty Award.
 2012 Invited Speaker: Brandeis University Distinguished High-Tech Alumni Event.
 2011 Microsoft Faculty Award.
 2011 Distinguished Speaker at Opening of Center Computer Engineering, Technion, Israel.
 2011 Kaufman Foundation Faculty Award.
 2010 Microsoft Faculty Award.
 2010 Google Faculty Award.
 2010 Yahoo! Faculty Award.
 2010 IEEE MICRO's Top Picks from 2010 award paper.
 2010 Distinguished Speaker Series, University of Toronto.
 2008-09 Intel Faculty Gift each year.
 2009 Invited Speaker: Columbia Applied Probability Day 2009.
 2009 Yahoo! Faculty Award.
 2008 Keynote Speaker: IEEE MASCOTS 2008.
 2008 Keynote Speaker: SIPEW 2008.
 2008 IBM Faculty Award.
 2007 Microsoft Faculty Award.
 2007 Program co-Chair ACM SIGMETRICS 2007.
 2007 Program co-Chair QEST 2007.
 2007 Plenary Speaker: Annual Lunteren Conference 2007.
 2007 Keynote Speaker: ORSIS 2007.
 2006 Plenary Speaker: Carnegie Mellon 50 Year Anniversary Symposium.
 2005 Invited to IFIP 7.3 WG on Computer Performance Modeling and Analysis.
 2001-04 IBM Faculty Award each year.
 2003 Herbert A. Simon Award for Teaching Excellence, Carnegie Mellon University.
 2003 ACM SIGMETRICS Best Student Paper Award.
 2003 IEEE International Teletraffic Congress Best Student Paper Award.
 2001 NSF CAREER Award.
 2001 EMC Faculty Award.
 2001-04 Anna McCandless Chair for Career Development, Carnegie Mellon University.
 1996-99 NSF Mathematical Sciences Postdoctoral Fellowship.
 1996 ACM SIGMETRICS Best Paper Award for Integrating Systems and Theory.
 1990-96 6-year National Physical Science Consortium Fellowship for Women and Minorities.
 1995 Demitri Angelakos Altruism Award, U.C. Berkeley.
 1994 Outstanding Graduate Student Instructor Award, Campus-Wide, U.C. Berkeley.
 1994 EECS Graduate Student Instructor of the Year Award, U.C. Berkeley.
 1994 Teaching Effectiveness Award, Campus-Wide, U.C. Berkeley.
 1990 Awarded National Science Foundation Graduate Fellowship.
 1990 Awarded Office of Naval Research Graduate Fellowship.

- 1988 Mitchom Prize in Computer Science, Brandeis University.
- 1988 Graduated Summa Cum Laude, Brandeis University.
- 1988 Undergraduate Thesis Highest Honors, Brandeis University.
- 1987 Phi Beta Kappa.
- 1984-88 Brandeis University 4-year Merit Scholarship.
- 1984-88 Lulu T. Briggs 4-year Scholarship.
- 1984 Valedictorian South Brunswick High School, New Jersey.

PUBLICATIONS

Books

- [24] Mor Harchol-Balter. *Introduction to Probability for Computing*. Cambridge University Press. 2024. *Based on 600 pages of lecture notes from 15-259 class at CMU.*
- [13] Mor Harchol-Balter. *Performance Analysis and Design of Computer Systems*. Cambridge University Press. 2013. *Based on 1000 pages of lecture notes from 15-849a,b and 15-857 classes at CMU.*

Thesis

- [96] Mor Harchol-Balter. *Network Analysis without Exponentiality Assumptions*. Ph.D. Thesis. University of California, Berkeley, December 1996.

Guest Editor

- [07] Mor Harchol-Balter. “New Perspectives on Scheduling.” *ACM Performance Evaluation Review, Special Issue*, vol. 34, no. 4, March 2007. Guest editor: Mor Harchol-Balter.

Chapters in Books

- [21] Ben Berg, Mor Harchol-Balter. “Optimal Scheduling of Parallel Jobs with Unknown Service Requirements.” *Handbook of Research on Methodologies and Applications of Supercomputing*, 2021. pp. 18-40.
- [99] Mor Harchol-Balter and Allen Downey. “Exploiting Process Lifetime Distributions for Load Balancing.” Book chapter in *Mobility: Processes, Computers and Agents (ACM Press Series)*, edited by Dejan Milojevic, Fred Douglass, and Richard Wheeler. Addison-Wesley Publishing Company and the ACM Press, April 1999, pp. 229-259.
- [97] Earlier version appeared in: *ACM Transactions on Computer Systems*, vol. 15, no. 3, August 1997, pp. 253-285.
- [96] Earlier version appeared in: *Proceedings of ACM SIGMETRICS 1996 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 96)*, May 1996, Philadelphia, PA, pp. 13-24. Winner of **Best Paper Award for Integrating Systems and Theory**.
- [95] Much earlier version appeared in: *Fifteenth ACM Symposium on Operating Systems Principles (SOSP) Poster Session*, Copper Mountain, CO, December 1995.

Refereed Journal Papers

- [23] Isaac Grosf, Yige Hong, Mor Harchol-Balter, Alan Scheller-Wolf. “The RESET and MARC Techniques, with Applications to Multiserver-Job Analysis.” In *Performance Evaluation*, vol. FILL , 2023, pp. FILL .
- [23] Conference version appeared in: *38th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (IFIP Performance 2023)*, Evanston, IL, November 2023.
- [23] Jalani Williams, Mor Harchol-Balter, and Weina Wang. “M/M/k with Deterministic Setup Times.” Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 6, Issue 3, Article 56 (Dec 2022), pp. 56:1 – 56:33.
- [23] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2023 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 23)*, Orlando, Florida. June 2023.
- [23] Isaac Grosf, Ziv Scully, Mor Harchol-Balter, and Alan Scheller-Wolf. “Optimal Scheduling in the Multiserver-job Model under Heavy Traffic.” Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 6, Issue 3, Article 51 (Dec 2022), pp. 51:1 – 51:32.
- [23] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2023 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 23)*, Orlando, Florida. June 2023.
- [22] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf. “WCFS: a new framework for analyzing multi-server systems.” *Queueing Systems*, vol. 102, June, 2022, pp. 143–174.
- [22] Mor Harchol-Balter. “The multiserver job queueing model.” *Queueing Systems*, vol. 100, 2022, pp. 201–203.
- [21] Mor Harchol-Balter. “Open Problems in Queueing Theory Inspired by Datacenter Computing.” *Queueing Systems*, vol. 97, no. 1, 2021, pp. 3–37.
- [21] Ziv Scully, Isaac Grosf, and Mor Harchol-Balter. “The Gittins Policy is Nearly Optimal in the M/G/k under Extremely General Conditions.” Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 4, Issue 3, Article 43 (Dec 2020), pp. 43:1 – 43:29.
- [21] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2021 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 21)*, Beijing, China. June 2021. Winner **INFORMS 2022 George Nicholson Prize**.
- [21] Mor Harchol-Balter, Takayuki Osogami, Alan Scheller-Wolf, Adam Wierman. “Correction to: Multi-server Queueing Systems with Multiple Priority Classes.” *Queueing Systems*.
- [21] Isaac Grosf, Kunhe Yang, Ziv Scully, and Mor Harchol-Balter. “Nudge: Stochastically improving upon FCFS.” Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 5, Issue 2, Article 21 (2021), pp. 21:1 – 21:29.
- [21] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2021 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 21)*, Beijing, China. June 2021. Winner **Best Paper Award**.
- [21] Rein Vesilo, Mor Harchol-Balter, Alan Scheller-Wolf. “Scaling properties of queues with time-varying load processes: extensions and applications..” *Probability in the Engineering and Informational Sciences*. 2021, pp. 1–42.
- [21] Weina Wang, Qiaomin Xie, and Mor Harchol-Balter. “Zero Queueing for Multi-Server Jobs.” Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 5, Issue 1, Article 7 (2021), pp. 7:1 – 7:25.
- [21] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2021 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 21)*, Beijing, China. June 2021.

- [21] Sherwin Doroudi, Thanassis Avgerinos, and Mor Harchol-Balter. “To Clean or Not to Clean: Malware Removal Strategies for Servers Under Load.” *European Journal of Operational Research*, vol. 292, 2021), pp. 596–609.
- [21] Isaac Grosf, Ziv Scully, Mor Harchol-Balter. “Optimal Multiserver Scheduling with Unknown Job Sizes in Heavy Traffic.” In *Performance Evaluation*, vol. 145, 2021, pp. 1-31.
- [20] Conference version appeared in: *38th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (IFIP Performance 2020)*, Milan, Italy, November 2020.
- [20] Benjamin Berg, Rein Vesilo, and Mor Harchol-Balter. “heSRPT: Parallel Scheduling to Minimize Mean Slowdown” In *Performance Evaluation*, vol. 144, 2020, pp. 1-20.
- [20] Conference version appeared in: *38th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (IFIP Performance 2020)*, Milan, Italy, November 2020.
- [20] Ziv Scully, Mor Harchol-Balter, Alan Scheller-Wolf. “Simple Near-Optimal Scheduling for the M/G/1.” *Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 4, Issue 1, Article 11 (March 2020)*, pp. 11:1 – 11:29.
- [20] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2020 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 20)*, Boston, MA. June 2020. Winner **Best Video Presentation Award**.
- [19] Isaac Grosf, Ziv Scully, Mor Harchol-Balter. “Load Balancing Guardrails: Keeping Your Heavy Traffic on the Road to Low Response Times.” *Proceedings of the ACM Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 3, Issue 2, Article 42 (June 2019)*, pp. 42:1 – 42:31, 2019.
- [19] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2019 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 19)*, Phoenix, AZ. June 2019. Winner **Best Student Paper Award**. Also **STOC 21 Theory Fest Short Plenary Pick**.
- [19] Weina Wang, Mor Harchol-Balter, Haotian Jiang, Alan Scheller-Wolf, and R. Srikant. “Delay asymptotics and bounds for multitask parallel jobs.” *Queueing Systems*, vol. 91, no. 3-4, March 2019, pp. 207–239.
- [18] Conference version appeared in: *36th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2018)*, Toulouse, France, December 2018.
- [18] Isaac Grosf, Ziv Scully, Mor Harchol-Balter. “SRPT for Multiserver Systems.” *Performance Evaluation*, vol. 127-128, Nov. 2018, pp. 154–175.
- [18] Conference version appeared in: *36th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2018)*, Toulouse, France, December 2018. Winner of **Best Student Paper Award**.
- [18] Daniel Berger, Nathan Beckmann, Mor Harchol-Balter. “Practical Bounds on Optimal Caching with Variable Object Sizes.” *Proceedings of the ACM on Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 2, Issue 2, June 2018*, pp. 32:1 – 32:38, 2018.
- [18] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2018 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 18)*, Los Angeles, CA. June 2018.
- [18] Ziv Scully, Mor Harchol-Balter, Alan Scheller-Wolf. “SOAP: One Clean Analysis of All Age-Based Scheduling Policies.” *Proceedings of the ACM on Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 2, Issue 1, March 2018*, pp. 16:1 – 16:30, 2018. **Applied Probability Society Best Student Paper Award Finalist**.

- [18] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2018 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 18)*, Los Angeles, CA. June 2018.
- [17] Ben Berg, Jan-Pieter Dorsman, Mor Harchol-Balter. “Towards Optimality in Parallel Job Scheduling.” *Proceedings of the ACM on Measurement and Analysis of Computer Systems - SIGMETRICS: Volume 1, Issue 2*, December 2017, pp: 40:1 – 40:30.
- [18] Conference version appeared in *Proceedings of ACM Sigmetrics/Performance 2018 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 18)*, Los Angeles, CA. June 2018.
- [17] Kristen Gardner, Mor Harchol-Balter, Alan Scheller-Wolf. “A Better Model for Job Redundancy: Decoupling Server Slowdown and Job Size.” *IEEE/ACM Transactions on Networking*, vol. 25, no. 6, pp. 3353-3367, 2017.
- [16] Earlier version appeared as: “A Better Model for Job Redundancy: Decoupling Server Slowdown and Job Size.” *IEEE Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2016)*, London, UK, September 2016, pp. 1–10. **First runner-up for Best Paper Award.**
- [17] Kristen Gardner, Mor Harchol-Balter, Esa Hyytia, Rhonda Righter. “Scheduling for efficiency and fairness in systems with redundancy.” *Performance Evaluation*, Vol. 116, November, 2017, pp. 1–25.
- [17] Kristen Gardner, Mor Harchol-Balter, Alan Scheller-Wolf, Mark Velednitsky, Samuel Zbarsky. “Redundancy-d: The Power of d Choices for Redundancy” *Operations Research*, vol. 65, no. 4, pp. 1078-1094, 2017.
- [16] Earlier version appeared as: “The Power of d Choices for Redundancy.” *Proceedings of ACM SIGMETRICS 2016 Conference on Measurement and Modeling of Computer Systems Poster Paper*, Antibes Juan-les-Pins, France, June 2016.
- [16] Sherwin Doroudi, Brian Fralix, Mor Harchol-Balter. “Clearing analysis on phases: Exact limiting probabilities for skip-free, unidirectional, quasi-birth-death processes” *Stochastic Systems*, vol. 6, no. 2, 2016, pp. 420-458.
- [16] Kristen Gardner, Sam Zbarsky, Sherwin Doroudi, Mor Harchol-Balter, Esa Hyytia, Alan Scheller-Wolf. “Reducing Latency via Redundant Requests: Exact Analysis” *Queueing Systems: Theory and Applications*, vol. 83, no. 3, 2016, pp. 227-259.
- [15] Earlier version appeared in *Proceedings of ACM Sigmetrics/Performance 2015 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 15)*, Portland, OR. June 2015, pp. 347–360. **Selected as one of best papers and invited to QUESTA journal.**
- [14] Sherwin Doroudi, Esa Hyytia, Mor Harchol-Balter. “Value Driven Load Balancing.” *Performance Evaluation*, vol. 79, 2014, pp. 306-327.
- [14] Earlier version appeared as: “Value Driven Load Balancing.” *32nd International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2014)*, Turin, Italy, November 2014.
- [14] Anshul Gandhi, Sherwin Doroudi, Mor Harchol-Balter, and Alan Scheller-Wolf. “Exact Analysis of the M/M/k/setup Class of Markov Chains via Recursive Renewal Reward.” *Queueing Systems: Theory and Applications*, vol. 77, no. 2, 2014, pp. 177–209.
- [13] Earlier version appeared in *Proceedings of ACM Sigmetrics 2013 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 13)*, Pittsburgh, PA June, 2013, pp. 153–166. **Selected as one of best papers and invited to Queueing Systems journal.**

- [14] Christian Borgs, Jennifer T. Chayes, Sherwin Doroudi, Mor Harchol-Balter and Kuang Xu. “The Optimal Admission Threshold in Observable Queues with State Dependent Pricing.” *Probability in the Engineering and Informational Sciences*, vol. 28, 2014, pp. 101-119.
- [13] Anshul Gandhi and Mor Harchol-Balter. “M/G/k with Staggered Setup.” *OR Letters*, vol. 41, no. 4, 2013, pp. 317-320.
- [12] Anshul Gandhi, Mor Harchol-Balter, Ram Raghunathan, and Mike Kozuch. “AutoScale: Dynamic, Robust Capacity Management for Multi-Tier Data Centers.” *ACM Transactions on Computer Systems*, vol. 30, no. 4, article 14, 2012, pp. 1 - 26.
- [12] Earlier version appeared as: Carnegie Mellon Technical Report CMU-CS-12-109, 2012.
- [12] ChangWoo Yang, Adam Wierman, Sanjay Shakkotti, and Mor Harchol-Balter. “Many flows asymptotics for SMART scheduling policies,” *IEEE Transactions on Automatic Control*, vol. 57, no. 2, 2012, pp. 376-391.
- [06] Earlier version appeared as: “Tail Asymptotics for Policies Favoring Short Jobs in a Many-Flows Regime.” *Proceedings of ACM Sigmetrics 2006 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 06)*, Saint Malo, France. June 2006, pp. 97-108.
- [11] Field Cady, Yi Zhuang, and Mor Harchol-Balter. “A Stochastic Analysis of Hard Disk Drives,” *International Journal of Stochastic Analysis*, vol. 2011, Article ID 390548, 21 pages.
- [11] Yoongu Kim, Michael Papamichael, Onur Mutlu, and Mor Harchol-Balter. “Thread Cluster Memory Scheduling”, *IEEE Micro, Special Issue: Micro’s Top Picks from 2010 Computer Architecture Conferences (MICRO TOP PICKS)*, vol. 31, no. 1, January/February, 2011, pp. 78-89.
- [10] Earlier version appeared as: “Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior.” *43rd International Symposium on Microarchitecture (MICRO-43)*, Atlanta, Georgia, December 2010, pp. 65-76. **Selected as Top Picks 2010 paper by IEEE Micro.**
- [11] Mor Harchol-Balter. “Queueing Disciplines.” *Wiley Encyclopedia Of Operations Research and Management Science*, 2011.
- [10] Varun Gupta, Michelle Burroughs, and Mor Harchol-Balter. “Analysis of Scheduling Policies under Correlated Job Sizes.” *Performance Evaluation*, vol. 67, no. 11, 2010, pp. 996-1013.
- [10] Earlier version appeared as: “Analysis of Scheduling Policies under Correlated Job Sizes.” *28th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2010)*, Namur, Belgium, November 2010.
- [10] Anshul Gandhi, Mor Harchol-Balter, and Ivo Adan. “Server Farms with Setup Costs.” *Performance Evaluation*, vol. 67, no. 11, 2010, pp. 1123-1138.
- [10] Earlier version appeared as: “Server Farms with Setup Costs.” *28th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2010)*, Namur, Belgium, November 2010.
- [10] Anshul Gandhi, Varun Gupta, Mor Harchol-Balter, and Michael Kozuch. “Optimality Analysis of Energy-Performance Trade-offs for Server Farm Management.” *Performance Evaluation*, vol. 67, no. 11, 2010, pp. 1155-1171.
- [10] Earlier version appeared as: “Optimality Analysis of Energy-Performance Trade-offs for Server Farm Management.” *28th International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2010)*, Namur, Belgium, November 2010.

- [10] Mor Harchol-Balter and Rein Vesilo. "To Balance or Unbalance Load in Size-Interval Task Allocation." *Probability in the Engineering and Informational Sciences*, vol. 24, 2010, pp. 219-244.
- [10] Varun Gupta, Mor Harchol-Balter, Jim Dai, and Bert Zwart. "On the Inapproximability of M/G/k: Why Two Moments of Job Size Distribution are Not Enough." *Queueing Systems: Theory and Applications*, vol. 64, no. 1, January, 2010, pp. 5- 48.
- [08] Atul Bhandari, Alan Scheller-Wolf, and Mor Harchol-Balter. "An Exact and Efficient Algorithm for the Constrained Dynamic Operator Staffing Problem for Call Centers." *Management Science*, vol. 54, no. 2, February 2008, pp. 339-353.
- [07] Varun Gupta, Mor Harchol-Balter, Karl Sigman, and Ward Whitt. "Analysis of Join-the-Shortest-Queue Routing for Web Server Farms." *Performance Evaluation*, vol. 64, no. 9-12, October 2007, pp. 1062-1081.
- [07] Earlier version appeared in *IFIP 7.3 International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2007)*, Cologne, Germany, October 2007.
- [06] Robert C. Hampshire, Mor Harchol-Balter, and William A. Massey. "Fluid and Diffusion Limits for Transient Sojourn Times of Processor Sharing Queues with Time Varying Rates." *Queueing Systems: Theory and Applications journal (QUESTA)*, vol. 53, no. 1/2, June 2006, pp. 19 - 30.
- [06] Takayuki Osogami and Mor Harchol-Balter. "Closed Form Solutions for Mapping General Distributions to Quasi-Minimal PH Distributions." *Performance Evaluation*, vol. 63(6), pp. 524-552 (2006). (Special Issue for best papers of Tools 03).
- [03] Earlier version appeared as: "A Closed-Form Solution for Mapping General Distributions to Minimal PH Distributions." *Lecture Notes in Computer Science 2794: The 13th International Conference on Modeling Techniques and Tools for Computer Performance Evaluation (Performance TOOLS 03)*. Urbana, Illinois, September 2003, pp. 200 - 217. **Selected as one of best papers.**
- [06] Bianca Schroeder and Mor Harchol-Balter. "Web Servers under Overload: How Scheduling can Help." *ACM Transactions on Internet Technologies*, vol. 6, no. 1, February, 2006, pp. 20-52.
- [03] Earlier version appeared in *Teletraffic Science and Engineering: Providing QoS in Heterogeneous Environments*, Elsevier publisher, vol. 5a, 2003. Based on *Proceedings of the 18th IEEE International Teletraffic Congress*, Berlin, Germany, September, 2003, pp. 171-180. Winner of **IEEE International Teletraffic Congress Best Student Paper Award.**
- [06] Adam Wierman, Takayuki Osogami, Mor Harchol-Balter, Alan Scheller-Wolf. "How Many Servers are Best in a Dual-Priority M/PH/k System?" *Performance Evaluation*, vol. 63, no. 12, 2006, pp. 1253-1272.
- [05] Mor Harchol-Balter, Takayuki Osogami, Alan Scheller-Wolf, and Adam Wierman. "Multi-Server Queueing Systems with Multiple Priority Classes." *Queueing Systems: Theory and Applications journal (QUESTA)*, vol. 51, no.3-4, 2005, pp. 331-360.
- [05] Takayuki Osogami, Mor Harchol-Balter, and Alan Scheller-Wolf. "Analysis of Cycle Stealing with Switching Cost and Thresholds." *Performance Evaluation*, vol. 61, no. 4, 2005, pp. 347-369.
- [03] Earlier version appeared as: "Analysis of Cycle Stealing with Switching Cost." *Proceedings of ACM Sigmetrics 2003 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 03)*. San Diego, CA. June 2003, pp. 184-195.
- [04] Bianca Schroeder and Mor Harchol-Balter. "Evaluation of Task Assignment Policies for Supercomputing Servers: The Case for Load Unbalancing and Fairness." Invited to *Cluster Computing*, vol. 7, no. 2, April 2004, pp. 151-161.

- [00] Earlier version appeared in: *Proceedings of 9th IEEE Symposium on High Performance Distributed Computing (HPDC '00)*, Pittsburgh, PA, August 2000, pp. 211–220.
- [04] Adam Wierman, Nikhil Bansal, Mor Harchol-Balter. “A Note on Comparing Response Times in M/GI/1/FB and M/GI/1/PS Queues.” *Operations Research Letters*, vol. 32, no. 1, January 2004, pp. 73-76.
- [03] Mor Harchol-Balter, Bianca Schroeder, Nikhil Bansal, and Mukesh Agrawal. “Size-based Scheduling to Improve Web Performance.” *ACM Transactions on Computer Systems*, vol. 21, no. 2, May 2003, pp. 207-233.
- [02] Mor Harchol-Balter. “Task Assignment with Unknown Duration.” *Journal of the ACM*, vol. 49, no. 2, March 2002, pp. 260-288.
- [00] Earlier version appeared in: *20th IEEE International Conference on Distributed Computing Systems (ICDCS '00)* Taipei, Taiwan, April 2000, pp. 214-223.
- [02] Mor Harchol-Balter, Karl Sigman, and Adam Wierman. “Asymptotic Convergence of Scheduling Policies with Respect to Slowdown.” *Performance Evaluation*, vol. 49, no. 1-4, 2002, pp. 241-256.
- [02] Earlier version appeared in: *IFIP 7.3 International Symposium on Computer Performance, Modeling, Measurements, and Evaluation (Performance 2002)*, Rome, Italy, September 2002.
- [00] Matthew Andrews, Antonio Fernandez, Mor Harchol-Balter, Tom Leighton, Lisa Zhang. “General Dynamic Routing with Per-Packet Delay Guarantees of $O(\text{distance} + 1/\text{session rate})$.” *SIAM Journal on Computing*, vol. 30, no. 5, 2000, pp. 1594-1623.
- [97] Earlier version appeared in: *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science (FOCS 97)*, Miami, FL, October 1997, pp. 294-303.
- [99] Mor Harchol-Balter, Mark Crovella, Cristina Murta. “On Choosing a Task Assignment Policy for a Distributed Server System.” *IEEE Journal of Parallel and Distributed Computing*, vol. 59, 1999, pp. 204-228.
- [98] Earlier version appeared in: *Lecture Notes in Computer Science, No. 1469: 10th International Conference on Modeling Techniques and Tools for Computer Performance Evaluation (Performance TOOLS 98)*, September 1998, pp. 231-242.
- [94] Peter Gemmel and Mor Harchol. “Tight Bounds on Expected Time to Add Correctly and Add Mostly Correctly.” *Information Processing Letters*, vol.49, January 1994, pp. 77-83.

Refereed Conference/Workshop Papers (not overlapping with above)

- [23] Isaac Grosf and Mor Harchol-Balter. “ServerFilling: A better approach to packing multiserver jobs.” *Proceedings of the 5th workshop on Advanced tools, programming languages, and platforms for implementing and evaluating algorithms for distributed systems (ApPLIED 2023)*. In conjunction with PODC 2023 Conference. Article No. 7, pp. 1–5, Orlando, CA, June 2023.
- [23] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf. “New Stability Results for Multiserver-job Models via Product-form Saturated Systems.” *ACM SIGMETRICS Performance Evaluation Review*, vol. 4, no. 6, pp. 1–3. (MAMA 2023).
- [21] Mor Harchol-Balter and Ziv Scully. “The Most Common Queueing Theory Questions Asked by Computer Systems Practitioners.” *1st International Workshop on Teaching Performance Analysis of Computer Systems (TeaPACS 2021)*. In conjunction with the IFIP Performance 2021 Conference. Milan, Italy, November 2021.

- [21] Ziv Scully and Mor Harchol-Balter. “The Gittins Policy in the M/G/1 Queue.” *19th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt 2021)*, Philadelphia, PA, October 2021.
- [21] Naveen Pai, Sean Zhang, Mor Harchol-Balter. “Computing the Death Rate of COVID-19.” In: *Byrski A., Czachurski T., Gelenbe E., Grochla K., Murayama Y. (eds) International Conference on Computer Science Protecting Human Society Against Epidemics (ANTICOVID 2021). IFIP Advances in Information and Communication Technology*, vol 616. Springer, Cham., pp 77-94, 2021.
- [21] Ben Berg, Justin Whitehouse, Ben Moseley, Weina Wang, Mor Harchol-Balter. “The Case for Phase-Aware Scheduling of Parallelizable Jobs.” *IFIP Performance Conference*, Milan, Italy, November 2021.
- [20] Ben Berg, Daniel Berger, Sara McAllister, Isaac Grosf, Sathya Gunasekar, Jimmy Lu, Michael Uhlar, Jim Carrig, Nathan Beckmann, Mor Harchol-Balter, Greg Ganger. “The CacheLib Caching Engine: Design and Experiences at Scale.” *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2020)*, Banff, CA, November 2020, pp. 769–786.
- [20] Benjamin Berg, Mor Harchol-Balter, Ben Moseley, Weina Wang, Justin Whitehouse. “Optimal Resource Allocation for Elastic and Inelastic Jobs.” *32nd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '20)*, Philadelphia, PA, July 2020, pp. 75–87.
- [20] Muhammad Tirmazi, Adam Barker, Nan Deng, Md Ehtesam Haque, Zhijing Gene Qin, Steven Hand, Mor Harchol-Balter, John Wilkes. “Borg: the Next Generation.” *Proceedings of the Fifteenth EuroSys Conference, (EUROSYS '20)*, Greece, April 2020, pp. 1–14.
- [19] Benjamin Berg, Rein Vesilo, Mor Harchol-Balter. “heSRPT: Optimal Scheduling of Parallel Jobs with Known Sizes.” *ACM SIGMETRICS Performance Evaluation Review*, vol. 47, no. 2, pp. 18-20. (MAMA 2019).
- [18] Ziv Scully, Mor Harchol-Balter. “SOAP Bubbles: Robust Scheduling under Adversarial Noise.” *56th Annual Allerton Conference on Communication, Control, and Computing (Allerton 2018)*, Urbana-Champaign, IL, October 2018, pp. 144-154.
- [18] Daniel S. Berger, Benjamin Berg, Timothy Zhu, Siddhartha Sen, Mor Harchol-Balter. “RobinHood: Tail Latency Aware Caching – Dynamic Reallocation from Cache-Rich to Cache-Poor.” *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2018)*, Carlsbad, CA, October 2018, pp. 195–212.
- [18] Isaac Grosf, Ziv Scully, Mor Harchol-Balter. “SRPT for Multiserver Systems” *ACM SIGMETRICS Performance Evaluation Review*, vol. 46, no. 3, pp. 8–9 (MAMA 2018).
- [17] Timothy Zhu, Mor Harchol-Balter, Michael Kozuch. “WorkloadCompactor: Reducing datacenter cost while providing tail latency SLO guarantees.” *ACM Symposium on Cloud Computing (SoCC'17)*, Santa Clara, Oct 2017, pp. 598–610.
- [17] Ziv Scully, Guy E. Blelloch, Mor Harchol-Balter, Alan Scheller-Wolf. “Optimally Scheduling Jobs with Multiple Tasks.” *ACM SIGMETRICS Performance Evaluation Review*, vol. 45, no. 2, pp. 36-38. (MAMA 2017).
- [17] Daniel Berger, Ben Berg, Timothy Zhu, Mor Harchol-Balter. “The Case of Dynamic Cache Partitioning for Tail Latency.” *14th USENIX Symposium on Networked Systems Design and Implementation Poster Paper (NSDI '17 Poster)*, Boston, MA, March 2017.
- [17] Daniel Berger, Ramesh Sitaraman, Mor Harchol-Balter. “AdaptSize: Orchestrating the Hot Object Memory Cache in a CDN.” *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI '17)*, Boston, MA, March 2017, pp. 483–498.

- [16] Timothy Zhu, Daniel Berger, Mor Harchol-Balter. “SNC-Meister: Admitting More Tenants with Tail Latency SLOs.” *ACM Symposium on Cloud Computing (SoCC’16)*, Santa Clara, October 2016, pp. 374–387.
- [16] Kristen Gardner, Samuel Zbarsky, Mark Velednitsky, Mor Harchol-Balter, Alan Scheller-Wolf. “Understanding Response Time in the Redundancy-d System.” *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, no. 2, pp. 33-35. (MAMA 2016).
- [16] Alexey Tumanov, Timothy Zhu, Jun Woo Park, Michael A. Kozuch, Mor Harchol-Balter, Gregory R. Ganger. “TetriSched: Optimistic Global Continuous Rescheduling in Dynamic Heterogeneous Clusters.” *Proceedings of EuroSys 2016*, London, UK, April 2016. Winner of **EuroSys Best Student Paper Award**.
- [15] Kristen Gardner, Sem Borst, Mor Harchol-Balter. “Optimal Scheduling for Jobs with Progressive Deadlines.” *IEEE INFOCOM 15*, Hong Kong, April, 2015.
- [14] Timothy Zhu, Alexey Tumanov, Michael A. Kozuch, Mor Harchol-Balter, Gregory R. Ganger. “PriorityMeister: Tail Latency QoS for Shared Networked Storage.” *ACM Symposium on Cloud Computing 2014 (SoCC’14)*, Seattle, WA, November, 2014, pp. 1–14.
- [12] Christian Borgs, Jennifer T. Chayes, Sherwin Doroudi, Mor Harchol-Balter, and Kuang Xu. “Pricing and queueing.” *ACM Performance Evaluation Review (MAMA 12)*, vol. 40, no.3, 2012, pp. 71-73.
- [12] Anshul Gandhi, Timothy Zhu, Mor Harchol-Balter, and Mike Kozuch. “SOFTScale: Stealing Opportunistically for Transient Scaling.” *ACM/IFIP/USENIX Middleware Conference (Middleware 12)*, Montreal, Canada, December, 2012, pp. 142-163.
- [12] Earlier version appeared as: Carnegie Mellon Technical Report CMU-CS-12-111, 2012.
- [12] Timothy Zhu, Anshul Gandhi, Mor Harchol-Balter, and Michael Kozuch. “Saving Cash by Using Less Cache.” *4th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 12)*, Boston, MA, June, 2012.
- [12] Anshul Gandhi, Mor Harchol-Balter, and Michael Kozuch. “Are sleep states effective in data centers?” *IEEE International Green Computing Conference (IGCC 2012)*, San Jose, CA, June 2012, pp. 1-10.
- [11] Anshul Gandhi, Mor Harchol-Balter, Ram Raghunathan and Michael Kozuch. “Distributed, Robust Auto-Scaling Policies for Power Management in Compute Intensive Server Farms.” *Open Cirrus Summit 11*, Georgia Tech, Atlanta, GA, October 2011.
- [11] Anshul Gandhi and Mor Harchol-Balter. “How Data Center Size Impacts the Effectiveness of Dynamic Power Management.” *49th Annual Allerton Conference on Communication, Control, and Computing (Allerton 11)*, Urbana-Champaign, IL, 2011.
- [11] Anshul Gandhi, Mor Harchol-Balter, and Mike Kozuch. “The case for sleep states in servers.” *SOSP 4th Workshop on Power-Aware Computing and Systems (HotPower 2011)*, Cascais, Portugal, October 2011, pp. 1–5.
- [10] Yoongu Kim, Dongsu Han, Onur Mutlu, and Mor Harchol-Balter. “ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers.” *16th IEEE International Symposium on High-Performance Computer Architecture (HPCA 10)*, 2010, pp. 1 - 12.
- [10] Anshul Gandhi, Mor Harchol-Balter, and Ivo Adan. “Analysis of an M/M/k System with Exponential Setup Times under Staggered Bootstrap.” *Madrid Conference on Queueing Theory (MCQT 10)*, Toledo, Spain, June 2010.
- [10] Anshul Gandhi, Mor Harchol-Balter, and Ivo Adan. “Decomposition Results for an M/M/k with Staggered Setup.” *ACM Performance Evaluation Review (MAMA 10)*, vol. 38, no.2, 2010, pp. 48-50.

- [09] Mor Harchol-Balter, Alan Scheller-Wolf, and Andrew Young. “Why Segregating Short Jobs from Long Ones under High Variability is Not Always a Win.” *47th Annual Allerton Conference on Communication, Control, and Computing (Allerton 09)*, Urbana-Champaign, Illinois, September, 2009, pp. 121–127.
- [09] Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, Jeff Kephart, and Charles Lefurgy. “Power Capping Via Forced Idleness.” *Workshop on Energy-Efficient Design (WEED 09)*, Austin, Texas, June 2009.
- [09] Mor Harchol-Balter, Alan Scheller-Wolf, and Andrew Young. “Surprising Results on Task Assignment in Server Farms with High-Variability Workloads.” *Proceedings of ACM Sigmetrics/Performance 2009 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 09)*, Seattle, WA. June 2009, pp. 287-298.
- [09] Anshul Gandhi, Mor Harchol-Balter, Rajarshi Das, and Charles Lefurgy. “Optimal Power Allocation in Server Farms.” *Proceedings of ACM Sigmetrics/Performance 2009 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 09)*, Seattle, WA. June 2009, pp. 157-168.
- [09] Varun Gupta and Mor Harchol-Balter. “Self-Adaptive Admission Control Policies for Resource-Sharing Systems.” *Proceedings of ACM Sigmetrics/Performance 2009 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 09)*, Seattle, WA. June 2009, pp. 311-322.
- [07] Varun Gupta, Jim Dai, Mor Harchol-Balter, and Bert Zwart. “The Effect of Higher Moments of Job Size Distribution on the Performance of an M/G/s Queueing System.” *ACM Performance Evaluation Review (MAMA 07)*, vol. 35, no.2, 2007, pp. 12-14.
- [07] Varun Gupta, Karl Sigman, Mor Harchol-Balter, and Ward Whitt. “Insensitivity for PS server farms with JSQ routing.” *ACM Performance Evaluation Review (MAMA 07)*, vol. 35, no.2, 2007, pp. 24-26.
- [06] Varun Gupta, Mor Harchol-Balter, Alan Scheller-Wolf, and Uri Yechiali. “Fundamental Characteristics of Queues with Fluctuating Load.” *Proceedings of ACM Sigmetrics 2006 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 06)*, Saint Malo, France. June 2006, pp. 203-215.
- [06] Bianca Schroeder, Adam Wierman, and Mor Harchol-Balter. “Closed versus Open System Models: a Cautionary Tale.” *Proceedings of Networked Systems Design and Implementation (NSDI 06)*, San Jose, CA. May 2006. pp. 239-252.
- [06] Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, Erich Nahum, and Adam Wierman. “How to Determine a Good Multi-Programming Level for External Scheduling.” *22nd IEEE International Conference on Data Engineering (ICDE 06)*, Atlanta, Georgia, April 2006, pp. 60-70.
- [06] Bianca Schroeder, Mor Harchol-Balter, Arun Iyengar, Erich Nahum. “Achieving Class-Based QoS for Transactional Workloads.” Poster paper in *22nd IEEE International Conference on Data Engineering (ICDE 06)*, Atlanta, Georgia, April 2006, pp. 153 - 154.
- [05] Adam Wierman, Mor Harchol-Balter, and Takayuki Osogami. “Nearly Insensitive Bounds on SMART Scheduling.” *Proceedings of ACM Sigmetrics 2005 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 05)*, Banff, Canada, June 2005, pp. 205-216.
- [05] Adam Wierman and Mor Harchol-Balter. “Classifying Scheduling Policies with Respect to Higher Moments of Conditional Response Time.” *Proceedings of ACM Sigmetrics 2005 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 05)*, Banff, Canada, June 2005, pp. 229-240.
- [05] Mor Harchol-Balter, Takayuki Osogami, and Alan Scheller-Wolf. “Robustness of Threshold Policies in Beneficiary-Donor Model.” *ACM Performance Evaluation Review (MAMA 05)*, vol. 33, no.2, 2005, pp. 36-38.
- [05] David McWherter, Bianca Schroeder, Natassa Ailamaki, and Mor Harchol-Balter. “Improving Preemptive Prioritization via Statistical Characterization of OLTP Locking.” *Proceedings of the 21st IEEE International Conference on Data Engineering (ICDE 05)*, San Francisco, CA, April 2005, pp. 446-457.

- [04] Takayuki Osogami, Mor Harchol-Balter, Alan Scheller-Wolf, and Li Zhang. "Exploring Threshold-based Policies for Load Sharing." *42nd Annual Allerton Conference on Communication, Control, and Computing (Allerton 04)*, University of Illinois, Urbana-Champaign, October, 2004, pp. 1012-1021.
- [04] Takayuki Osogami, Adam Wierman, Mor Harchol-Balter, Alan Scheller-Wolf. "A Recursive Analysis Technique for Multi-Dimensionally infinite Markov Chains." *ACM Performance Evaluation Review (MAMA 04)*, vol. 32, no. 2, 2004, pp. 3-5.
- [04] Adam Wierman and Mor Harchol-Balter. "Formalizing SMART Scheduling." *ACM Performance Evaluation Review (MAMA 04)*, vol. 32, no. 2, 2004, pp. 12-13.
- [04] David McWherter, Bianca Schroeder, Natassa Ailamaki, and Mor Harchol-Balter. "Priority Mechanisms for OLTP and Transactional Web Applications." *20th IEEE International Conference on Data Engineering (ICDE 04)*, Boston, MA, April 2004, pp. 535-546.
- [03] Takayuki Osogami and Mor Harchol-Balter. "Necessary and Sufficient Conditions for Representing General Distributions by Coxians." *Lecture Notes in Computer Science 2794. 13th International Conference on Modeling Techniques and Tools for Computer Performance Evaluation (Performance TOOLS 03)*, Urbana, Illinois, September 2003, pp. 182-199. **Selected as one of best papers.**
- [03] Mor Harchol-Balter, Cuihong Li, Takayuki Osogami, Alan Scheller-Wolf, and Mark Squillante. "Cycle Stealing under Immediate Dispatch Task Assignment." *Fifteenth ACM Annual Symposium on Parallel Algorithms and Architectures (SPAA 03)*, San Diego, CA, June, 2003, pp. 274-285.
- [03] Adam Wierman and Mor Harchol-Balter. "Classifying Scheduling Policies with respect to Unfairness in an M/G/1." *Proceedings of ACM Sigmetrics 2003 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 03)*, San Diego, CA, June 2003, pp. 238-249. Winner of **Sigmetrics Best Student Paper Award.**
- [03] Mor Harchol-Balter, Cuihong Li, Takayuki Osogami, Alan Scheller-Wolf, Mark Squillante. "Task Assignment with Cycle Stealing under Central Queue." *23rd IEEE International Conference on Distributed Computing Systems (ICDCS 03)*, Providence, RI, May 2003, pp. 628 - 637.
- [01] Nikhil Bansal and Mor Harchol-Balter. "Analysis of SRPT Scheduling: Investigating Unfairness." *Proceedings of ACM Sigmetrics 2001 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 01)*, June 2001, Cambridge, MA, pp. 279-290.
- [01] Mor Harchol-Balter. "Job Placement with Unknown Duration and No Preemption" *ACM Performance Evaluation Review (MAMA 01)*, vol. 28, no. 4, March 2001.
- [01] Nikhil Bansal and Mor Harchol-Balter. "Approximate Analysis of M/G/1/PS and M/G/1/SRPT under Transient Overload." *ACM Performance Evaluation Review (MAMA 01)*, vol 29, no. 3, December, 2001. Extended version available as Technical report Number CMU-CS-01-134.
- [01] Mor Harchol-Balter, Nikhil Bansal, Bianca Schroeder. "SRPT Scheduling for Web Servers." (invited paper) *Lecture Notes in Computer Science, No. 2221. 7th International Workshop, Job Scheduling Strategies for Parallel Processing (JSSPP 01)*, Cambridge, MA, June 2001, pp. 11-20.
- [99] Mark Crovella, Bob Frangioso, and Mor Harchol-Balter. "Connection Scheduling in Web Servers." *USENIX Symposium on Internet Technologies and Systems (USITS 99)*, Boulder, Colorado, October 1999, pp. 243-254.
- [99] Mor Harchol-Balter, Tom Leighton, and Daniel Lewin. "Resource Discovery in Distributed Networks." *18th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC 99)*, Atlanta, May 1999, pp. 229-238.

- [99] Mor Harchol-Balter. “The Effect of Heavy-Tailed Job Size Distributions on Computer System Design.” *Proceedings of the ASA-IMS Conference on Applications of Heavy Tailed Distributions in Economics, Engineering and Statistics*, Washington, DC, June 1999.
- [98] Mark Crovella, Mor Harchol-Balter, Cristina Murta. “Task Assignment in a Distributed System: Improving Performance by Unbalancing Load.” *Proceedings of ACM Sigmetrics 1998 Conference on Measurement and Modeling of Computer Systems (SIGMETRICS 98) Poster Session*, June 1998, Madison, WI, pp. 268-269.
- [95] Mor Harchol-Balter and David Wolfe. “Bounding Delays in Packet-Routing Networks.” *Proceedings of the 27th Annual ACM Symposium on the Theory of Computation (STOC 95)*, Las Vegas, May 1995, pp. 248-275.
- [94] Micah Adler, Peter Gemmell, Mor Harchol-Balter, Richard M. Karp, Claire Kenyon. “Selection in the Presence of Noise: The Design of Playoff Systems.” *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 94)*, Arlington, VA, January 1994, pp. 564-572.
- [94] Mor Harchol-Balter and Paul E. Black. “Queueing analysis of oblivious packet-routing algorithms.” *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 94)*, Arlington, VA, January 1994, pp. 583-592.

Technical reports (not overlapping with above)

- [21] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf. “The Finite-Skip Method for Multiserver Analysis.” ArXiv:2109.12663. September, 2021.
- [20] Isaac Grosf, Mor Harchol-Balter, Alan Scheller-Wolf. “Stability for Two-class Multiserver-job Systems.” ArXiv:2010.00631. 2020.
- [13] Sherwin Doroudi, Mustafa Akan, Mor Harchol-Balter, Jeremy Karp, Christian Borgs and Jennifer T. Chayes. “Priority Pricing in Queues with a Continuous Distribution of Customer Valuations.” CMU-CS-13-109, May 2013.
- [09] Anshul Gandhi, Mor Harchol-Balter, and Ivo Adan “M/G/k with Exponential Setup.” CMU-CS-09-166, September 2009.
- [98] Mor Harchol-Balter, Mark Crovella, SungSim Park. “The Case for SRPT Scheduling in Web Servers.” MIT-LCS-TR-767, October 1998.
- [97] Mor Harchol-Balter, Mark Crovella, Cristina Murta. “To Queue or Not to Queue: When Queueing is Better Than Timesharing in a Distributed System.” BUCS-TR-1997-017, October 1997.
- [95] Mor Harchol-Balter. “Bounding Delays in Packet-Routing Networks with Light Traffic.” University of California at Berkeley Computer Science Division Tech Report Number UCB/CSD-95-885. October 1995.
- [95] Allen Downey and Mor Harchol-Balter. “A Note on ‘The Limited Performance Benefits of Migrating Active Processes for Load Balancing’.” University of California at Berkeley Computer Science Division Tech Report Number UCB/CSD-95-888, November 1995.
- [94] Mor Harchol-Balter and David Wolfe. “In Network of Queues, M/M/1 Can Outperform M/D/1.” University of California at Berkeley Computer Science Division Tech Report Number UCB/CSD-94-841, November 1994.

Patents Issued

- [15] Mor Harchol-Balter, Anshul Gandhi, Rajarshi Das, and Jeffrey Kephart. “Power Budget Allocation in Multi-Processor Systems.” U.S. Patent Serial No. 9,052,895. Filed: April 7, 2011. Issued June 9, 2015.

- [15] Anshul Gandhi, Varun Gupta, Mor Harchol-Balter, and Mike Kozuch. “Dynamic Capacity Management of Multiple Parallel Connected Computing Resources.” U.S. Patent Serial Number: 8,806,018. Filed March 30, 2012. Issued Aug. 12, 2014.
- [13] Mor Harchol-Balter and Anshul Gandhi. “Systems and Methods for Managing Power Consumption and Performance of a Processor.” U.S. Patent Serial No. 8,589,709. Filed: July 22, 2010. Issued Nov. 19, 2013.
- [10] Mor Harchol-Balter and Tom Leighton and Daniel Lewin. “System and Method for Resource Discovery.” U.S. Patent Serial No. 7,808,926. Filed Oct. 2007. Issued October 5, 2010. *Licensed to Akamai Technologies, Inc.*
- [07] Mor Harchol-Balter and Tom Leighton and Daniel Lewin. “System and Method for Resource Discovery.” U.S. Patent Serial No. 7,295,523. Filed May 1999. Issued November 13, 2007. *Licensed to Akamai Technologies, Inc.*
- [01] Mor Harchol-Balter and Mark Crovella. “Method and Apparatus for Assigning Tasks in a Distributed Server System.” U.S. Patent Serial No. 6,223,205. Filed Feb 13, 1998. Issued April 24, 2001.

Patents Pending

- [05] Arun Iyengar, Erich Nahum, Bianca Schroeder, and Mor Harchol-Balter. “Method and Apparatus for Dynamically Scheduling Requests.” U.S. Patent Application Serial No. 10/804,516. Filed: March 19, 2004.

TALKS

- Nov 2023 “Queueing Models for Today’s Compute Centers.” Keynote talk at IFIP Performance Workshop: Queueing Theory for Emerging Classical and Quantum Systems.
Invited by: Arpan Mukhopadhyay and Thirupathaiiah Vasantam
- Oct 2023 “Queueing Models for Today’s Compute Centers.” Tutorial at APS Informs.
Invited by: Christina Lee Yu and Andrew Daw
- Sept 2023 “Scheduling Multiserver Compute Jobs.” ISyE seminar at Georgia Tech.
Invited by: Debankur Mukherjee, Eunhye Song, Diego F. Cifuentes Pardo
- May 2023 “Queueing Models for Today’s Compute Centers.” Invited talk at SNAPP Seminar, Online.
Invited by: Yuan Hong, Rouba Ibrahim, Chang-Han Rhee, Vsevolod Shneer, and Kuang Xu.
- Oct 2022 “Can Hybrid Priority Increase Revenue”
Invited talk at INFORMS 2022, Indianapolis, IN. Invited by Nasser Barjesteh.
- Sept 2022 “Recent Breakthroughs in Stochastic Scheduling Theory”
Keynote presentation at International Teletraffic Congress 2022. Invited by Minghua Chen and Longbo Huang.
- Sept 2022 “Optimal Scheduling for Multi-Server Systems”
Invited Speaker at Allerton Conference.
Urbana-Champaign, IL. Hosts: R. Srikant and Bruce Hajek.
- June 2022 “Markov Chains Inspired by Data Center Computing”
Keynote presentation at Matrix-Analytic Methods Conference 2022. Invited by Barbara Margolius.
- Dec 2021 “The Most Common Queueing Theory Questions Asked by Computer Systems Practitioners”
James F. Kelly Distinguished Lecture Series
Penn State University, Computer Science Department. Invited by Timothy Zhu.
- Dec 2021 “Recent Breakthroughs in Stochastic Scheduling Theory”
IEEE UEMCON Keynote Speaker.
12th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference.

- Invited by Rajashree Paul.
- Nov 2021 “The Most Common Queueing Theory Questions Asked by Computer Systems Practitioners”
Presentation at 1st International Workshop on Teaching Performance Analysis
of Computer Systems (TeaPACS 2021). Invited by Y.C. Tay.
- Oct 2021 “Overview of Multi-Server Queueing Problems Inspired by Today’s Datacenters”
INFORMS Annual Conference, Anaheim, CA.
- Sept 2021 “New Directions in Stochastic Scheduling”
Keynote Presentation: Data-driven Queueing Challenges (DDQC 2021). Invited by Peter Taylor.
- Sept 2021 “Optimal Scheduling for Multi-Server Systems”
CMU CSD Theory Lunch Seminar.
- June 2021 “Optimal Scheduling in Multi-Server Systems”
STOC Theory Fest Plenary Talk. Invited by Nicole Immorlica.
- Apr 2021 “Recent Breakthroughs in Stochastic Scheduling Theory”
Scheduling Seminar. Invited by Michael Pinedo.
- Nov 2020 “New Breakthroughs in Scheduling Theory”
Columbia University. Applied Probability and Risk Seminar. NYC, NY.
Invited by: Vineet Goyal and Henry Lam.
- Oct 2020 “New Breakthroughs in Scheduling Theory”
MIT-Laboratory of Information and Decision Systems, Boston, MA.
Invited by: Eytan Modiano.
- Nov 2019 “New Breakthroughs in Scheduling”
NYU-Stern School of Business, NY, NY.
Invited by: Zhengyuan Zhou and Michael Pinedo.
- Nov 2019 “Queueing Theory Gets a MakeOver”
MIT-LIDS 80th Anniversary Celebration, Boston, MA.
Invited by: John Tsitsiklis.
- Oct 2019 “Scheduling in a Parallel World”
INFORMS Annual Conference, Seattle, WA.
Invited by: Ishai Menache.
- Aug 2019 “New Breakthroughs in Scheduling”
Queueing Theory Network Applications (QTNA 2019) Keynote Talk, Ghent, Belgium.
Invited by: Sabine Wittevrongel.
- July 2019 “Improving Batch BE”
Google Tech Talk, Sunnyvale, CA.
Invited by: Steve Hand.
- June 2019 “The Power of SOAP Scheduling”
Tutorial speaker at ACM SIGMETRICS 2019 conference, Phoenix, AZ.
- May 2019 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
ORIE Colloquium at Cornell University.
Invited by: Jamol Pender.
- Apr 2019 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
Joint OR/OM Seminar at Tepper School of Business, CMU.
Host: Alan Scheller-Wolf.
- Dec 2018 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
Keynote Speaker at Young European Queueing Theorists conference 2018, Toulouse France.
Invited by: Onno Boxma, Manu Kumar Gupta, and Eyal Castiel.
- Dec 2018 “Delay Asymptotics and Bounds for Multi-task Parallel Jobs”
Speaker at IFIP Performance 2018 Conference, Toulouse France.
- Nov 2018 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
INFORMS Annual Conference speaker. Phoenix, AZ.
- Oct 2018 “Resource Management in the Face of Stochastic Uncertainty”
Distinguished Lecture Series Speaker. University of California, San Diego.

- Host: Yoav Freund.
- Oct 2018 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
Invited Speaker at Joint ISL-OR Seminar, Stanford University, Palo Alto, CA.
Host: Balaji Prabhakar.
- Oct 2018 “SOAP Bubbles: Robust Scheduling Under Adversarial Noise”
Invited Speaker at Allerton Conference.
Urbana-Champaign, IL. Hosts: R. Srikant and Bruce Hajek.
- July 2018 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
Invited Speaker at TTIC Workshop on Data Center Scheduling from Theory to Practice.
TTIC, Chicago, IL. Hosts: Leana Golubchik, Samir Khuller, Barna Saha, Cliff Stein.
- July 2018 “Introductory Queueing Theory Tutorial”
Invited Speaker at Workshop on Operations Research of Biological Systems.
International Center for Theoretical Physics (ICTP), Trieste, Italy. Host: Rami Pugatch.
- June 2018 “SOAP: One Clean Analysis of all Age-Based Scheduling Policies”
Distinguished Speaker at IMA CCS Information Modeling and Control of Complex Systems Workshop.
Ohio State University. Host: Ness Shroff.
- May 2018 “Reducing Latency via Redundant Requests.”
Plenary speaker at 40th Annual Dutch Queueing Colloquium. Hosts: Bert Zwart and Peter van de Ven.
- Mar 2018 Panelist Google Women in Tech Summit.
Host: Kara Boutselis.
- Oct 2017 “Queueing-Theoretic Models for Data Center Performance.”
Talk at 2017 INFORMS Annual Conference, Houston, TX.
- Feb 2017 “Reducing Latency via Redundant Requests.”
Keynote Talk at 2017 MIT LIDS Student Conference, Boston, MA. Host: Martin Zubeldia.
- Oct 2016 “Queueing with Redundant Requests: A More Realistic Model”
University of Illinois, Urbana-Champaign, IL. Host: Yi Lu.
- Sep 2016 “Queueing with Redundant Requests: A More Realistic Model”
IBM T.J. Watson Research Lab, Yorktown Heights, NY. Host: Siva Theja.
- Sep 2016 “Queueing with Redundant Requests: A More Realistic Model”
Columbia University, New York City, NY. Host: Ward Whitt, Karl Sigman.
- Sep 2016 “Queueing with Redundant Requests: A More Realistic Model”
CMU INFORMS Student Chapter presentation. Host: Thiago Serra.
- Aug 2016 “Smarter Task Assignment for Server Farms via Replication”
CanQueue Keynote Presentation.
- July 2016 “Queueing with Redundant Requests: A More Realistic Model”
IBM Spark Center, San Francisco, CA. Host: Nimrod Megiddo.
- June 2016 “Smarter Task Assignment for Server Farms via Replication”
ACM Sigmetrics Keynote Presentation.
- May 2016 “Smarter Task Assignment for Server Farms via Replication”
Simons Institute, U.C. Berkeley. Women in Theory conference.
- May 2016 “Queueing with Redundant Requests: First Exact Analysis”
Distinguished Speaker at IMA CCS Information Modeling and Control of Complex Systems Workshop.
Ohio State University. Host: Ness Shroff.
- Apr 2016 “Towards Better Dispatching Policies for Server Farms: How Replication can Help”
University of California at Berkeley. Host: Rhonda Righter.
- Apr 2016 “Queueing with Redundant Requests: First Exact Analysis”
Stanford University. Host: Peter Glynn.
- Mar 2016 “Towards Better Dispatching Policies for Server Farms: How Replication can Help”
Google, Mountainview. Host: John Wilkes.
- Mar 2016 “Queueing with Redundant Requests: First Exact Analysis”
University of Southern California. Host: Amy Ward.
- Feb 2016 “The Power of d Choices for Job Replication”

- Information Theory and Applications Workshop (ITA), San Diego, CA.
- Nov 2015 “Queueing with Redundant Requests: First Exact Analysis”
Rotman School of Business, U. Toronto. Hosts: Opher Baron and Oded Berman.
- Jul 2015 “Queueing with Redundant Requests: First Exact Analysis”
INFORMS APS Conference, Istanbul.
- Jul 2015 “What Queueing Theory Teaches us about Distributed Systems Design.”
Keynote talk. ICDCS 2015 Conference, Columbus, OH. Hosts: Xiaodong Zhang and Tarek F. Abdelzaher.
- Jun 2015 “What Queueing Theory Teaches us about Systems Design.”
Three-hour tutorial talk at ACM/IEEE ISCA 2015 Conference, Portland, OR.
- May 2015 “Queueing with Redundant Requests: First Exact Analysis.”
Stanford ISL Colloquium. Host: Neal Master.
- May 2015 “Waiting for the Loo and other Queueing Tales.”
Manuel Blum Magic 77 Celebration at Carnegie Mellon University. Host: Lenore Blum.
- Feb 2015 “Queueing with Redundant Requests: First Exact Analysis.”
U.C. Berkeley, IEOR Department. Host: Rhonda Righter.
- Feb 2015 “Queueing with Redundant Requests: First Exact Analysis.”
Information Theory and Applications Workshop (ITA), San Diego, CA.
- Oct 2014 “Dynamic Power Management of Data Centers: Theory and Practice.”
15-719 Lecture in Garth Gibson’s class (CMU).
- July 2014 “The Chicken McNuggets Problem: Solving Linear Diophantine Equations.” Leap talk, CMU. Host: Steven Rudich.
- June 2014 “Dynamic Power Management of Data Centers: Theory and Practice.”
Keynote Presentation for GreenMetrics Conference, Austin, TX. Invited by Adam Wierman.
- May 2014 “Dynamic Power Management of Data Centers: Theory and Practice.”
LCCC Workshop in Cloud Control, Lund University, Sweden. Invited by: Maria Kihl.
- Apr 2014 “Value Driven Load Balancing.” MSR-CMU Mind Swap in NYC.
- Apr 2014 “Dynamic Power Management of Data Centers.” Industry-Academia Partnership (IAP) Cloud Workshop at CMU.
Invited by Jim Ballingall.
- Dec 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
DIMACS Working Group on Algorithms for Green Data Storage. Invited by: Stark Draper.
- Oct 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
EECS Department, Northwestern University, Distinguished Speakers Seminar. Host: Peter Dinda.
- Sept 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
Computer Science Department, Carnegie Mellon University. Host: Guy Blelloch.
- July 2013 “Markov Chains and their use in Computer Systems Analysis.” Andrew’s Leap talk, CMU. Host: Steven Rudich.
- July 2013 “Matching with Flexibility.” INFORMS-APS conference talk.
- July 2013 “Queueing in Computer Science.” INFORMS-APS conference talk.
- July 2013 “Power Management in Data Centers: Recursive Renewal Reward.” INFORMS-APS conference talk.
- May 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
Computer Science Department, Georgia Institute of Technology. Host: Jun Xu.
- May 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
Computer Science Department, Columbia University. Host: Vishal Misra.
- May 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
Computer Science Department, University of Southern California. Host: Leana Golubchik.
- May 2013 “Dynamic Power Management of Data Centers: Theory and Practice.”
Tepper School of Business, CMU. Host: Mustafa Akan.
- Nov 2012 “Solving M/M/k/Setup and Other Open Variants.”
MIT LIDS. Invited by: Devavrat Shah and Eytan Modiano.
- Nov 2012 “Power Management in Data Centers.”
Brandeis University High Tech Alumni Leadership Conference.
Invited by Jordan Pollack.
- Aug 2012 “Solving M/M/k/Setup and Other Open Variants via Recursive Renewal Reward.”
EURANDOM. Eindhoven, Netherlands. Stochastic Activity Month.

- Invited by Onno Boxma.
- June 2012 “Pricing and Queueing.”
MAMA Workshop affiliated with Sigmetrics 2012.
- Feb 2012 “Optimal Threshold-based Pricing for Cloud Services.”
Microsoft-New England, Cambridge, MA. Invited by: Jennifer Chayes.
- July 2011 “AutoScale: Dynamic Power Management in Data Centers.”
Microsoft-New England, Cambridge, MA. Invited by: Jennifer Chayes.
- July 2011 “AutoScale: Dynamic Power Management in Data Centers.”
Microsoft, Redmond, WA. Invited by: Tom McMail.
- July 2011 “Power Management in Data Centers: A Queueing Theorist’s Perspective.”
Basque Center Applied Mathematics (BCAM), Basque County, Spain.
Invited by: Urtzi Ayesta.
- June 2011 “Power Management in Data Centers.”
Distinguished Speaker for Opening of Computer Engineering Center at Technion, Haifa, Israel.
Invited by: Ron Kimmel, Adam Schwartz, Eli Biham.
- April 2011 “Dynamic Power Management in Data Centers.”
HP Labs. Invited by: Kim Keeton.
- April 2011 “Dynamic Power Management in Data Centers.”
Facebook. Invited by: Adam Ward.
- Nov 2010 “Open Problems in Power Management.”
2010 IFIP WG 7.3 Workshop, Namur, Belgium.
- Nov 2010 “Server Farms with Setup Costs.”
2010 Performance Conference, Namur, Belgium.
- Nov 2010 “Power Management in Server Farms.”
2010 INFORMS Conference, Austin, TX, November 2010. Session chair. Invited by: Cathy Xia.
- Oct 2010 “Analytical Modeling of Data Centers to Optimize Performance and Power.”
Distinguished Speaker Series. Computer Science Dept., University of Toronto.
Invited by: Marsha Chechik.
- June 2010 “Why Segregating Short Jobs from Long Ones under High Variability is not always a Win.”
2010 INFORMS Conference, Austin, TX, November 2010. Invited by: Doug Down.
- Feb 2010 “Scheduling in Server Farms: Approaches and Open Problems.” Tutorial speaker.
2010 Information Theory and Applications Workshop, UCSD. Invited by: Tara Javidi.
- Feb 2010 “Why Segregating Short Jobs from Long Ones under High Variability is not always a Win.”
Stanford Management Sciences and Engineering Colloquium. Host: Peter Glynn.
- Nov 2009 “Why Segregating Short Jobs from Long Ones under High Variability is not always a Win.”
Invited speaker at NET-COOP 2009. EURANDOM Institute. Netherlands. Invited by: Onno Boxma.
- Oct 2009 “Surprising Results on Task Assignment in Server Farms under High Variability Workloads.”
Invited Speaker at Operations Research Seminar in ISYE Department, Georgia Tech. Host: Jim Dai.
- Sep 2009 “Why Segregating Short Jobs From Long Ones Under High Variability Is Not Always a Win.”
Allerton 2009 Conference talk, Urbana, Illinois.
- July 2009 “Surprising Results on Task Assignment with High-Variability Workloads.”
INFORMS Applied Probability Society, Ithaca, NY.
- July 2009 “Surprising Results on Task Assignment with High-Variability Workloads.”
Columbia Applied Probability Day, Columbia University, NY. Invited by: Karl Sigman.
- June 2009 “Surprising Results on Task Assignment with High-Variability Workloads.”
SIGMETRICS 2009 conference talk.
- Sep 2008 “Scheduling for Server Farms: Approaches and Open Problems.” Keynote talk, MASCOTS 2008,
Baltimore, MD. Invited by: Carey Williamson and Ethan Miller.
- Jul 2008 “Immigration and Emigration at CMU.” CRA-Snowbird conference.
Invited by: Kim Bruce and Hank Korth.
- Jun 2008 “Scheduling for Server Farms: Approaches and Open Problems.” Keynote talk, SIPEW 2008,
Darmstadt, Germany. Invited by: Kai Sachs.

- May 2008 “Analysis of Join-the-Shortest-Queue Scheduling.” Hisashi Kobayashi Retirement Workshop. Princeton University. Invited by: Brian Mark.
- Oct 2007 “What Performance Modeling Teaches us About Computer System Design.”
Craw Distinguished Lecture Series Speaker. Ohio State University. Host: Bettina Bair.
- Jun 2007 “Analysis of Scheduling in Server Farms.”
Mathematics Department, University of San Diego. Host: Ruth Williams.
- May 2007 “Analysis of Scheduling in Server Farms.”
Keynote Speaker, Operations Research Society of Israel (ORSIS) Conference.
- Apr 2007 “Analyzing Multi-Server Systems with Dependencies:
Cycle Stealing, Priority Queueing, Affinity Scheduling, and More.”
Computer Science Department, Stanford University. Host: Balaji Prabhakar.
- Apr 2007 “Analysis of Web Server Farms with Join-the-Shortest-Queue Routing.”
Operations Research Group, Stanford University. Host: Balaji Prabhakar.
- Mar 2007 “Auction-Based Scheduling of the TeraGrid.”
CMU/Pitt/Intel Workshop on Next-Generation Computing. Host: Peter Lee.
- Jan 2007 “Analysis of Scheduling in Server Farms.”
Plenary Speaker, Lunteren Conference, Netherlands.
- Jan 2007 “Works-in-Progress: From Web Server Farms to Limited Resource Sharing.”
EURANDOM, Netherlands. Host: Onno Boxma.
- Sep 2006 “What Performance Modeling Teaches Us About Computer Systems Design.”
CMU Student Theory Lunch.
- Aug 2006 “Analysis of Scheduling in Server Farms.”
Invited talk. IBM Watson Research Labs. Hosts: Asser Tantawi and Mark Squillante.
- Jul 2006 “What Performance Modeling Teaches us about Computer Systems Design.”
Invited talk. Microsoft Research, Redmond, WA. Hosts: Jennifer Chayes and Ken Church.
- Jun 2006 “Analysis of Scheduling in Server Farms.”
Invited tutorial. Sigmetrics 2006, San Malo, France. Host: Thomas Bonald.
- May 2006 “Nearly Insensitive Bounds on SMART Scheduling.”
Invited talk at Uri Yechiali’s retirement workshop, Shefayim, Israel.
- Apr 2006 “Task Assignment for Server Farms.”
Guest lecturer in Satya’s class, CMU.
- Apr 2006 “What Performance Modeling Teaches Us About Computer Systems Design.”
CS 50 Anniversary Celebration, CMU.
- May 2006 “Dimensionality Reduction Analysis of Cycle Stealing, Priority Queueing and More.”
Microsoft, Cambridge, England. Host: Dushyanth Narayanan.
- Mar 2006 “Dimensionality Reduction Analysis of Cycle Stealing, Threshold Policies and More.”
Imperial College, London, England. Host: Peter Harrison.
- Jan 2006 “Dimensionality Reduction Analysis of Cycle Stealing, Threshold Policies and More.”
University of California, Berkeley. Host: Rhonda Righter.
- Sep 2005 “Scheduling Your Network Connections (the SYNC Project).”
Polytechnic University, NYC. Host: Keith Ross and Torsten Suel.
- Jul 2005 “Nearly Insensitive Bounds on SMART Scheduling.”
INFORMS Applied Probability conference, Ottawa, Canada.
- June 2005 “Robustness of Threshold Policies in Beneficiary-Donor Model.”
MAMA Workshop, Banff, Canada.
- Jun 2005 “Nearly Insensitive Bounds on SMART Scheduling.”
ACM Sigmetrics Conference, Banff, Canada.
- Feb 2005 “Recursive Dimensionality Reduction Applied to the Analysis of Multiserver Scheduling.”
Lehigh University. Host: Andrew Ross.
- Feb 2005 “Recursive Dimensionality Reduction Techniques Applied to the Analysis of Multiserver Scheduling.”
Columbia University – Industrial Engineering and Operations Research department. Host: Cliff Stein.
- Feb 2005 “Scheduling Your Network Connections (the SYNC Project).”

- Columbia University – CS department. Host: Vishal Misra.
- Feb 2005 “Recursive Dimensionality Reduction Applied to the Analysis of Multiserver Scheduling.”
University of Pittsburgh – Industrial Engineering and Operations Research. Host: Jay Rajgopal.
- Dec 2004 “Recursive Dimensionality Reduction Techniques Applied to the Analysis of Multiserver Scheduling.”
Tel Aviv University – Mathematics department. Host: Uri Yechiali.
- Dec 2004 “Scheduling Your Network Connections (the SYNC Project).”
Tel Aviv University – CS department. Host: Hanoch Levy.
- Jul 2004 “Recursive Dimensionality Reduction Techniques Applied to the Analysis of Multiserver Scheduling including Cycle Stealing, Priority Queueing, and Threshold Policies”
Invited speaker. Stochastic Networks Conference, Montreal, Canada.
- Jul 2004 “Recursive Dimensionality Reduction Techniques Applied to the Analysis of Multiserver Scheduling including Cycle Stealing, Priority Queueing, and Threshold Policies.”
MIT – EECS and LIDS departments. Host: Charles Leiserson.
- Jul 2004 “Scheduling Your Network Connections (the SYNC Project).”
MIT – EECS and LIDS departments. Host: Dina Katabi.
- Jun 2004 “Recursive Dimensionality Reduction Techniques Applied to the Analysis of Multiserver Scheduling including Cycle Stealing, Priority Queueing, and Threshold Policies.”
Invited talk. EURANDOM. Eindhoven, Netherlands. Host: Onno Boxma.
- Jun 2004 “Exact closed-form analysis of Queues with Fluctuating Load.”
Invited talk. EURANDOM. Eindhoven, Netherlands. Host: Geert-jan van Houtum.
- Nov 2003 “Scheduling Static and Dynamic Web Requests.”
Invited talk. Microsoft Research. Host: John Dunagan.
- Nov 2003 “Scheduling Your Network Connections (the SYNC Project).”
Invited talk. University of Washington. Host: Susan Eggers.
- Nov 2003 “Task Assignment Policies for Server Farms: Including New Analysis of Cycle Stealing.”
Invited talk. University of Washington. Host: Susan Eggers.
- Oct 2003 “Scheduling Your Network Connections (the SYNC Project).”
Invited talk. Brown University. Host: Anna Lysyanskaya.
- Oct 2003 “Task Assignment Policies for Server Farms: Including New Analysis of Cycle Stealing.”
Invited talk. Brown University. Host: Anna Lysyanskaya.
- Jul 2003 “Scheduling in Multiserver Systems.”
Talk at Aladdin SITE Review.
- Jun 2003 “Applying to Graduate School in Computer Science.”
Talk at Aladdin first College Workshop.
- May 2003 “Task Assignment Policies for Server Farms: Including New Analysis of Cycle Stealing.”
Invited talk. EMC² Corporation. Hosts: Peter Lauterbach, Ric Wheeler.
- Apr 2003 “Scheduling Static and Dynamic Web Requests.”
Invited talk. IBM T.J. Watson Research Lab. Host: Arun Iyengar.
- Nov 2002 “Load balancing algorithms in Networks of Workstations.”
Invited talk. Carnegie Mellon, Concert Group Meeting. Host: Karl Crary.
- Oct 2002 “Reduction of 2D-infinite Markov Chains to 1D-infinite Markov Chains.”
Invited talk. Carnegie Mellon, Aladdin Meeting. Host: Guy Blelloch.
- Jun 2002 “Task Assignment with Immediate Dispatch.”
SPAA 03 Conference.
- May 2002 “Task Assignment with Cycle Stealing under Central Queue.”
ICDCS 03 Conference.
- Jul 2002 “Investigating Unfairness in Scheduling Algorithms.”
Invited talk. Conference on Stochastic Networks. Stanford University.
- Jul 2002 “Scheduling Your Network Connections.”
Invited talk. Computer Science Seminar. University of California Berkeley. Host: Matt Welsh.
- Nov 2001 “Scheduling Your Network Connections.”
Invited talk. Computer Networking Seminar. Caltech University. Host: Steven Low.

- Sep 2001 “Scheduling Your Network Connections.”
Invited talk. Computer Systems Seminar. University of Illinois Urbana-Champaign.
Host: Michael Loui.
- Sep 2001 “Scheduling Solutions for Coping with Transient Overload.”
Invited talk. Communications Seminar. University of Illinois Urbana-Champaign.
Host: Michael Loui.
- Jun 2001 “Scheduling Your Network Connections.”
LIRMM – Laboratoire d’Informatique, de Robotique et de Microelectronique de Montpellier, France.
Host: Alain Jean-Marie.
- Jun 2001 “Scheduling Your Network Connections.”
Invited talk. Instituto Superior Tecnico, Lisbon. Host: Carlos Filipe Gomes Bisbo.
- Jun 2001 “Task Assignment with Unknown Duration.”
Invited talk. Instituto Superior Tecnico, Lisbon. Host: Carlos Filipe Gomes Bisbo.
- Jun 2001 “Analysis of M/G/1/PS and M/G/1/SRPT under Transient Overload.”
Invited talk at 3rd Workshop on Mathematical performance Modeling and Analysis (MAMA 2001).
Invited by: Mark Squillante.
- Jun 2001 “Scheduling Your Network Connections.”
M.I.T. Lab for Computer Science. Hosts: Vincent Chen and John Guttag.
- Jun 2001 “How Heavy Tails Impact Job Scheduling.”
Invited talk at Job Scheduling Strategies for Parallel Processing Workshop. Cambridge, MA.
Invited by: Larry Rudolph.
- May 2001 “Scheduling Your Network Connections.”
Invited talk. IBM Corporation T.J. Watson Research Lab.
- Mar 2001 “Scheduling Your Network Connections.” Invited talk at DIMACS Workshop on Resource Management
and Scheduling in Next Generation Networks. Invited by: Rajmohan
- Mar 2001 “Scheduling Your Network Connections.” Duke University. Host: Ron Parr.
- Jan 2001 “Scheduling Your Network Connections.” Invited talk. HP Labs. Host: John Wilkes.
- Jan 2001 “Scheduling Your Network Connections.” Invited talk. Stanford University. Host: Mary Baker.
- Jun 2000 “Task Assignment with Unknown Duration.”
Invited talk at 2nd Workshop on Mathematical performance Modeling and Analysis (MAMA 2000).
Invited by: Mark Squillante.
- Oct 1999 “Connection Scheduling in Web Servers.”
Presented at USITS ’99.
- Jun 1999 “The Case for SRPT Scheduling in Web Servers.”
End-to-End Research Group Meeting. MIT Lab for Computer Science. Host: David Clark.
- Jun 1999 “Task Assignment with Unknown Duration.”
Invited paper at ASA-IMS Conference on Applications of Heavy Tailed Distributions
in Economics, Engineering and Statistics. Washington DC.
- Apr 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: Carnegie Mellon University. Host: Hui Zhang.
- Apr 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: Caltech Institute of Technology. Host: Mani Chandy.
- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: AT&T Research Labs. Host: Jennifer Rexford.
- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: Johns Hopkins. Host: Gerald Mason.
- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: Massachusetts Institute of Technology. Host: Charles Leiserson.
- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: University of Texas, Austin. Host: Mike Dahlin.
- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: New York University. Host: Richard Cole.

- Mar 1999 “The Effect of Heavy-Tailed Job Size Distributions on System Design.”
Interview talk: UMASS-Amherst. Host: Don Towsley.
- Feb 1999 “Task Assignment in a Distributed Server System.”
Brandeis University. Host: Liuba Shrira.
- Jan 1999 “Task Assignment in a Distributed Server System.”
Invited talk. University of California, Berkeley. Host: Manuel Blum.
- Oct 1998 “Task Assignment in a Distributed Server System.”
M.I.T., Theory Group. Host: Tom Leighton.
- Sep 1998 “Task Assignment in a Distributed Server System.” Presented at the 10th International
Conference on Modeling Techniques and Tools for Computer Performance Evaluation.
- Mar 1998 “Task Assignment in a Distributed Web Server.” Cisco Systems.
- Mar 1998 “Task Assignment in a Distributed Supercomputing Server.”
MIT Xolas Supercomputing Group. Host: Charles Leiserson.
- Mar 1998 “Task Assignment in a Distributed Web Server.” Lucent Technologies.
- Feb 1998 “Task Assignment in a Distributed Web Server.” Arrowpoint Communications.
- Feb 1998 “Task Assignment for a Distributed Supercomputing Server.”
Cornell University. Host: Sid Resnick.
- Feb 1998 “Task Assignment in a Distributed Server System.” IBM.
- Jan 1998 “Task Assignment in a Distributed Server System.” Bay Networks.
- Jun 1997 “The Impact of Heavy-Tailed Distributions on Load Balancing in Web Servers.”
Invited to the INFORMS Applied Probability Conference, Cambridge, MA. Invited by Mark Crovella.
- Jan 1997 “Exploiting Process Lifetime Distributions for Dynamic Load Balancing.”
AT&T Research Lab. Host: Gisli Hjalmysson.
- Oct 1996 “Exploiting Process Lifetime Distributions for Dynamic Load Balancing.”
Boston University. Host: Murad Taqqu.
- May 1996 “Exploiting Process Lifetime Distributions for Dynamic Load Balancing.” Presented at
the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems.
- May 1996 “Exploiting Process Lifetime Distributions for Dynamic Load Balancing.”
University of California, San Diego. Hosts: Fran Berman and Russell Impagliazzo.
- Jan 1996 “Bounding Delays in Packet-Routing Networks.” IBM Almaden. Host: Eli Upfal.
- May 1995 “Bounding Delays in Packet-Routing Networks.” Presented at the 27th Annual Symposium
Theory of Computing.
- Jan 1994 “Queueing Analysis of Oblivious Packet-Routing Networks.” Presented at the Fifth Annual
ACM-SIAM Symposium on Discrete Algorithms.

PROFESSIONAL SERVICE

SIG Chair:

2023-27 SIG Chair for ACM Sigmetrics.

Fellow of: ACM, IEEE

Member of: ACM-Sigmetrics, Informatics, IFIP W.G. 7.3.

Conference Chair:

2013 Conference Chair SIGMETRICS 2013.

2005 Conference Co-Chair and Program Co-Chair CMU/Eindhoven Collaborative Workshop on Queueing.
See www.aladdin.cs.cmu.edu/workshops/eindhoven/.

2004 Conference Co-Chair and Program Co-Chair WORKshop on Multiserver Scheduling (WORMS04).
See www.cs.cmu.edu/~harchol/WORMS04/.

Technical Program Chair:

2007 Program co-Chair SIGMETRICS 2007.

2007 Program co-Chair QEST 2007.

SIG Officer:

2005-2007 Secretary/Treasurer ACM SIGMETRICS.

Editor:

2019 - 2022 Associate Editor *Queueing Systems: Theory and Applications (QUESTA)*.

2014 - 2016 Associate Editor *ACM Transactions on Modeling and Performance of Computer Systems*.

2008 - 2010 Associate Editor *Performance Evaluation* journal.

2007 Guest Editor Performance Evaluation Review, Special Issue on Scheduling, March 2007.

Technical Program Committee and Related:

2023 Technical Program Committee: IFIP PERFORMANCE'23.

2022 Technical Program Committee: Stochastic Network Conference '22.

2021 Technical Program Committee: SIGMETRICS '21, IFIP PERFORMANCE '21.

2020 Technical Program Committees IFIP PERFORMANCE '20, Stochastic Network Conference '20.

Meet the Stars event PERFORMANCE '20.

2018 Technical Program Committees: SIGMETRICS '18.

2017 Sigmetrics Rising Star Award Committee.

2017 Technical Program Committees: SIGMETRICS '17.

2015 Sigmetrics Rising Star Award Committee.

2015 Technical Program Committees: SIGMETRICS '15.

2014 Served on Review Board for Heidelberg Laureate Forum .

2014 Technical Program Committees: SIGMETRICS '14, IFIP PERFORMANCE '14, ECQT '14.

2013 Technical Program Committees: IC2E '13.

2012 Technical Program Committees: SIGMETRICS '12, IEEE IGCC '12.

2011 Technical Program Committees: PERFORMANCE '11.

2010 Technical Program Committees: HOTMETRICS '10, PERFORMANCE '10.

2009 Technical Program Committees: SIGMETRICS '09, MAMA '09, HotPower '09.

2009 Sigmetrics Rising Star Award Committee.

2008 Technical Program Committees: SIGMETRICS '08, SIPEW '08, HotPower '08, MAMA '08.

2008 Sigmetrics Rising Star Award Committee.

2007 Technical Program Committees: SIGMETRICS '07, QEST '07, MAMA '07.

2006 Technical Program Committees: SIGMETRICS '06, WWW '06, QEST '06.

2005 Technical Program Committees: SIGMETRICS '05, WWW '05.

2004 Technical Program Committees: SIGMETRICS '04, WWW '04.

2003 Technical Program Committee: SIGMETRICS '03.

2002 Technical Program Committees: SIGMETRICS '02, Performance '02.

2002 Publicity Chair: Performance '02.

2001 Tutorial Chair: SIGMETRICS '01.

2000 Technical Program Committee: SIGMETRICS '00.

Panelist:

Apr 2019 Formal Methods in the Field Operating and Distributed Systems. Headed by Samee Khan.
Jun 2014 DOE Panel on Analytical Modeling of Distributed Systems. Headed by Richard Carlson.
Jan 2011 DOE ASCR Early Career Research Panel. Headed by Sonia Sachs.
Jan 2002 NSF CCR-Operating Systems and Compilers Panel. Headed by Randy Chow.
Feb 2000 NSF ITR Panel. Headed by Darlene Fisher.

UNIVERSITY SERVICE

2023-pres Created Supplemental Mentor Program for Female and Non-Binary Students.
2022-pres Created weekly SCS Graduate Student Womens Lunch. Helping with organization and funding.
2022-pres Managing CSD ACO Program.
2022-pres Heavily involved in SCS Undergrad Performance.
– Tracking SCS Undergrad Performance across classes.
– Proposing solutions via changes to both Admissions and Curriculum.
– Researching solutions from other universities.
2018-pres Mentor for faculty member Weina Wang.
2012-pres Mentor for faculty member David Kosbie.
2020-pres Active member CSD DEI Committee.
– Heading Sub-Committee on Reporting/Tracking Infractions.
– Member Sub-Committee on Increasing Diversity in PhD Applications
2020-pres Active member of SCS Women Faculty group.
2020-2022 SCS-wide R&P Committee.
2022 Hosted faculty recruiting events at my home.
2021-2022 Member Undergraduate Math Preparedness Committee, led by Tom Cortina.
2021 Chair for Dave Kosbie Promotion.
Aug 2021 Host for CSD IC Party at my home.
2021 Organized Theory Open House. Revised full slide deck.
2010-2020 AoPS Math Jam for Recruiting strong Math Undergrads.
2020 Leading Theory Consortium Initiative.
2020 Chair for Weina Wang R&P.
2020 CSD representative for SCS Council.
2020 TartanHacks Judge.
2019 Mentorship of faculty member Jean Yang.
2019 CSD representative for SCS Council.
2018-19 SCS Dean Search Committee.
2018 Managing CSD ACO Program.
Aug 2018 Host for CSD IC Party at my home.
2018 Faculty Member Representative for Omega Rho – Undergrad Honor Society of OR.
2018 Mentorship of faculty member Jean Yang.
2018 CSD representative for SCS Council.
Apr 2018 Women@SCS Speaker and Interviewee.
Feb 2018 Hosted Women@SCS Brunch.
2017-18 CSD Hiring Committee.
2017 Mentorship of faculty member Jean Yang.
Jul 2017 Hosted Goodbye Party for Avrim Blum.
Apr 2017 Celebration of Diversity Recruiting.
2017 Managing CSD ACO Program.
2017 CSD representative for SCS Council.
2017 Faculty Member Representative for Omega Rho – Undergrad Honor Society of OR.

2017 SCS Undergraduate Stress Committee.
 2017 Managing CSD ACO Program.
 2016-17 CSD Admissions Committee.
 2016-17 CSD Hiring Committee.
 2016 Managing CSD ACO Program.
 2016 Mentorship of faculty member Bernhard Haeupler.
 2016 Mentorship of faculty member Jean Yang.
 2016 Faculty Member Representative for Omega Rho – Undergrad Honor Society of OR.
 2016 SCS Undergraduate Stress Committee.
 2016 CSD representative for SCS Council.
 Aug 2016 Host CSD IC party at my home.
 2015 CSD representative for SCS Council.
 2015 Mentorship of faculty member Bernhard Haeupler.
 2015 Managing CSD ACO Program.
 Aug 2015 Host CSD IC party at my home.
 Aug 2015 Organized CSD Faculty & Grad Student Picnic Honoring Graduating PhD students.
 June 2015 Initiating CSD-wide Grad Student Teaching Award. Ran CSD-wide award proces.
 May 2015 Member of SCS-wide Perlis Award Committee.
 May 2015 Hosted Women@SCS brunch at my home.
 Mar 2015 Presentation to undergraduates on “Applying to Grad School.”
 2014/15 Initiating CSD Phone Interview process.
 2014 Managing the reinstatement of the ACO-CS PhD program.
 2014/15 Mentorship of faculty member Bernhard Haeupler.
 2014 Chairing promotion committee for David Kosbie.
 July 2014 Presentation at Andrews Leap to High School students.
 April 2014 Presentation to undergraduates on “Applying to Grad School.”
 Aug 2013 Talk at CSD IC on Finding an Advisor.
 Aug 2013 Host CSD IC party at my home.
 2013 Served on promotion committee for Ariel Proccaccia.
 July 2013 Presentation at Andrews Leap to High School students.
 Oct 2012 Presentation to 100 undergraduates on “Applying to Grad School.”
 Aug 2012 Host CSD IC party at my home.
 2012 Teaching Track Search Committee.
 Jan 2012 Host CSD Women’s party at my home for female grad students and faculty.
 Nov 2011 Host CSD Women’s party at my home for female grad students and faculty.
 Aug 2011 Host CSD IC party at my home.
 Jan 2011 Exam questioner/grading for Tepper School qualifiers.
 2011 CSD Masters Planning Committee.
 2008 - 2011 Associate Department Head for Graduate Education, Computer Science.
 Primary responsibilities included:
 – Structuring Computer Science PhD Program and Creating all Rules/Regulations for PhD Program
 – Management of Graduate Admissions
 – Management of Graduate Recruiting & Open House
 – Management of Immigration Course and Emigration Seminars
 – Running Black Friday meetings, every semester
 – Assignment of graduate students to faculty
 – Mentoring of graduate students and meeting with them
 – Managing Doctoral Review Committee
 2009/10 Selection Committee for CSD Department Head.
 Aug 2010 Host CSD IC party at my home.
 Jan 2010 Exam questioner/grading for Tepper School qualifiers.
 2010 Doctoral Review Committee (DRC).

Aug 2009 Host CSD IC party at my home.
 Jan 2009 Exam questioner/grading for Tepper School qualifiers.
 2009 Doctoral Review Committee (DRC).
 Nov 2008 Presentation to 100 undergraduates on "Applying to Grad School."
 Spr 2008 CSD Retreat Planning Committee.
 Jan 2008 Exam questioner/grading for Tepper School qualifiers.
 2008 Doctoral Review Committee (DRC).
 2007 CSD Finances committee.
 2006/07 Gates classroom committee.
 2007 Doctoral Review Committee (DRC).
 Jun 2007 Speaker at Breakfast for Volunteer Alumni Leaders.
 Jan 2007 Exam questioner/grading for Tepper School qualifiers.
 Nov 2006 Presentation to 100 undergraduates on "Applying to Grad School."
 Aug 2006 Host CSD IC party at my home.
 Apr 2006 Presentation to 100 undergraduates on "Applying to Grad School."
 Apr 2006 Speaker at Carnegie Mellon 50 Year Anniversary Symposium.
 Jan 2006 Exam questioner/grading for Tepper School qualifiers.
 2006 Doctoral Review Committee (DRC).
 Apr 2005 Chaired committee for Allen Newell Graduate Teaching Award.
 Apr 2005 Chaired committee for Herbert A. Simon Award.
 Apr 2005 Presentation to 100 undergraduates on "Applying to Grad School."
 Mar 2005 Presentation at Open House Graduate Student Recruiting.
 Spr 2005 CS hiring committee: Performance evaluation area.
 2005 Doctoral Review Committee (DRC).
 2004-2005 Instrumenting SCS Perlis Teaching Assistant Award.
 Jan 2005 Exam questioner/grading for Tepper School qualifiers.
 Dec 2004 Undergraduate math colloquium speaker. Host: Deborah Brandon.
 Sep 2004 Women in SCS Road Show.
 Aug 2004 Host CSD IC party at my home.
 May 2004 Judge for Herbert A. Simon Award.
 Apr 2004 Presentation to 100 undergraduates on "Applying to Grad School."
 Jan 2004 Exam questioner/grading for Tepper School qualifiers. Also, reader on some summer papers.
 2004 Doctoral Review Committee (DRC).
 Aug 2003 Host CSD IC party at my home.
 May 2003 Judge for Newell Award.
 Apr 2003 Presentation to 100 undergraduates on "Applying to Grad School."
 Apr 2003 Heavy involvement in faculty recruiting. Host parties.
 Jan 2003 Exam questioner/grading for Tepper School qualifiers.
 2003 Doctoral Review Committee (DRC).
 Nov 2002 Presentation to 100 undergraduates on "Applying to Grad School."
 Aug 2002 Host CSD IC party at my home.
 Aug 2002 Provide interview for undergraduate recruitment tape.
 July 2002 Committee to evaluate department chair position.
 May 2002 Joined SCS Web advisory committee.
 May 2002 Heavy involvement in faculty recruiting. Host parties.
 Apr 2002 Presentation to 150 undergraduates on "Applying to Grad School."
 Mar 2002 Presentation at Open House Graduate Student Recruiting.
 Jan 2002 Exam questioner/grading for Tepper School qualifiers.
 Nov 2001 Sent letters to top 130 CMU alumni encouraging them to apply to grad school at CMU.
 Nov 2001 Presentation to CMU Women Graduate Organization on "Advising and Being Advised."
 Aug 2001 Host CSD IC party at my home.
 Aug 2001 Presentation at Andrews Leap to High School students.

Apr 2001 Heavy involvement in faculty recruiting. Host parties/candidates.
 Mar 2001 Presentation at Open House Graduate Student Recruiting.
 Nov 2000 Presentation at Freshman IC.
 Nov 2000 Led Panel at CS Immigration Course titled
 "How to be a Great Advisor/ How to be a Great Grad Student."
 Nov 2000 Sent letters to top 40 CMU alumni encouraging them to apply to grad school at CMU.
 Aug 2000 Host CSD IC party at my home.
 July 2000 Presentation at Andrews Leap to High School students.
 Apr 2000 Presentation to prospective undergraduates at Undergrad Open House.
 Mar 2000 Presentation at Open House Graduate Student Recruiting.
 Oct 1999 Presentation to Women in SCS group on "Getting the Most Out of a Conference."
 Aug 1999 Presentation at Andrews Leap to High School students.

COMMUNITY SERVICE/OUTREACH

2018 Google Women in Tech Summit speaker.

2008-pres Created first-ever Western PA American Regional Mathematics League (ARML) team.
 Includes 25 high schoolers from all over Western PA.
 Our team ranked 3rd in the nation in the B division in the 2012 competition and is now in the A division.
 Organize weekly 3-hour team practices and trips to national math competitions throughout the year.
 Receive annual funding from Jump Trading, Chicago, IL.

2004-2005 Volunteer teacher of Elementary School National Math Olympiad Contest at Winchester-Thurston school.
 Teach every Monday afternoon.

2002-2004 Volunteer math teacher at Winchester-Thurston School, Pittsburgh, PA.
 Teach advanced mathematics to 2nd graders. Wednesday mornings, every other week.

1988-1990 Volunteer Teacher at City Year, Stilling St., Boston, MA.
 Volunteer Teacher Boston Partners in Education.
 Teach math to disadvantaged students one evening per week.

CONTRACT AND GRANT SUPPORT

Current:

Title: NSF-CMMI: New Approaches to Multiserver Scheduling
 PI : Mor Harchol-Balter.
 co-PI: Ziv Scully.
 Agency: NSF-CMMI.
 Duration: Sept '23 - Aug '26.
 Amount: \$473,545.00.
 Support: 25%SU.
 Award No: 1123095
 PM: Georgia-Ann Klutke.

Past:

Title: NSF-CMMI: Optimal Scheduling of Parallelizable Jobs in Cloud Computing Environments
PI : Mor Harchol-Balter.
co-PI: Ben Moseley.
Agency: NSF-CMMI.
Duration: Jan '20 - Dec '22.
Amount: \$549,470.00.
Support: 25%SU.
Award No: 1938909
PM: Georgia-Ann Klutke.

Title: NSF-CSR: Medium: Collaborative Research: Foundations of Cache Network Operations for Content Delivery
PIs : Mor Harchol-Balter and Ramesh Sitaraman.
Agency: NSF-CSR.
Duration: Sept '18 - Sept '22.
Amount: \$1,200,000.00.
Support: 33%SU.
Award No: 1763701 (originally 180341).
PM: Samee Khan.

Title: Predictable Scheduling for Borg
PI : Mor Harchol-Balter.
Agency: Google Faculty Award
Duration: August '20 - unrestricted
Amount: \$70,000
Award No: Google

Title: Facebook Graduate Fellowship.
PI : Benjamin Berg (my graduate student).
Duration: Jan '19 - Jan '21.

Title: NSF-XPS: FULL: Bridging Parallel and Queueing-Theoretic Scheduling – SUPPLEMENT
PIs : Mor Harchol-Balter.
Agency: NSF-XPS – SUPPLEMENT.
Duration: July '19 - July '20
Amount: \$150,000.
Support: 33%SU.
Award No: 1629444.
PM: Tracy Kimbrel.

Title: RobinHood: Cache Sharing System
PI : Mor Harchol-Balter.
Agency: Microsoft Faculty Award
Duration: April '18 - unrestricted
Amount: \$30,000
Award No: Microsoft

Title: NSF-XPS: FULL: Bridging Parallel and Queuing-Theoretic Scheduling.
PIs : Guy Blleloch, Umut Acar, Mor Harchol-Balter.
Agency: NSF-XPS.
Duration: July '16 - July '19
Amount: \$825,000.
Support: 33%SU.
Award No: 1629444.
PM: Tracy Kimbrel.

Title: NSF Graduate Fellowship.
PI : Ziv Scully (my graduate student).
Duration: September '16 - September '19.

Title: NSF-CMMI:Reducing Latency by Replicating Jobs.
PI : Mor Harchol-Balter.
Agency: NSF-CMMI.
Duration: Sept '15 - Sept '18
Amount: \$299,711.
Support: 33%SU.
Award No: 1538204.
PM: Michael Fu.

Title: Performance Analysis and Design of Computer Systems.
PI : Mor Harchol-Balter.
Agency: Facebook Faculty Award
Duration: Sept '15 - unrestricted
Amount: \$30,000
Award No: Facebook

Title: When Many Workloads Share Networked Storage: How to Guarantee Tail Latency SLOs.
PI : Mor Harchol-Balter.
Agency: Google Faculty Research Award: Unrestricted Gift
Duration: February '15 - February '16.
Amount: \$67,000.
Award No: Google – TBA.

TITLE: NSF REU Supplement
PI: Mor Harchol-Balter.
Agency: NSF-CSR Supplement to NSF-CSR-1116282.
Duration: Summer 2015.
Amount: \$8000.
Award No: 25851.2.1122199.
PM: Mimi McClure.

Title: NSF-CMMI:Priority Pricing for Profit Maximization Given Strategic, Delay-Sensitive Customers with a Continuum of Types.
PI : Mor Harchol-Balter (with Mustafa Akan, co-PI).
Agency: NSF-CMMI
Duration: Aug '13 - Aug '16
Amount: \$290,000.
Support: 25%SU.
Award No: 1334194.
PM: Sheldon Jacobson.

Title: Intel Pittsburgh Cloud Computing ISTC Fellowship for Student.
PI : Timothy Zhu (my graduate student).
Agency: Intel Pittsburgh.
Duration: September '13 - '16
Amount: Yearly support for 1 Ph.D. student

Title: NSF-CSR:Small Simple Dynamic Traffic-Oblivious Power Management for Multi-Tier Web Clusters.
PI : Mor Harchol-Balter.
Agency: NSF-CSR
Duration: Aug '11 - Aug '15
Amount: \$373,472.
Support: 33%SU.
Award No: 1116282.
PM: Krishna Kant.

Title: NSF Graduate Fellowship.
PI : Kristy Gardner (my graduate student).
Duration: September '12 - September '15.

Title: Intel Pittsburgh Cloud Computing ISTC Fellowship for Student.
PI : Anshul Gandhi (my graduate student).
Agency: Intel Pittsburgh.
Duration: September '11 - '13
Amount: Yearly support for 1 Ph.D. student

Title: Western PA ARML Team Funding.
PI : Mor Harchol-Balter
Agency: Jump Trading, Inc.
Duration: Sept '12 - Sept '13.
Amount: \$12,000.
Award No: ARML

Title: Maximizing Revenue in a Multi-Class Environment with Unknown Customer Valuations and Unknown Delay Sensitivities.
PI : Mor Harchol-Balter.
Agency: Microsoft-CMU Center for Computational Thinking

Duration: Aug '12 - Aug '13.
Amount: \$80,000.
Award No: Microsoft

Title: NSF-CSR: Student Travel Support for Sigmetrics 2013.
PI : Mor Harchol-Balter.
Agency: NSF-CSR
Duration: Jan '13 - Jan '14
Amount: \$10,000.
Award No: 1300202.
PM: Krishna Kant.

Title: Maximizing Pricing Based on Current Delay.
PI : Mor Harchol-Balter.
Agency: Microsoft-CMU Center for Computational Thinking
Duration: Aug '11 - Aug '12.
Amount: \$74,736.
Award No: Microsoft

Title: Western PA ARML Team Funding.
PI : Mor Harchol-Balter.
Agency: Jump Trading, Inc.
Duration: Sept '11 - Sept '12.
Amount: \$7,500.
Award No: ARML

Title: Kauffman Foundation Commercialization Grant.
PI : Awarded to my student: Anshul Gandhi.
Agency: Kauffman Foundation
Duration: Jan '11 - Jan '12.
Amount: \$25,000.
Award No: Kauffman

Title: Performance Analysis and Design of Computer Systems.
PI : Mor Harchol-Balter.
Agency: Yahoo! Faculty Research Award
Duration: Sept '10 - Sept '11.
Amount: \$30,000
Award No: Yahoo

Title: Energy Efficient and Proportional Datacenter Computing.
PI : Dave Andersen and Mor Harchol-Balter.
Agency: Google Focused Research Award
Duration: February '10 - February '11.
Amount: \$100,000.
Award No: Google 22796.1-2.5003506

Title: Power Management in Data Centers.
PI : Mor Harchol-Balter.
Agency: Microsoft-CMU Center for Computational Thinking
Duration: Aug '10 - Aug '11.
Amount: \$75,000.
Award No: Microsoft 5004487 Project 24592

Title: Building a Power-Efficient Load Balancer.
PI : Mor Harchol-Balter.
Agency: Yahoo! Faculty Research Award
Duration: Sept '09 - Sept '10.
Amount: \$25,000
Award No: 5004676

Title: IBM Faculty Award
PI : Mor Harchol-Balter.
Agency: IBM Faculty Research Award
Duration: June '08 - May '09.
Amount: \$30,000
Award No: 5002343

Title: POW! Project
PI : Mor Harchol-Balter.
Agency: Innovation Works
Duration: June '09 - Sept '09.
Amount: \$15,000
Award No: N/A

Title: Optimizing Server Farm Performance in Power-Constrained Environments.
PI : Mor Harchol-Balter.
Agency: The Technology Collaborative of Pittsburgh.
Duration: January '09 - June '10.
Amount: \$124,000.
Support: 66%SU.
Award No: 1010970

Title: Optimizing Server Farm Performance in Power-Constrained Environments.
PI : Mor Harchol-Balter.
Agency: Matching funds (in kind) from Project Olympus and the Idea Foundry, Pittsburgh.
Duration: January '09 - June '10.
Amount: \$30,000.

Title: Efficient Power Management by Alternating Between "Extreme" Power States.
PI : Mor Harchol-Balter.

Agency: CMU Technology Transfer Center & Project Olympus.
Duration: January '09 - June '10.
Amount: \$40,000.
Award No: 2001801

Title: Intel Pittsburgh Lab Fellowship for Student.
PI : Anshul Gandhi (my graduate student).
Agency: Intel Pittsburgh.
Duration: September '08 - May '10.
Amount: 1.5 years graduate student support.

Title: SMA/PDOS Collaborative Research: Design, Analysis, and Control of Adaptive Sharing Mechanisms.
co-PI : Mor Harchol-Balter (with Vishal Misra, Dan Rubenstein, Ed Coffman, and Predrag Jelenkovic).
Agency: NSF.
Grant No: NSF CCR-0615262. Division of Computer and Network Systems. Computer Systems.
Duration: August '06 - July '09.
Amount: \$1,000,000.
Support: 33%SU.
Center No: 15653.1.1120828.

Title: SGER: Collaborative Research: CSR-SMA:
New Breakthrough in Analyzing Limited Resource Sharing Systems.
PI : Mor Harchol-Balter (with Jim Dai and Bert Zwart).
Agency: NSF.
Grant No: NSF CNS-0719106.
Duration: August '07 - August '08.
Amount: \$100,000.
Support: 33%SU.
Center No: 17666.1.1121054

Title: Scheduling the TeraGrid.
PI : Mor Harchol-Balter.
Company: Microsoft Research. Award Number: 16101.
Duration: May '07 - May '08.
Amount: \$75,000.
Support: Winner of Microsoft Breakthrough Research Grant. Unrestricted Gift.

Title: CAREER: The Impact of Resource Scheduling on Improving Server Performance.
PI : Mor Harchol-Balter.
Agency: NSF.
Grant No: NSF CCR-0133077. CAREER Award. Division of Computer and Network Systems. Distributed Systems.
Duration: March '02 - February '08.
Amount: \$350,000.
Support: 33%SU.
Center No: 8061.1.1120231.

Title: ITR: Improving the Performance of Web Servers under Overload.
PI : Mor Harchol-Balter.
Agency: NSF.
Grant No: ITR-0313148. Division of Computer and Communication Foundations.
Duration: August '03 - July '07.
Amount: \$280,000.
Support: 8%SU.
Center No: 10575.1.1120449.

Title: Siebel Fellowship.
PI : Adam Wierman (my graduate student).
Duration: September '06 - September '07.

Title: External QoS Management System for Backend Database Servers.
PI : Mor Harchol-Balter.
Agency: The Technology Collaborative of Pittsburgh.
Duration: August '05 - August '06.
Amount: \$188,550.
Support: 33%SU.
Center No: 13821.1a.1010629.

Title: Analysis of Cycle Stealing and other Problems via new Dimensionality Reduction Approach.
PI : Mor Harchol-Balter.
Agency: NSF.
Grant No: CCR-0311383. Division of Computer and Communication Foundations. Theory of Computing.
Duration: July '03 - July '06.
Amount: \$150,000.
Support: 25%SU.
Center No: 10372.1.1120427.

Title: NSF PhD Graduate Fellowship for Student.
PI : Adam Wierman (my graduate student).
Agency: NSF.
Duration: July '03 - July '06.
Amount: 3 years graduate student support.

Title: QOS for Online Shopping: Providing Priority by Scheduling the Database.
PI : Mor Harchol-Balter (with Natassa Ailamaki, Co-PI)
Agency: Pittsburgh Digital Greenhouse.
Duration: August '03 - August '05.
Amount: \$265,192.
Support: 33%SU.
Center No: 10363.1.1010366.

Title: IBM Faculty Award with PhD Graduate Fellowship for Student.
PI : Bianca Schroeder (my graduate student).

Agency: IBM.
Duration: September '03 - September '05.
Amount: 2 years graduate student support.

Title: ITR: Exploiting Remote Infrastructure for Mobile Information Access.
PI : M. Satyanarayanan (Mor Harchol-Balter, Co-PI).
Agency: NSF.
Grant No: ITR ANI-0081396. Division of Computer and Network Systems.
Duration: September '00 - September '05.
Amount: \$448,820.
Support: 33%SU.
Center No: 6748.1.1120092.

Title: Running Web Servers under Overload.
PI : Mor Harchol-Balter.
Agency: Pittsburgh Digital Greenhouse.
Duration: August '01 - August '02.
Amount: \$200,500.
Support: 50-100%AY.
Center No: 6737.1.1010099.

Title: Connection Scheduling in Web Servers.
PI : Mor Harchol-Balter.
Agency: Pittsburgh Digital Greenhouse.
Duration: August '00 - August '01.
Amount: \$200,075.
Support: 67-100%SU.
Center No: 4737.1.1010099.

Title: Faculty Gift.
PI : Mor Harchol-Balter.
Agency: EMC² Corporation.
Duration: April '01 - March '04.
Amount: \$100,000.
Support: 0-22%AY.
Center No: 8236.1.1010272.

Title: IBM Faculty Award with PhD Graduate Fellowship for Student.
PI : Nikhil Bansal (my graduate student).
Agency: IBM.
Duration: September '01 - September '03.
Amount: 2 years graduate student support.

Title: Postdoctoral Fellowship Mathematical Sciences.
PI : Mor Harchol-Balter.
Agency: NSF.

Grant No: 9627445.
Duration: September '96 - August '00.
Support: Postdoctoral Research Fellowship Stipend.

INDUSTRY EXPERIENCE

- 2019 - pres Collaborating with Google on Borg Scheduling.
- 2020 - pres Collaborating with Google on Execution Economy.
- 2021 - 2022 Collaborating with Voltron on Database Scheduling.
- 2019 - 2021 Collaborating with Facebook on Caching Systems.
- 2017 - 2019 Collaborating with Microsoft on Web Caching.
- 2016 - 2018 Collaborating with Akamai on CDN Caching Algorithms.
- 2015 - 2016 Collaborating with Google on Meeting Tail QoS Guarantees for Shared Networked Storage.
- 2008 - 2016 Collaborating with Intel Research-Pittsburgh on Data Center Power Management (Mike Kozuch).
- 2012 - 2014 Collaborating with Microsoft Research-Cambridge on Prioritizing Mixed Workloads (Eno Thereska).
- 2011 - 2014 Collaborating with Microsoft Research-New England on Pricing and Queueing (Jennifer Chayes).
- 2009 - 2011 Collaborating with IBM on Cycle Stealing Analysis.
- 2009 - 2010 Collaborating with Yahoo on Building a Power-Efficient Load Balancer.
- 2009 - 2010 Collaborating with Google on Co-scheduling batch and production workloads.
- 2007 - 2009 Collaborating IBM T. J. Watson Lab on Optimal Power Management (Jeff Kephart).
- 2005 - 2006 Supported by IBM Corporation, Seagate Technologies, Spinnaker Networks, Maya Design, and Cisco Technologies through The Technology Collaborative of Pittsburgh (TTC) grant. Quarterly meetings.
- 2003 - 2005 Supported by IBM Corp. and Seagate Technologies through Pittsburgh Digital Greenhouse grant. Quarterly meetings.
- 2001 - 2003 Supported by Spinnaker Networks, Inc. through Pittsburgh Digital Greenhouse grant. Quarterly Meetings.
- 2001 - 2003 Supported by gift from EMC² Corporation. Contacts: Peter Lauterbach, Ric Wheeler.
- 2001 - 2005 My students supported by IBM graduate fellowships. Collaborate with many people at IBM on research.

- 2000 - 2001 Supported by Cisco Systems, through PDG grant. Contact: John Wakerley. Quarterly meetings.
- 1990 - 1991 Employed at NASA Ames Research Center, Moffit Field, CA.
Used Newton's Method for Fractal Approximation to generate Iterated Function System (IFS) for piecewise linear graphs. Developed rotation and size invariant filters for image recognition using simulated annealing.
- 1988 - 1990 Employed at GTE Government Systems, Needham, MA.
Mathematician for Advanced Machine Intelligence Technology group. Responsible for theory, design, and implementation of the Intelahedron™, a pattern classification system based on linear and non-linear frame-of-reference transformations. Work experience included:
1. Researching and developing eigendecomposition transformations to force separation of classes, reduce dimensionality, and identify the characteristic features of a class.
 2. Researching fast algorithms for obtaining the eigensystem of an autocorrelation matrix for a class of input signals.
 3. Applying the Intelahedron™ to: seismic data classification, underground oil detection, recognition of mixed signals with noise, decision making in a tactical battlefield.
- Co-authored with Dr. Angel DeCegama several hundred page report, *The Intelahedron™* (GTE Proprietary), detailing mathematical theory, algorithms, and implementation.
- Summer 87 Employed at AT&T Bell Laboratories, Murray Hill, NJ.
Developed object-oriented software tools in SCHEME for constraint-based propagation networks, including graphics and debugging tools.

TEACHING EXPERIENCE (Links to classes are available from my web page.)

- 15-829 **Performance Modeling Tools for Computer Systems Researchers**
NEW graduate class. This class is aimed at computer systems PhD students who are already involved in doing systems research, where the goal is to improve the performance of the system. Improving performance could involve reducing response times, providing class-based response time differentiation, improving tail behavior, scheduling to favor certain jobs, reducing loss/drop rate, increasing throughput, increasing revenue, reducing power or other costs, load balancing, etc. This class requires zero math background. The focus of the class is on learning how to translate the research being done by the CS PhD student into the appropriate queueing network framework. There is a major emphasis on guided research performance improvements.
Fall 2022 – Reviews were all 5.0 out of 5.0.
- 15-359 **Probability and Computing**
NEW Undergraduate class in probability for computer science majors. Probability theory has become indispensable in computer science. In areas such as artificial intelligence and computer science theory, probabilistic methods and ideas based on randomization are central. In other areas such as networks and systems, probability is becoming an increasingly useful framework for handling uncertainty and modeling the patterns of data that occur in complex systems. This course gives an introduction to probability as it is used in computer science theory and practice, drawing on applications and current research developments as motivation and context. Topics include combinatorial probability and random graphs, heavy tail distributions, concentration inequalities, various randomized algorithms,

sampling random variables and computer simulation, the “probabilistic method” applied to graph theory problems, and Markov chains and their many applications, from Web search engines to models of network protocols.

Fall 2004 – Co-taught with John Lafferty. My reviews averaged 4.7 out of 5.0.

Spring 2006 – Co-taught with John Lafferty. My reviews averaged 4.8 out of 5.0.

Fall 2007 – Co-taught with Ryan O’Donnell. My reviews averaged 4.9 out of 5.0.

Fall 2008 – Co-taught with Ryan O’Donnell. My reviews averaged 4.9 out of 5.0.

Spring 2012 – Co-taught with Klaus Sutner. My reviews averaged 4.6 out of 5.0.

Spring 2013 – Co-taught with Victor Adamchik. My reviews averaged 4.7 out of 5.0.

Fall 2016 – Taught alone. Reviews averaged 4.95 out of 5.0.

Spring 2015 – Taught alone. Reviews averaged 4.8 out of 5.0.

Fall 2016 – Taught alone. Reviews averaged 4.7 out of 5.0.

Spring 2018 – Co-taught with Rashmi Vinayak. My reviews averaged 5.0 out of 5.0.

15-259 **Fall 2018** – **15-359 class changes name to 15-259. Enrollment jumps to 75-100 students.**

Spring 2019 – Taught alone. Reviews averaged 4.9 out of 5.0.

Spring 2020 – Co-taught with Weina Wang. My reviews averaged 4.9 out of 5.0.

Spring 2021 – Co-taught with Weina Wang. My reviews averaged 4.8 out of 5.0.

Spring 2022 – Co-taught with Weina Wang. My reviews averaged 4.9 out of 5.0.

15-849a **Performance Analysis and Design of Computer Systems**

Highly-intensive graduate class for computer systems students. The course emphasizes techniques to help one quickly make design decisions, such as: Which device is worth speeding-up? Is it better to use one fast machine or several slow ones? How should the machines be configured? Which load balancing algorithm will maximize throughput? Which task assignment policy will minimize delays? For minimizing packet loss, should we increase the buffer size, or buy a faster device? How do various scheduling policies compare? Techniques taught include: back-of-the-envelope analysis, Markov chains, networks of queues, queueing theory basics measurement of systems, creating models from measurements, confidence intervals, simulation techniques, working with heavy-tailed distributions, as well as task assignment policies and scheduling for server farms. Student reviews write: “Take this class at all cost!”

Fall 1998 – I developed this class at MIT and passed out my hand-written lecture notes.

Reviews averaged over 6.6 out of 7.0.

Highest rating out of the 134 EECS classes taught at MIT during the past year.

Fall 1999 – First time taught at CMU. I produced 24 latexed lectures, each 15-20 pages.

Produced 7 latexed homeworks (each 10 problems) with complete solutions.

All materials are available on my web page, and are used by 2000 students and faculty worldwide. I also answer people’s questions daily regarding the notes.

Reviews averaged 5.0 out of 5.0 in every category.

Fall 2000 – Reviews averaged 4.8 out of 5.0.

Fall 2002 – Reviews averaged 4.8 out of 5.0.

Course adopted as REQUIREMENT for all Tepper OM students starting 2002.

15-857a **Fall 2006** – **15-849a class changes name to 15-857a.**

Reviews averaged 4.8 out of 5.0.

Fall 2007 – Reviews averaged 4.9 out of 5.0.

Fall 2009 – Reviews averaged 5.0 out of 5.0.

Fall 2011 – Reviews averaged 4.8 out of 5.0.

Fall 2012 – Reviews averaged 5.0 out of 5.0.

Fall 2014 – Reviews averaged 5.0 out of 5.0.

Fall 2015 – Reviews averaged 4.8 out of 5.0.

Fall 2017 – Reviews averaged 4.9 out of 5.0.

Fall 2019 – Reviews averaged 4.8 out of 5.0.

Fall 2021 – Reviews averaged 4.9 out of 5.0. Co-taught with Weina Wang.

15-849b **Theory of Performance Modeling.**

More theoretical version of 15-849a, geared towards students wishing to do research in queueing theory. The goal of queueing theory has always been to improve the design/performance of systems, e.g. networks, servers, memory, distributed systems, etc., by finding smarter schemes for allocating resources to jobs. In this class we will study the beautiful mathematical techniques used in queueing theory, including stochastic analysis, discrete-time and continuous-time Markov chains, renewal theory, BCMP product-form, Laplace transforms, supplementary random variables, fluid theory, virtual work, intensive scheduling theory, matrix-geometric techniques, and more. Throughout we emphasize realistic workloads, in particular heavy-tailed workloads. The techniques studied in this class are useful to students in Computer Science, Mathematics, ACO, OR, Statistics and Engineering. This course is packed with open problems – problems which if solved are not just interesting theoretically, but which have huge applicability to the design of computer systems today.

Fall 2001 – First introduction of this class. Created new set of 24 lectures, with about 50% overlap to 15-849a. Lecture notes are on web page and are used by researchers world-wide. Course reviews averaged 4.8 out of 5.0.

Fall 2003 – Reviews averaged 5.0 out of 5.0.

Spring 2005 – Reviews averaged 4.7 out of 5.0.

15-858a **Advanced Topics in Stochastic Processes**

Advanced topics include: Stochastic Orderings, Heavy-tailed Workloads, Rate Conservation Law, Transform-based Analysis of Queueing Theory, Scheduling Theory, Fluid Analysis, Tail Behaviors, Matrix-Analytic Methods, and more.

Spring 2007 – Co-taught with Alan Scheller-Wolf (Tepper School). Reviews: 5.0 out of 5.0.

Spring 2017 – Co-taught with Alan Scheller-Wolf (Tepper School) under 47-776.

15-441 **Computer Networking**

Course emphasizes basic performance and engineering tradeoffs in the design and implementation of computer networks. To make the issues more concrete, the class includes several multi-week projects requiring significant design and implementation. The goal is for students to learn not only what computer networks are and how they work today, but also why they are designed the way they are and how they are likely to evolve in the future. Examples are drawn from the Internet. Topics include: congestion/flow/error control, routing, addressing, naming, multi-casting, switching, internetworking, and network security.

Spring 2001 – co-taught with Peter Steenkiste. Reviews averaged 4.59 out of 5.0.

Spring 2002 – co-taught with Peter Steenkiste. Reviews averaged 4.83 out of 5.0.

Fall 2002 – co-taught with Srinu Seshan. Reviews averaged 4.8 out of 5.0.

21-127 **Concepts of Mathematics**

This course introduces the basic concepts, ideas and tools involved in doing mathematics. As such, its main focus is on presenting informal logic, and the methods of mathematical proof. These subjects are closely related to the application of mathematics in computer science. Topics discussed include a basic introduction to elementary number theory, induction, the algebra of sets, equivalence relations, congruences, recurrence equations and functions.

Spring 2014 – co-taught with Mike Picollelli. My reviews averaged 4.7 out of 5.0.

Enrollment of 360 students in our two sections, plus another 75 in an additional section.

- 15-820b **SQUALL Seminar – Scheduling and Queuing Around Lunchtime**
Graduate-level research seminar on topics in scheduling theory. Students read and present recent papers as well as their own research results.
Taught for the benefit of graduate students, so that they can get feedback on their speaking skills.
Attended by students, multiple faculty, and postdocs.
2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009
2010, 2011, 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022, 2023.

TEACHING ASSISTANT EXPERIENCE

- Fall 1993 Graduate Student Instructor, U.C. Berkeley.
Undergraduate Computer Science Algorithms
All reviews available upon request. Reviews averaged 5.0 out of 5.0.
Winner 3 teaching awards.
- Fall 1987 Teaching Assistant, Brandeis University.
Undergraduate Theory of Computation.
All reviews available upon request. Reviews averaged 5.0 out of 5.0.

STUDENT ADVISING

Postdocs:

1. **Daniel Berger** – 2017-2019. *Researcher at Microsoft Research, Seattle, WA. Also Affiliate Professor at U. Washington.*

My Ph.D. Students:

1. **Bianca Schroeder** – entered Fall 1999. Proposed May 2003. Graduated December 2005. *Professor at University of Toronto, Computer Science Department.*
2. **Takayuki Osogami** – entered Fall 2001. Proposed April 2004. Graduated August 2005. *Researcher at IBM Tokyo Research Labs.*
3. **Adam Wierman** – entered Fall 2001. Proposed Feb 2005. Graduated May 2007. Winner of 2007 CMU SCS Distinguished Dissertation Award. Winner SIGMETRICS Rising Star Award. *Professor and Department Head at Caltech, Computer Science Department.*
4. **David McWherter** – entered Fall 2001. Proposed April 2006. Graduated Aug 2008. *Researcher at Google, Seattle Lab.*
5. **Varun Gupta** – entered Fall 2004. Proposed October 2010. Graduated May 2011. *Associate professor at University of Chicago Booth School of Business.*
6. **Anshul Gandhi** – entered Fall 2007. Proposed November 2010. Graduated July 2013. Winner SPEC Dissertation Award. Winner SIGMETRICS Rising Star Award. *Associate professor at SUNY Stony Brook, Computer Science Department.*

7. **Sherwin Doroudi** – entered Fall 2010. Proposed October 2015. Graduated May 2016. Winner 2016 William W. Cooper Doctoral Dissertation Award in Management or Management Science. *Assistant professor at University of Minnesota, ISyE Department.*
8. **Kristen Gardner** – entered Fall 2012. Proposed December 2015. Graduated May 2017. *Assistant professor at Amherst College, MA, Computer Science Department.*
9. **Timothy Zhu** – entered Fall 2011. Proposed November 2015. Graduated May 2017. *Assistant professor at Penn State University, State College, PA, Computer Science Department.*
10. **Ziv Scully** – entered Fall 2016. Proposed 2021. Graduated Aug 2022. *Assistant professor at Cornell University, Ithaca, NY, IEOR Department.*
11. **Ben Berg** – entered Fall 2016. Proposed 2021. Graduated Aug 2022. *Assistant professor at UNC-Chapel Hill, Computer Science Department.*
12. **Isaac Grosf** – entered Fall 2017. Proposed 2022. Graduated Aug 2023. *Assistant professor at Northwestern University, ISyE Department.*

My MS Students:

1. Yi Zhuang – entered Fall 2008. Graduated with MS December 2010.

Ph.D. Thesis Committee Member:

1. Chris Lumb – Graduated 2004 – ECE.
2. Atul Bhandari – Graduated 2006 – Tepper School (Operations Management).
3. Eno Thereska – Graduated 2007 – ECE.
4. Carol Frieze – Graduated 2007 – CS.
5. Ying Xu – Graduated 2015 – Tepper School (Operations Management).
6. Alexey Tumanov – Graduated 2016 – ECE.
7. Leela Nageswaran – Graduated 2018 – Tepper School (Operations Management).
8. Daniel Berger – Graduated 2018 – CS, University of Kaiserslautern, Germany.
9. Stephen Mueller – Graduated 2018 – CS.
10. Igor Kadota – Graduated 2020 – CS, MIT.
11. Naama Ben-David – Graduated 2020 – CS, CMU.

Undergraduate Researchers Supervised:

1. Vanshika Chowdhary – 2022.
2. Naveen Pai – 2020-2021.
3. Sean Zhang – 2020-2021.
4. Kunhe Yang – 2019-2021.

5. Jacky (Sheng) Zhou – 2018-2019.
6. Corwin de Boor – 2017-2018.
7. Tanay Vakharia – 2017.
8. Haotian Jiang – 2017.
9. Sam Zbarsky – 2014-2015.
10. Adrian Trejo Nunez – 2011-2012.
11. Kevin Zhao – 2010-2011.
12. Victor Marmol – 2010-2011.
13. Ram Raghunathan – 2010-2011. Thesis won Alumni Award for Undergraduate Excellence.
14. Justin Hurley – 2008-2009.
15. Michelle Burroughs – 2008-2009.
16. Gwendolyn Stockman – 2005-2006.
17. Alex Fishman – 2005-2006.
18. Sonesh Surana – 1999-2000.