

Siddhartha Sen

<http://sidsen.org>

530 W 113th St., Apt. 7B
New York, NY 10025

sidsen@microsoft.com
(617) 834-4885

Education

- Princeton University** Princeton, NJ
Ph.D. in Computer Science June 2013
Thesis title: *New Systems and Algorithms for Scalable Fault Tolerance*
M.A. in Computer Science September 2009
Advisors: Michael J. Freedman, Robert E. Tarjan
GPA: 4.0/4.0
- Massachusetts Institute of Technology** Cambridge, MA
M.Eng. in Computer Science and Engineering June 2004
Thesis title: *Dynamic Processor Allocation for Adaptively Parallel Jobs*
Advisor: Charles E. Leiserson
GPA: 5.0/5.0
- B.S. in Computer Science and Engineering and B.S. in General Mathematics,
Minor in Spanish June 2003
GPA: 4.8/5.0
- International School Manila High School** Manila, Philippines
Graduated valedictorian and class speaker June 1999
GPA: 4.0/4.0

Research Interests

Distributed systems, machine learning systems, data structures and algorithms

Honors

- Window & Devices Group Universal Storefronts Innovation Award (joint with MSN product group, 2016)
- Best Student Paper award, PODC 2012
- First Google Ph.D. Fellowship in Fault-Tolerant Computing (2009-2012)
- Award paper, LADIS 2011
- Invited papers to Journal of the ACM (2012), Operating Systems Review (2012, 2011), Theoretical Computer Science (2009), Algorithmica (2009)
- Princeton University Graduate Fellowship, 2007
- Member of Phi Beta Kappa, Tau Beta Pi, and Eta Kappa Nu, National Society of Collegiate Scholars (NSCS), and Sigma Xi national honor societies
- International Baccalaureate Diploma Programme total 44/45

Research Experience

- 2013– **Microsoft Research** (New York City, Silicon Valley)
Researcher
Working with machine learning researchers on a service for optimizing online decisions; the service has been released (<http://www.aka.ms/mwt>) and open-sourced, and has been deployed in internal product groups (e.g., MSN) as well as external customers. Working with computational social science and prediction markets researchers to build a platform for real-time social media analysis (e.g., of Twitter data). Working with colleagues at Princeton, Columbia, NYU, and Yale on projects related to caching, scalable data structures in multicore machines, data protection in machine learning systems, distributed storage in ultra-low-latency networks, wide-area video streaming and analytics, and other topics.
- 2012–2013 **Princeton University** (Network Systems Group), **New York University** (Networking and Wide-Area Systems Group)
Junior Research Scientist, Postdoctoral Researcher
Primary project is designing and building JetStream [2, 4], a system for analyzing global-scale data streams (big data) in real-time. JetStream combines rich, online analytical processing (OLAP) databases with stream-processing techniques to enable bandwidth-aware data adaptation. Secondary project is designing distributed storage indices with better load balance in fast (Infiniband) datacenter networks, using tree structures from my research [19], skip lists, and multiple-choice hashing.
- 2007–2012 **Princeton University** (Network Systems Group, Theory Group)
Graduate Research Assistant, Google Fellow
Designed, built, and analyzed several distributed systems that are provably fast, reliable, and scalable. Research resulted in new systems [7, 10, 5] and new theoretical results proving their efficiency and robustness. Notable projects include new data structures and analyses [19, 11] for a database indexing technique that failed in red-black trees (inciting a lawsuit) but succeeds in B-trees; Prophecy [10], a system that makes Byzantine fault tolerance (BFT) practical for Internet cloud services, and new techniques for increasing the fault-resilience of BFT [6] and scaling out BFT to global-scale peer-to-peer systems [20]; GreenUp [7], a system for saving energy in the enterprise while keeping machines highly available, and a new balls-in-bins analysis of its load balancing scheme; and LocalFlow [5], an optimal, local flow routing scheme for datacenter networks based on software-defined networking (SDN).
- 6/10–8/10 **Microsoft Research Redmond** (Distributed Systems Group)
Research Intern
Worked on two research projects related to high availability: 1) GreenUp, a provably-efficient, decentralized system for keeping machines highly available even as they sleep [7]; and 2) A characterization of Byzantine fault tolerance using partial broadcast channels and hypergraph coloring, leading to improved and tunable fault resilience [6].

- 6/08–8/08 **Microsoft Research SVC** (Distributed Computing Group)
Research Intern
 Worked on two research projects related to routing: 1) Preprocessing heuristics for shortest-path algorithms used by online maps; and 2) Provably-efficient virtual ring routing schemes for overlay networks [12]. Made extensive use of the Dryad cluster and DryadLINQ language to parallelize a graph preprocessing algorithm, and discovered several bugs in both systems.
- 2003–2004 **MIT CSAIL** (Supercomputing Technologies Group)
Research Assistant
 Designed and implemented a dynamic processor-allocation system for adaptively parallel jobs, focusing on jobs scheduled with the randomized work-stealing algorithm [30]. Extended the Cilk multithreaded language’s runtime environment to include this processor allocation system.
- 6/02–9/02 **IBM T. J. Watson Research Center** (Scalable Systems Group)
Co-op Pre-Professional Programmer / Research Assistant
 Researched and analyzed various technologies and architectures of grid computing, focusing on the areas of resource metering, logging, and problem determination in heterogeneous grid environments. Designed a Resource Metering System (RMS) for grid services based on the Application Response Measurement and Common Information Model standards.

Industry Experience

- 9/07–12/07 **Dial Directions**
 (part-time) *Software Engineer*
 Architected and developed batch recognition systems for tuning grammar generation and voice recognition in production servers.
- 2004–2007 **Microsoft Corporation** (Clustering & High Availability Group)
Software Design Engineer
 Designed, developed, and tested network load balancing software for the next release of Windows Server (codename “Longhorn”). Drove the generation of ideas and intellectual property [34–40] for Microsoft’s high-availability and scale-out technologies, collaborating with Microsoft Research when appropriate.
- 7/03–8/03, **Microsoft Corporation** (Windows Network Load Balancing Group)
 9/03–1/04 *Software Design Engineer Intern*
 (part-time) Worked on research projects related to load balancing and distributed problem diagnosis. Designed and implemented a Distributed Assertions Framework (DAF) based on abstract state machines for programmatically verifying the execution of distributed objects [37]. Designed and implemented a prototype for the Distributed Atom Manager (DAM), a distributed state-maintenance system for next-generation Windows load balancing software.
- 6/01–8/01 **Microsoft Corporation** (PocketPC Group)
Software Design Engineer Intern
 Developed an HTTP mail transport for the PocketPC Inbox application to comple-

ment the already existing IMAP and POP3 transports. The transport performed web-based DAV authentication and provided full mail functionality, including a “connected” mode that processes transactions instantaneously.

Teaching Experience

- 6/16–8/16 **Microsoft Research NYC** (Data Science Summer School)
Instructor
Lectured on computational tractability, data structures, and algorithms. Mentored a final project on analyzing traffic flow and carpooling potential using public NYC taxi ride data.
- 2008–2009 **Princeton University** (COS 423—Theory of Algorithms)
Guest Lecturer
Lectured on P vs. NP, NP-completeness, and using approximation algorithms to cope with NP-completeness.
- 9/08–1/09 **Princeton University** (COS 423—Theory of Algorithms)
Teaching Assistant
Independently designed and taught a semiweekly, hour-long recitation lecture; course material included results from my research [13, 14]. Graded assignments, proposed and supervised final projects, and helped students with proofs through on-site office hours and email.
- 2/09–5/09 **Princeton University** (COS 126—General Computer Science)
Teaching Assistant
Taught a semiweekly, hour-long recitation lecture, prepared and graded assignments and exams, and helped students with coding/design problems through on-site office hours and email.
- 2003–2004 **MIT Chemical Engineering Department** (I-Lab Heat Exchanger Project)
Research Assistant
The I-Lab is an interactive heat exchanger experiment run remotely by students taking courses in control and heat/mass transfer. Responsibilities included developing interfaces to new experiments using LabVIEW, designing and maintaining the heat exchanger website, and coordinating the use of the experiment by other universities. Work was published in an education journal [23].
- 2002–2003 **MIT CSAIL** (6.170—Laboratory in Software Engineering)
Laboratory Assistant
Assisted professors and teaching assistants in the required software engineering lab for computer science majors. Responsibilities included previewing problem sets and assignments, helping students with coding/design problems, and performing on-site duty in computer clusters (minimum of 5 hours/week).
- 6/00–7/00 **MIT Minority Introduction to Engineering and Science (MITES)**
Resident Advisor in Physics
Supervised rising high school seniors in this highly-selective, fully-funded summer program. Graded assignments and exams and held office hours for physics.

Student Advising

- Graduate Amy Tai (Princeton), Riley Spahn (Columbia), Mathias Lecuyer (Columbia), Aaron Blankstein (Princeton), Matvey Arye (Princeton), Christopher Mitchell (NYU), Daniel Larkin (Princeton)
- Undergrad Kay Ousterhout (Princeton, now at UC Berkeley), Jarett Schwartz (Princeton, now at UC Berkeley), Hong Kyun Kim (Princeton, now at U. Chicago)

Professional Activities

- Program comm. CoNEXT '16, NSDI '16, SOCC '15, P2P '15, SSS'13, SIROCCO '13
- External reviews ICALP '15, ESA '14, FOCS '13, NSDI '13, SODA '13, NSDI '12, CATS '12, SODA '12, SODA '11, ICALP '10, LATIN '10, STACS '10
- Journal reviews Theoretical Computer Science (TCS), Transactions on Knowledge and Data Engineering (TKDE), Software: Practice and Experience (SPE), Information Processing Letters (IPL)

Service and Leadership

- 2000–2008 Salsa and Hip-Hop dance instructor
- 2001–2003 President of MIT Casino Rueda (Salsa) Group
- 2000–2001 President/Treasurer of MIT International Students Association
- 2001–2002 Rush Chair/Alumni Secretary of Nu Delta Fraternity

Refereed Conference and Workshop Publications

(Author list of theory papers is alphabetical in this and subsequent sections)

- [1] Christopher Mitchell, Kate Montgomery, Lamont Nelson, **Siddhartha Sen**, and Jinyang Li. Balancing cpu and network in the cell distributed b-tree store. In *Proc. 2016 USENIX Annual Technical Conference (ATC)*, 2016. 14 pages.
- [2] Ariel Rabkin, Matvey Arye, **Siddhartha Sen**, Vivek S. Pai, and Michael J. Freedman. Aggregation and degradation in jetstream: Streaming analytics in the wide area. In *Proc. 11th Symposium on Networked Systems Design and Implementation (NSDI)*, 2014. 14 pages.
- [3] Daniel Larkin, **Siddhartha Sen**, and Robert E. Tarjan. A back-to-basics empirical study of priority queues. In *Proc. 16th Workshop on Algorithm Engineering and Experiments (ALENEX)*, 2014. 12 pages.
- [4] Ariel Rabkin, Matvey Arye, **Siddhartha Sen**, Vivek Pai, and Michael J. Freedman. Making every bit count in wide-area analytics. In *Proc. 14th Conference on Hot Topics in Operating Systems (HotOS)*, 2013. 6 pages.
- [5] **Siddhartha Sen**, David Shue, Sunghwan Ihm, and Michael J. Freedman. Scalable, optimal flow routing in datacenters via local link balancing. In *Proc. 9th Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2013. 12 pages.

- [6] Alexander Jaffe, Thomas Moscibroda, and **Siddhartha Sen**. On the price of equivocation in Byzantine agreement. In *Proc. 31st Principles of Distributed Computing (PODC)*, 2012. 10 pages.
- [7] **Siddhartha Sen**, Jacob R. Lorch, Richard Hughes, Carlos Garcia Jurado Suarez, Brian Zill, Weverton Cordeiro, and Jitendra Padhye. Don't lose sleep over availability: The GreenUp decentralized wakeup service. In *Proc. 9th Networked Systems Design and Implementation (NSDI)*, 2012. 14 pages.
- [8] **Siddhartha Sen**, Sunghwan Ihm, Kay Ousterhout, and Michael J. Freedman. Brief announcement: Bridging the theory-practice gap in multi-commodity flow routing. In *Proc. 25th International Symposium on Distributed Computing (DISC)*, 2011. 2 pages.
- [9] **Siddhartha Sen** and Michael J. Freedman. Commensal cuckoo: Secure group partitioning for large-scale services. In *Proc. 5th Workshop on Large Scale Distributed Systems and Middleware (LADIS)*, 2011. 6 pages.
- [10] **Siddhartha Sen**, Wyatt Lloyd, and Michael J. Freedman. Prophecy: Using history for high-throughput fault tolerance. In *Proc. 7th Networked Systems Design and Implementation (NSDI)*, 2010. 16 pages.
- [11] **Siddhartha Sen** and Robert E. Tarjan. Deletion without rebalancing in balanced binary trees. In *Proc. 21st Symposium on Discrete Algorithms (SODA)*, 2010. 10 pages.
- [12] Dahlia Malkhi, **Siddhartha Sen**, Kunal Talwar, Renato Fonseca F. Werneck, and Udi Wieder. Virtual ring routing trends. In *Proc. 23rd International Symposium on Distributed Computing (DISC)*, 2009. 15 pages.
- [13] Bernhard Haeupler, **Siddhartha Sen**, and Robert E. Tarjan. Rank-pairing heaps. In *Proc. 17th European Symposium on Algorithms (ESA)*, 2009. 12 pages.
- [14] Bernhard Haeupler, **Siddhartha Sen**, and Robert E. Tarjan. Rank-balanced trees. In *Proc. 11th International Symposium on Algorithms and Data Structures (WADS)*, 2009. 12 pages.
- [15] **Siddhartha Sen** and Robert E. Tarjan. Deletion without rebalancing in multiway search trees. In *Proc. 20th International Symposium on Algorithms and Computation (ISAAC)*, 2009. 10 pages.
- [16] Bernhard Haeupler, Telikepalli Kavitha, Rogers Mathew, **Siddhartha Sen**, and Robert E. Tarjan. Faster algorithms for incremental topological ordering. In *35th International Colloquium on Automata, Languages and Programming (ICALP)*, 2008. 13 pages.

Refereed Journal Publications

- [17] **Siddhartha Sen**, Robert E. Tarjan, and David Hong Kyun Kim. Deletion without rebalancing in binary search trees. *ACM Transactions on Algorithms*, 2016. To appear.
- [18] Bernhard Haeupler, **Siddhartha Sen**, and Robert E. Tarjan. Rank-balanced trees. *ACM Transactions on Algorithms*, 11(4), 2015. 26 pages.

- [19] **Siddhartha Sen** and Robert E. Tarjan. Deletion without rebalancing in multiway search trees. *ACM Transactions on Database Systems*, 39(1), 2014. 14 pages.
- [20] **Siddhartha Sen** and Michael J. Freedman. Commensal cuckoo: Secure group partitioning for large-scale services. *Operating Systems Review*, 46(1), 2012. 7 pages.
- [21] Bernhard Haeupler, Telikepalli Kavitha, Rogers Mathew, **Siddhartha Sen**, and Robert E. Tarjan. Incremental cycle detection, topological ordering, and strong component maintenance. *ACM Transactions on Algorithms*, 8(1), 2012. 33 pages.
- [22] Bernhard Haeupler, **Siddhartha Sen**, and Robert E. Tarjan. Rank-pairing heaps. *SIAM Journal on Computing*, 40(6), 2011. 23 pages.
- [23] Anders Selmer, Mike Goodson, Markus Kraft, **Siddhartha Sen**, V. Faye McNeill, Barry S. Johnston, and Clark Colton. Process dynamics and control experiment performed across the atlantic. *Chemical Engineering Education*, 2005. 6 pages.

In Submission

- [24] Aaron Blankstein, **Siddhartha Sen**, and Michael J. Freedman. Hyperbolic caching: Flexible caching for web applications, 2016.
- [25] Mathias Lecuyer, Riley Spahn, Roxana Geambasu, Tzu-Kuo Huang, and **Siddhartha Sen**. Pyramid: Enhancing rigor and selectivity in big data protection with count-based featurization, 2016.
- [26] Matvey Arye, **Siddhartha Sen**, and Michael J. Freedman. Poor video streaming performance explained (and fixed), 2016.
- [27] Irina Calciu, **Siddhartha Sen**, Mahesh Balakrishnan, and Marcos K. Aguilera. Black-box concurrent data structures for NUMA architectures, 2016.
- [28] Alekh Agarwal, Sarah Bird, Markus Cozowicz, Luong Hoang, John Langford, Stephen Lee, Jiaji Li, Dan Melamed, Gal Oshri, Oswaldo Ribas, **Siddhartha Sen**, and Alex Slivkins. A multiworld testing decision service. <https://arxiv.org/abs/1606.03966>, 2016.

Theses

- [29] **Siddhartha Sen**. *New Systems and Algorithms for Scalable Fault Tolerance*. PhD thesis, Princeton University, June 2013.
- [30] **Siddhartha Sen**. Dynamic processor allocation for adaptively parallel jobs. Master's thesis, Massachusetts Institute of Technology, September 2004.

Invited Publications

- [31] **Siddhartha Sen**. Review of PODC 2012. *SIGACT News*, 43(4), 2012. 8 pages.
- [32] Matvey Arye, **Siddhartha Sen**, Ariel Rabkin, and Michael J. Freedman. Towards efficient stream processing in the wide area. In *Proc. 6th Workshop on Large Scale Distributed Systems and Middleware (LADIS)*, 2011. 2 pages.

Patents

- [33] Alex Slivkins, Alekh Agarwal, John Langford, Sarah Bird, **Siddhartha Sen**, Lihong Li, Miro Dudik, and Leon Bottou. Application testing. *United States Patent Application*. Filed Oct. 13, 2014.
- [34] **Siddhartha Sen**, Vladimir Petter, and Andrea D’Amato. Client affinity in distributed load balancing systems. *United States Patent* #8176495. Issued May 8, 2012.
- [35] Vladimir Petter, **Siddhartha Sen**, Andrea D’Amato, and Jimmy C. Yu. Maintaining client affinity in network load balancing systems. *United States Patent* #8046467. Issued Oct. 25, 2011.
- [36] **Siddhartha Sen** and Amit Date. Software testing techniques for stack-based environments. *United States Patent* #7827531. Issued Nov. 2, 2010.
- [37] **Siddhartha Sen**, Joseph M. Joy, Nikolai Tillman, Colin L. Campbell, and Margus Veanes. Diagnosing problems in distributed systems. *United States Patent* #7548911. Issued June 16, 2009.
- [38] Jacob R. Lorch, **Siddhartha Sen**, Jitendra D. Padhye, Richard L. Hughes, and Carlos G. J. Suarez. Decentralized sleep management. *United States Patent Application* #12/940,806. Filed Nov. 5, 2010.
- [39] **Siddhartha Sen**. Dynamic network load balancing using roundtrip heuristic. *United States Patent Application* #11/518,673. Filed Sept. 11, 2006.
- [40] **Siddhartha Sen**. Load balancing via rotation of cluster identity. *United States Patent Application* #11/276,761. Filed Mar. 13, 2006.

Refereed Conference Talks

- [41] Scalable, optimal flow routing in datacenters via local link balancing. In *9th Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, December 2013.
- [42] On the price of equivocation in Byzantine agreement. In *31st Principles of Distributed Computing (PODC)*, July 2012.
- [43] Don’t lose sleep over availability: The GreenUp decentralized wakeup service. In *9th Networked Systems Design and Implementation (NSDI)*, April 2012.
- [44] Security in peer-to-peer systems revisited (work in progress). In *23rd Symposium on Operating Systems Principles (SOSP)*, October 2011.
- [45] Bridging the theory-practice gap in multi-commodity flow routing (brief announcement). In *25th International Symposium on Distributed Computing (DISC)*, September 2011.
- [46] Commensal cuckoo: Secure group partitioning for large-scale services. In *5th Workshop on Large Scale Distributed Systems and Middleware (LADIS)*, September 2011.
- [47] Prophecy: Using history for high-throughput fault tolerance. In *7th Networked Systems Design and Implementation (NSDI)*, April 2010.

- [48] Deletion without rebalancing in balanced binary trees. In *21st Symposium on Discrete Algorithms (SODA)*, January 2010.
- [49] Prophecy: Using history for high-throughput fault tolerance (work in progress). In *22nd Symposium on Operating Systems Principles (SOSP)*, October 2009.
- [50] Rank-pairing heaps. In *17th European Symposium on Algorithms (ESA)*, September 2009.
- [51] Rank-balanced trees. In *11th International Symposium on Algorithms and Data Structures (WADS)*, August 2009.
- [52] Deletion without rebalancing in multiway search trees. In *20th International Symposium on Algorithms and Computation (ISAAC)*, December 2009.
- [53] Faster algorithms for incremental topological ordering. In *35th International Colloquium on Automata, Languages and Programming (ICALP)*, July 2008.

Other/Invited Talks

- [54] Coalescing balls and bins. Bertinoro Workshop on Algorithms and Data Structures, June 2013.
- [55] Localflow: Simple, local flow routing in datacenters. DIMACS Workshop on Systems and Networking Advances in Cloud Computing, December 2011.
- [56] The price of equivocation: Characterizing Byzantine agreement via hypergraph coloring. Bertinoro Workshop on Algorithms and Data Structures, June 2011.
- [57] New balanced search trees. Microsoft Research Redmond Colloquium, August 2010.
- [58] New balanced search trees. China Theory Week, September 2010.
- [59] Simpler, better balanced search trees. University of Vienna, May 2010.
- [60] Simpler, better balanced search trees. NYU Polytechnic, April 2010.
- [61] Balanced search trees simplified. Dagstuhl Seminar on Data Structures, March 2010.
- [62] Simpler, better balanced search trees. Princeton University, February 2010.

References

Prof. Jennifer Chayes

Distinguished Scientist and Managing Director, Microsoft Research
 jchayes@microsoft.com

Prof. Michael J. Freedman

Assistant Professor, Princeton University
 mfreed@cs.princeton.edu

Prof. Jinyang Li

Associate Professor, New York University
jinyang@cs.nyu.edu

Dr. Jacob R. Lorch

Researcher, Microsoft Research Redmond
lorch@microsoft.com

Dr. Thomas Moscibroda

Senior Researcher, Microsoft Research Asia
Chair Professor, Tsinghua University
moscitho@microsoft.com

Dr. Jitendra Padhye

Principal Researcher, Microsoft Research Redmond
padhye@microsoft.com

Prof. Robert E. Tarjan

James S. McDonnell Distinguished University Professor, Princeton University
ret@cs.princeton.edu