Michael A. Bender

Professor
Department of Computer Science
Stony Brook University
Stony Brook, NY 11794-2424 USA

Phone: +1 631-632-7835
Fax: +1 631-632-8334
bender@cs.stonybrook.edu
http://www.cs.stonybrook.edu/∼bender

Personal Data

Married, two children.
U.S. citizen.
Foreign languages: Fluent in French, Spanish.

Employment

Professor. Dept of Computer Science, Stony Brook University, 2015–present.
Associate Professor. Dept of Computer Science, Stony Brook University, 2004–2015.
Founder and Chief Scientist. Tokutek, 2012–2015.
Founder and Chief Technical Officer. Tokutek, 2010–1012.
Founder and VP of Engineering. Tokutek, 2006–2010.
Visiting Research Fellow. Dept of Computer Science, Kings College London, 2004–2005.
Assistant Professor. Dept of Computer Science, Stony Brook University, 1998–2004.
Visiting Scientist. Computer Science and AI Laboratory, MIT, 2003 – 2004.
Research Associate. Bell Laboratories, Lucent Technologies, Murray Hill, summer 1996.

Education

Harvard University. PhD 1998 and SM 1995 in Computer Science.
Advisor: Michael O. Rabin. Thesis: *New Algorithms and Metrics for Scheduling.*

Ecole Normale Supérieure de Lyon. Diplôme d'Etudes Approfondies (DEA) d'Informatique Fondamentale, Magistère d'Informatique et de Modelisation. *Mention bien*, 1993.

Harvard University. AB *Magna Cum Laude* with Highest Honors in Applied Mathematics, 1992.

Honors

1. David R. Smith Leading Scholar in Computer Science, 2016-2020.
2. FAST Best Paper Award, 2016.
3. Chancellor's Award for Excellence in Teaching, 2015.
4. IPDPS Best Paper Award, 2015.
5. Major Contributions to Graduate Education and Research, CS Dept, Stony Brook, 2012.
6. Imre Simon Test-of-Time Award, 2012.
7. Undergraduate Teaching Award, CS Dept, Stony Brook, 2006.
8. R&D 100 Award, 2006.

9. PODS Best Newcomer Award, 2006.
10. Dean's Award for Excellence in Graduate Teaching, Stony Brook, 2005.
11. Graduate Teaching Award, Dept of Computer Science, Stony Brook, 2000.
12. Rotary Fellowship at the Ecole Normale Supérieure de Lyon, France, 1992–1993.

Program Committees

1. 10th Annual Fall Workshop on Computational Geometry 2000.
2. Genetic and Evolutionary Computation Conference (GECCO) 2001.
3. 1st International Workshop on Efficient Algorithms (WEA) 2001.
4. Latin American Theoretical Informatics (LATIN) 2002.
5. Genetic and Evolutionary Computation Conference (GECCO) 2002.
6. 5th Workshop on Algorithm Engineering and Experiments (ALENEX) 2003.
7. 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2003,
8. 10th International Conference on High Performance Computing (HiPC) 2003.
9. Latin American Theoretical Informatics (LATIN) 2004.
10. 3rd International Conference on FUN with Algorithms (FUN) 2004.
11. 11th International Conference on High Performance Computing (HiPC) 2004.
12. Genetic and Evolutionary Computation Conference (GECCO) 2004.
13. 15th Annual International Symposium on Algorithms and Computation (ISAAC) 2004.
14. 10th Annual International Computing and Combinatorics Conference (COCOON) 2004.
15. 13th Annual European Symposium on Algorithms (ESA) 2005.
16. 2nd Multidisciplinary International Conference on Scheduling (MISTA) 2005.
17. 12th International Conference on High Performance Computing (HiPC) 2005 (**Vice-Chair, Algorithms Track**).
18. 11th European Conference on Parallel Processing (Euro-Par) 2005 (**Vice-Chair, Scheduling and Load Balancing**).
19. 20th IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2006.
20. 12th International Conference on Parallel and Distributed Systems (ICPADS), 2006.
21. 18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) 2006.
22. 13th String Processing and Information Retrieval (SPIRE) 2006.
23. 21st IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2007.
24. Workshop on Programming Models for Grid Computing (PMGC) 2007.
25. Latin American Theoretical Informatics (LATIN) 2008.
26. 9th Workshop on Algorithm Engineering and Experiments (ALENEX) 2008.
27. 35th International Colloquium on Automata, Languages, and Programming (ICALP) 2008.
28. 4th International Workshop on Algorithmic Aspects of Wireless Sensor Networks (ALGOSENSORS) 2008.
29. 21st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) 2009. (**PC Chair**)
30. 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2010.
31. 29th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC) 2010.
32. 17th International Conference on High Performance Computing (HiPC) 2010.
33. 31st International Conference on Distributed Computing Systems (ICDCS) 2011 (**Vice-Chair, Algorithms Track**).

34. 19th Annual European Symposium on Algorithms, Engineering and Applications Track (ESA) 2011.
35. 10th Latin American Theoretical Informatics (LATIN) 2012.
36. 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2013.
37. 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) 2013.
38. ASE/IEEE International Conference on Big Data, 2013.
39. 11th Latin American Theoretical Informatics (LATIN) 2014.
40. 21st International Conference on High Performance Computing (HiPC) 2014.
41. Sixth Workshop on Massive Data Algorithmics (MASSIVE) 2014.
42. 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2015.
43. 12th Latin American Theoretical Informatics (LATIN) 2016.
44. 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) 2016.
45. Eighth International Conference on Fun with Algorithms (FUN) 2016.
46. 36th ACM SIGACT-SIGOPS Symposium on the Principles of Distributed Computing (PODC) 2017.
47. 46th International Conference on Parallel Processing (ICPP) 2017. **Algorithms Track Co-Chair**.
48. 13th Latin American Symposium on Theoretical INformatics (LATIN) 2018. **PC Chair**.
49. 27th Annual European Symposium on Algorithms (Track B) 2019. **PC Chair.**
50. 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS) 2019.

Selected Conference Organization

1. Publicity Chair. Symposium on Parallelism in Algorithms and Architectures (SPAA), 2000–2007.
2. Co-organizer. Dagstuhl Seminar 04301: Cache-Oblivious and Cache-Aware Algorithms, 2004.
3. Co-organizer. CIRM Center at Marseille-Luminy: Scheduling Algorithms for New Emerging Applications, 2006.
4. Co-organizer. CIRM Center at Marseille-Luminy: New Challenges in Scheduling Theory, 2008.
5. Program Committee Chair. 21st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2009.
6. Co-organizer. Centre CNRS "La Villa Clythia," Frejus, France. New Challenges in Scheduling Theory, 2010.
7. Co-organizer. 44th ACM Symposium on Theory of Computing (STOC), Tutorial on Algorithms for Memory-Sensitive Computing, 2012.
8. Program Committee Chair (joint with M. Mosteiro). 13th Latin American Symposium on Theoretical INformatics (LATIN) 2018.

Editorial Boards

1. Guest Editor. *Journal of Algorithms* Special Issue on SODA 2003.
2. Editor. *Journal of Discrete Algorithms,* 2004–2014.
3. Guest Editor. *Journal of Scheduling* Special Issue on Scheduling Algorithms for New Emerging Applications, 2006-2008.

4. Guest Editor. *Journal of Scheduling* Special Issue on New Challenges in Scheduling Theory, 2008–2014.
5. Guest Editor. *Theory of Computing Systems* Special Issue on SPAA 2009.

Selected Patents

1. United States Patent 6,112,221. A System and Method for Scheduling Webservers with a Quality-of-Service Guarantee for Each User. Michael A. Bender, Soumen Chakrabarti, and Shanmugavelayutham Muthukrishnan. Awarded: August 29, 2000.

2. United States Patent 7,565,657. Allocating Application to Group of Consecutive Processors in Fault-Tolerant Deadlock-Free Routing Path Defined by Routers Obeying Same Rules for Path Selection. Vitus Leung, Michael A. Bender, David Bunde, and Cynthia Phillips. Awarded: July 21, 2009.

3. United States Patent 8,185,551. Disk-Resident Streaming Dictionary. Bradley C. Kuszmaul, Michael A. Bender, and Martin Farach-Colton. Awarded: May 21, 2012.

4. United States Patent 8,489,638. Disk-Resident Streaming Dictionary. Bradley C. Kuszmaul, Michael A. Bender, and Martin Farach-Colton. July 16, 2013.

5. United States Patent 8,996,563. High-performance streaming dictionary. Michael A. Bender, Martin Farach-Colton, Yonatan R. Fogel, Zardosht Kasheff, Bradley C. Kuszmaul, Vincenzo Liberatore, Barry Perlman, Rich Prohaska, and David S. Wells. Awarded: March 31, 2015.

Grants

1. co-PI: Algorithms in Support of Scalable Tactical Imagery eXploitation. Subcontract from ISX Corporation on a grant from DARPA. $40,000 (with E. Arkin, AMS, and J. Mitchell, AMS). Award effective: 8/1/1999–12/25/2000.

2. co-PI: Algorithms in Support of Pheromone Robotics. Subcontract from HRL Laboratories on a grant from DARPA. $171,062 (with E. Arkin, AMS, and J. Mitchell, AMS). Award effective: 9/1/1999–9/1/2002.

3. co-PI: EIA 0112849 ITR/SY(CISE): Cache-Oblivious Data Structures. National Science Foundation (with L. Arge, Duke, and E. Demaine, MIT). $449,571 ($161,529 SBU portion). Award effective: 9/1/2001–8/31/2004.

4. PI: Algorithmic Support for Cplant Scheduling. Sandia National Laboratories. $46,132 (with E. Arkin, AMS). Award effective: 6/11/2001–9/28/2001.

5. co-PI: CCF 0103059 Nanoscale Single Electron Switching Arrays for Self-Evolving Neuromorphic Networks. National Science Foundation. $599,980 (with K. Likarev, Physics, J. Lukens, Physics, and A. Mayr, Physics). Award effective: 7/1/2001–6/30/2003.

6. PI: CCF 0208670 Data Structures and Algorithms for Maintaining Data Locality. National Science Foundation. $149,937. Award effective: 7/15/2002–6/30/2005.

7. PI: Algorithmic Support for Cplant Scheduling. Sandia National Laboratories. $20,351. Award effective: 6/11/2002–9/20/2002.

8. PI: Algorithmic Support for Cplant Scheduling. Sandia National Laboratories. $34,680. Award effective: 5/1/2003–1/1/2003.

9. co-PI: ACI 0324974 ITR: Transactions Everywhere. National Science Foundation (with B. Kuszmaul, MIT, and C. Leiserson, MIT). $650,000 ($100,000 SBU portion). Award effective: 9/2003–8/2004.

10. PI: CCF 0540897 Collaborative Research: High-Performance Data Access through Memory Abstraction. National Science Foundation (with M. Farach-Colton, Rutgers). $399,033 ($199,836 SBU portion). Award effective: 8/1/2006–7/31/2009.

11. co-PI: CNS 0627645 Authenticating Reality. National Science Foundation (with R. Johnson and D. Samaras, CS Dept). $350,000. Award effective: 10/1/2006–9/30/2009.

12. PI: CCF 0621439 HECURA: Collaborative Research: Techniques for Streaming File Systems and Databases. National Science Foundation (with M. Farach-Colton, Rutgers). $341,400, ($141,378 SBU portion). Award effective: 8/1/2006–1/31/2008.

13. PI: Techniques for Massive-Data Collection in Wireless and Sensor Networks. Center of Excellence in Wireless and Information Technology (CEWIT). $10,000. Award effective: 1/2007–12/2007.

14. Participating Scientist: NSF IIP-0740298 SBIR Phase I: A Storage Engine for High-Volume Data (M. Farach-Colton, PI and B. Kuszmaul, Participating Scientist). National Science Foundation. $100,000. Award effective: 1/1/2008–12/31/2008.

15. co-PI: Algorithms in Support of Fractal Tree Databases. Tokutek, Inc./Sensor CAT (with J. Mitchell, AMS). $33,512.50. Award effective: 9/1/2009–2/1/2010.

16. co-PI: Continuation of Algorithms in Support of Fractal Tree Databases. Tokutek, Inc./Sensor CAT (with J. Mitchell, AMS). $57,962.50. Award effective: 1/1/2010–7/1/2010.

17. Participating Scientist: NSF IIP-0945687 SBIR Phase I: A Multithreaded Storage Engine using Highly-Concurrent Fractal Trees (with B. Kuszmaul, PI, and M. Farach-Colton, Participating Scientist). National Science Foundation. $150,000. Award effective: 1/1/10–12/31/10.

18. Participating Scientist: NSF IIP-1058565 SBIR Phase II: A Multithreaded Storage Engine Using Highly-Concurrent Fractal Trees (with B. Kuszmaul, PI, and M. Farach-Colton, Participating Scientist). National Science Foundation. $500,000. Award effective: 2/1/2011-1/31/2013.

19. PI: CCF 0634793 Collaborative Research: Adversarial Contention Resolution. National Science Foundation (with M. Farach-Colton). $299,499 ($166,999 SBU portion). Award effective: 3/1/2007–2/28/2011.

20. co-PI: DOE DE-FG02-08ER25853 Efficient Databases for HPC Clusters. DOE Small Business Grant (with M. Farach-Colton and B. Kuszmaul). $1,750,000. Award effective: 2008–2011.

21. PI: CCF 0937822 HECURA: Collaborative: Multidimensional and String Indexes for Streaming Data. National Science Foundation (with M. Farach-Colton, Rutgers, and C. E. Leiserson and B. Kuszmaul, MIT). 1,000,000 ($199,993 SBU portion). Award effective: 9/1/2009–8/31/2013.

22. PI: CCF 1114809 AF: Small: Collaborative Research: Algorithms for Reallocation Problems. National Science Foundation (with M. Farach-Colton, Rutgers). $499,989 ($249,999 SBU portion). Award effective: 7/1/2011–6/30/2014.

23. PI: CCF 1217708 AF: SMALL: Collaborative Research: Data Structures for Parallel Algorithms. National Science Foundation (with K. Agrawal, Wash U, and J. T. Fineman, Georgetown). $449,939 ($138,999 SBU portion). Award effective: 8/1/2012–7/31/2015.

24. PI: IIS 1247726 BIGDATA: Mid-Scale: DCM: Collaborative Research: Eliminating the Data Ingestion Bottleneck in Big Data Applications. National Science Foundation (with M. Farach-Colton, Rutgers, and R. Johnson, SBU). $1,200,000, ($800,000 SBU portion). Award effective: 2/1/2013–1/31/2017.

25. co-PI: IIS 1251137 BIGDATA: Small: DCM: Collaborative Research: An Efficient, Versatile, Scalable, and Portable Storage System for Scientific Data Containers. National Science Foundation (with E. Zadok and R. Johnson, SBU, L. Shrira, Brandeis, and W. Benger, LSU). $746,289 ($444,267 SBU portion). Award effective: 7/1/2013–6/30/2017.

26. co-PI: SPIR: High Performance Lock Trees for Concurrent Data Structures. Tokutek (with R. Zhao, SBU). $15,128. Award effective: 6/1/2013–1/1/2014.

27. PI: CNS 1408695 CSR: Medium: Collaborative Research: FTFS: A Read/Write-optimized Fractal Tree File System. National Science Foundation (with Rob Johnson and Don Porter, SBU, Martin Farach-Colton, Rutgers, Bradley C. Kuszmaul, MIT). $1,103,993 ($623,993 SBU Portion). Award effective: 7/1/2014–6/30/2017.

28. co-PI: CCF 1439084 XPS: FULL: CCA: Collaborative Research: Cache-Adaptive Algorithms: How to Share Core among Many Cores. National Science Foundation (with Rob Johnson and Rezaul A Chowdhury, SBU). $800,000. Award effective: 12/1/2014–11/30/2017.

29. PI: Advanced Data Structures for Improved Cyber Resilience and Awareness in Untrusted Environments. Sandia National Laboratories (with Rob Johnson, SBU). $208,506. Award effective: 6/9/2015–9/30/2016.

30. PI: CCF 1617618. AF: Small: Collaborative Research: Maintaining Order. National Science Foundation (with Jeremy Fineman, Georgetown) $449,999 ($210,296 SBU portion). Award effective: 9/1/2016–8/31/2019.

31. co-PI: CCF 1725543 SPX: Collaborative Research: Harnessing the Power of High-Bandwidth Memory via Provably Efficient Parallel Algorithms. National Science Foundation (with Michael Ferdman, SBU and Benjamin Mosely, CMU). $750,000 ($500,000 SBU portion). Award effective: 9/15/2017–8/31/2021.

32. PI: CCF-BSF: 1716252 AF: Small: Collaborative Research: The Dictionary Problem Considered. National Science Foundation (with Martin Farach-Colton, Rutgers and Rob Johnson, VMware Research SBU). $500,000 ($250,000 SBU portion). Award effective: 9/1/2017–8/31/2020.

Refereed Journal Publications

1. Michael A. Bender and Howard A. Stone. An Integral Equation Approach to the Study of the Steady-State Current at Surface Microelectrodes. *Journal of Electroanalytical Chemistry and Interfacial Chemistry*, 351:29–55, 1993.

2. Michael A. Bender, Michel Gastaldo, and Michel Morvan. Parallel Interval Order Recognition and Construction of Interval Representations. *Theoretical Computer Science*, 143(1):73–91, 1995.

3. Yonatan Aumann, Michael A. Bender, and Lisa Zhang. Efficient Execution of Nondeterministic Parallel Programs on Asynchronous Systems. *Information and Computation*, 139(1):1–16, 1997.

4. Michael A. Bender and Chandra Chekuri. Performance Guarantees for the TSP with a Parameterized Triangle Inequality. *Information Processing Letters*, 73:17-21, 2000.

5. Mie Sato, Ingmar Bitter, Michael A. Bender, Arie E. Kaufman, and Masayuki Nakajima. Tree-Structure Extraction Algorithm for Accurate and Robust Skeletons (in Japanese). *The Journal of the Institute of Image Information and Television Engineers*, 2000.

6. Chandra Chekuri and Michael A. Bender. An Efficient Approximation Algorithm for Minimizing Makespan on Uniformly Related Machines. *Journal of Algorithms*, 41:212–224, 2001.

7. Michael A. Bender and Dana Ron. Testing Properties of Directed Graphs: Acyclicity and Connectivity. *Random Structures and Algorithms*, 20(2): 184–205, 2002.

8. Matthew Andrews, Michael A. Bender, and Lisa Zhang. New Algorithms for the Disk Scheduling Problem. *Algorithmica*, 32(2): 277–301, 2002.

9. Michael A. Bender and Michael O. Rabin. Online Scheduling of Parallel Programs on Heterogeneous Systems with Applications to Cilk. *Theory of Computing Systems*, 35: 289–304, 2002. Special Issue on *SPAA '00*.

10. Michael A. Bender, Antonio Fernández, Dana Ron, Amit Sahai, and Salil P. Vadhan. The Power of a Pebble: Exploring and Mapping Directed Graphs. *Information and Computation*, 176(1):1–21, 2002.

11. Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Steven Skiena. The Lazy Bureaucrat Scheduling Problem. *Information and Computation*, 184(1):129–146, 2003.

12. Carl M. Bender, Michael A. Bender, Erik D. Demaine, and Sándor P. Fekete. What is the Optimal Shape of a City? *Journal of Physics A: Mathematical and General*, 37:147–159, 2004.
    Journal of Physics A #1 Most Downloaded Article in 2004.

13. Michael A. Bender and Martin Farach-Colton. The Level Ancestor Problem Simplified. *Theoretical Computer Science—Special Issue on LATIN 2002*, 321(1):5–12, 2004.

14. Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S. B. Mitchell, Saurabh Sethia, and Steven Skiena. When Can You Fold a Map? *Computational Geometry: Theory and Applications (CGTA)*, 29(1):23–46, 2004. Special issue of selected papers from the 10th Annual Fall Workshop on Computational Geometry, 2000.

15. Marcelo O. Sztainberg, Esther M. Arkin, Michael A. Bender, and Joseph S. B. Mitchell. Theoretical and Experimental Analysis of Heuristics for the Freeze-Tag Robot Awakening Problem. *IEEE Transactions on Robotics and Automation*, 20(4):691–701, 2004.

16. Michael A. Bender, S. Muthukrishnan, and Rajmohan Rajaraman. Approximation Algorithms for Average Stretch Scheduling. *Journal of Scheduling*, 7(3):195–222, 2004. Special Issue on *SODA 02*.

17. Michael A. Bender, Ziyang Duan, John Iacono, and Jing Wu. A Locality-Preserving Cache-Oblivious Dynamic Dictionary. *Journal of Algorithms*, 3(2):115-136, 2004.
Journal of Algorithms Hottest Article, #1 Most Downloaded in 2004.

18. Michael A. Bender, Saurabh Sethia, and Steven Skiena. Data Structures for Maintaining Set Partitions. *Random Structures and Algorithms*, 25:43–67, 2004.

19. Yonatan Aumann and Michael A. Bender. Efficient Low-Contention Asynchronous Consensus with the Value-Oblivious Adversary Scheduler. *Distributed Computing*, 17(3): 191–207, 2005.

20. Michael A. Bender, Erik D. Demaine, and Martin Farach-Colton. Cache-Oblivious B-Trees. *SIAM Journal on Computing*, 35(2): 341–358, 2005.

21. Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Sándor P. Fekete, Joseph S. B. Mitchell, and Saurabh Sethia. Optimal Covering Tours with Turn Costs. *SIAM Journal on Computing*, 35(3): 531–566, 2005.

22. Michael A. Bender, Martin Farach-Colton, Giridhar Pemmasani, Steven Skiena, and Pavel Sumazin. Lowest Common Ancestors in Trees and Directed Acyclic Graphs. *Journal of Algorithms*, 57(2): 75–94, 2005.

23. Michael A. Bender, Martin Farach-Colton, and Miguel A. Mosteiro. Insertion Sort is $O(n \log n)$. *Theory of Computing Systems* 39(3): 391–397, 2006. Special Issue on *FUN '04*.

24. Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, Joseph S. B. Mitchell, and Martin Skutella. The Freeze-Tag Problem: How to Wake Up a Swarm of Robots. *Algorithmica*, 46(2): 193–221, 2006.

25. Lars Arge, Michael A. Bender, Erik D. Demaine, Bryan Holland-Minkley, and J. Ian Munro. Cache-Oblivious Priority Queue and Graph Algorithm Applications. *SIAM Journal on Computing*, 36(6): 1672–1695, 2007.

26. Michael A. Bender, Bryan Bradley, Geetha Jagannathan, and Krishnan Pillaipakkamnatt. Sum-of-Squares Heuristics for Bin Packing and Memory Allocation. *ACM Journal of Experimental Algorithms*, 12: 2.3, 2007.

27. Carl M. Bender and Michael A. Bender. Optimal Shape of a Blob. *Journal of Mathematical Physics*, 2007. 48, 073518, 2007.

28. Michael A. Bender and Haodong Hu. An Adaptive Packed-Memory Array. *Transactions on Database Systems*, 32(4): 2007. Special Issue on *PODS '06*.

29. Michael A. Bender, Raphaël Clifford, and Kostas Tsichlas. Scheduling Algorithms for Procrastinators. *Journal of Scheduling*, 11(2):95-104, 2008.

30. Michael A. Bender, Dongdong Ge, Simai He, Haodong Hu, Ron Y. Pinter, Firas Swidan, Steven Skiena, and Firas Swidan. Improved Bounds on Sorting with Length-Weighted Reversals. *Journal of Computer and Systems Sciences*, 74(5):744-774, 2008.

31. Michael A. Bender, David P. Bunde, Erik D. Demaine, Sándor P. Fekete, Vitus J. Leung, Henk Meijer, and Cynthia A. Phillips. Communication-Aware Processor Allocation for Supercomputers: Finding Point Sets of Small Average Distance. *Algorithmica*, 50(2): 279-298, 2008. Special Issue on *WADS '05*.

32. Kunal Agrawal, Michael A. Bender, and Jeremy T. Fineman. The Worst Page-Replacement Policy. *Theory of Computing Systems*, 44(2): 175-185, 2009. Special Issue on *FUN '07*.

33. Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Riko Jacob, and Elias Vicari. Optimal Sparse Matrix Dense Vector Multiplication in the I/O-Model. *Theory of Computing Systems*, 47(4): 934-962, 2010. Special Issue on *SPAA '07*.

34. Michael A. Bender, Bradley C. Kuszmaul, Shang-Hua Teng, and Kebin Wang. Optimal Cache-Oblivious Mesh Layouts. *Theory of Computing Systems*, 48(2): 269–296, 2011.

35. Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Dongdong Ge, Simai He, Haodong Hu, John Iacono, and Alejandro López-Ortiz. The Cost of Cache-Oblivious Searching. *Algorithmica*, 61(2): 463–505, 2011.

36. Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Valentin Polishchuk. The Snowblower Problem. *Computational Geometry*, 44(8): 370–384, 2011.

37. Michael A. Bender, Martin Farach-Colton, Rob Johnson, Russell Kraner, Bradley C. Kuszmaul, Dzejla Medjedovic, Pablo Montes, Pradeep Shetty, Richard P. Spillane, and Erez Zadok. Don't Thrash: How to Cache Your Hash on Flash. *PVLDB*, 5(11): 1627–1637, 2012.

38. Michael A. Bender, Ritwik Bose, Rezaul Chowdhury, and Samuel McCauley. The Kissing Problem: How to End a Gathering When Everyone Kisses Everyone Else Goodbye. *Theory of Computing Systems Special Issue on FUN12*, Theory of Computing Systems. 54(4): 715–730, 2014.

39. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Reallocation Problems in Scheduling. *Algorithmica*, 73: 389–409, 2015

40. William Jannen, Jun Yuan, Yang Zhan, Amogh Akshintala , John Esmet, Yizheng Jiao, Ankur Mittal, Prashant Pandey, Phaneendra Reddy, Leif Walsh, Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. BetrFS: Write-Optimization in a Kernel File System. *Transactions on Storage—Special Issue on USENIX FAST 2015*, 11(4): 18:1–18:29, 2015.

41. Michael A. Bender, Sándor P. Fekete, Alexander Kröller, Vincenzo Liberatore, Joseph S.B. Mitchell, Valentin Polishchuk, and Jukka Suomela. The Minimum Backlog Problem. *Theoretical Computer Science*, 605(9): 51–61, 2015.

42. Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Robert E. Tarjan. A New Approach to Incremental Cycle Detection and Related Problems. *ACM Transactions on Algorithms*, 12(2), 14:1–14:22, 2016.

43. Michael A. Bender, Roozbeh Ebrahimi, Haodong Hu, and Bradley C. Kuszmaul. B-trees and Cache-Oblivious B-trees with Different-Sized Atomic Keys. *Transactions on Database Systems*, 41(3), 19:1–19:33, July 2016.

44. Michael A. Bender, Rezaul A. Chowdhury, Pramod Ganapathi, Samuel McCauley, and Yuan Tang. The range 1 query (R1Q) problem. *Theoretical Computer Science*, available online, January 2016.

45. Michael A. Bender, Jonathan W. Berry, Simon D. Hammond, K. Scott Hemmert, Samuel McCauley, Branden Moore, Benjamin Moseley, Cynthia A. Phillips, David S. Resnick, and Arun Rodrigues. Two-level main memory co-design: Multi-threaded algorithmic primitives, analysis, and simulation. *Journal of Parallel and Distributed Computing.*, 102: 213–228, 2017.

46. Jun Yuan, Yang Zhan, William Jannen, Prashant Pandey, Amogh Akshintala, Kanchan Chandnani, Pooja Deo, Zardosht Kasheff, Leif Walsh, Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. Writes Wrought Right, and Other Adventures in File System Optimization. *Transactions on Storage—Special Issue on USENIX FAST 2016*, 13(1), 3:1–3:21, March 2017.

47. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Cost-Oblivious Storage Reallocation. *ACM Transactions on Algorithms*, Volume 13, Issue 3, Article 38, August 2017.

48. Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro. deBGR: an efficient and near-exact representation of the weighted de Bruijn graph. *Bioinformatics*, 33 (14): i133—i141, July 2017, `https://doi.org/10.1093/bioinformatics/btx261`.

49. Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro. Squeakr: an exact and approximate k-mer counting system. *Bioinformatics*, 34(4): 568–575, 2018. `https://doi.org/10.1093/bioinformatics/btx636`.

50. Prashant Pandey, Fatemeh Almodaresi, Michael A. Bender, and Michael Ferdman, Rob Johnson, and Rob Patro. Mantis: A Fast, Small, and Exact Large-Scale Sequence-Search Index. *Cell Systems*, 2018. In press. `http://dx.doi.org/10.1016/j.cels.2018.05.021`.

Refereed Conference Publications

51. Michael A. Bender and Donna K. Slonim. The Power of Team Exploration: Two Robots Can Learn Unlabeled Directed Graphs. *Proc. 35th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 75–85, 1994.

52. Yonatan Aumann and Michael A. Bender. Efficient Asynchronous Consensus with the Value-Oblivious Adversary Scheduler. *Proc. 23rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 622–633, 1996.

53. Matthew Andrews, Michael A. Bender, and Lisa Zhang. New Algorithms for the Disk Scheduling Problem. *Proc. 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 580–589, 1996.

54. Yonatan Aumann, Michael A. Bender, and Lisa Zhang. Efficient Execution of Nondeterministic Parallel Programs on Asynchronous Systems. *Proc. 8th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 270–276, 1996.

55. Yonatan Aumann and Michael A. Bender. Fault-Tolerant Data Structures. *Proc. 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 580–589, 1996.

56. Michael A. Bender, Soumen Chakrabarti, and S. Muthukrishnan. Flow and Stretch Metrics for Scheduling Continuous Job Streams. *Proc. 9th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 270–279, 1998.

57. Michael A. Bender, Antonio Fernández, Dana Ron, Amit Sahai, and Salil P. Vadhan. The Power of a Pebble: Exploring and Mapping Directed Graphs. *Proc. 30th Annual ACM Symposium on Theory of Computing (STOC)*, pages 269–278, 1998.

58. Chandra Chekuri and Michael A. Bender. An Efficient Approximation Algorithm for Minimizing Makespan on Uniformly Related Machines. *Proc. 6th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 383–393, 1998.

59. Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Steven Skiena. The Lazy Bureaucrat Scheduling Problem. *Proc. 6th Workshop on Discrete Algorithms (WADS)*, pages 80–85, 1999.

60. Michael A. Bender and Chandra Chekuri. Performance Guarantees for the TSP with a Parameterized Triangle Inequality. *Proc. 6th Workshop on Discrete Algorithms (WADS)*, pages 122–133, 1999.

61. Michael A. Bender and Martin Farach-Colton. The LCA Problem Revisited. *Latin American Theoretical INformatics (LATIN) 2000*, pages 88–94, 2000. **Imre Simon Test of Time Award.**

62. Michael A. Bender, Saurabh Sethia, and Steven Skiena. Data Structures for Maintaining Set Partitions. *Proc. 7th Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 83–96, 2000.

63. Michael A. Bender and Dana Ron. Testing Acyclicity of Directed Graphs in Sublinear Time. *Proc. 27th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 809–820, 2000.

64. Michael A. Bender and Michael O. Rabin. Scheduling Cilk Multithreaded Computations on Processors of Different Speeds. *Proc. 12th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 13–21, 2000.

65. Ingmar Bitter, Mie Sato, Michael A. Bender, Kevin T. McDonnell, Arie E. Kaufman, and Min Wan. CEASAR: A Smooth, Accurate, and Robust Centerline Extraction Algorithm. *IEEE Visualization*, pages 45–52, 2000.

66. Mie Sato, Ingmar Bitter, Michael A. Bender, and Arie E. Kaufman. TEASAR: Tree-Structure Extraction Algorithm for Accurate and Robust Skeletons. *Proc. 8th Pacific Conference on Computer Graphics and Applications Graphics*, pages 281–287, 2000.

67. Michael A. Bender, Erik D. Demaine, and Martin Farach-Colton. Cache-Oblivious B-Trees. *Proc. 41st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 399–409, 2000.

68. Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Sándor P. Fekete, Joseph S. B. Mitchell, and Saurabh Sethia. Optimal Covering Tours with Turn Costs. *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 138–147, 2001.

69. Michael A. Bender, Giridhar Pemmasani, Steven Skiena, and Pavel Sumazin. Finding Least Common Ancestors in Directed Acyclic Graphs. *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 845–854, 2001.

70. Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S. B. Mitchell, Saurabh Sethia, and Steven Skiena. When Can You Fold a Map? *Proc. 7th Workshop on Discrete Algorithms (WADS)*, pages 401–413, 2001.

71. Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, Joseph S. B. Mitchell, and Martin Skutella. The Freeze-Tag Problem: How to Wake Up a Swarm of Robots. *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 568–577, 2002.

72. Michael A. Bender, Ziyang Duan, John Iacono, and Jing Wu. A Locality-Preserving Cache-Oblivious Dynamic Dictionary. *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 29–38, 2002.

73. Michael A. Bender, S. Muthukrishnan, and Rajmohan Rajaraman. Improved Algorithms for Stretch Scheduling. *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 762–771, 2002.

74. Michael A. Bender and Martin Farach-Colton. The Level Ancestor Problem Simplified. *Latin American Theoretical INformatics (LATIN)*, pages 508–515, 2002.

75. Lars Arge, Michael A. Bender, Erik D. Demaine, Bryan Holland-Minkley, and J. Ian Munro. Cache-Oblivious Priority Queue and Graph Algorithm Applications. *Proc. 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 268–276, 2002.

76. Michael A. Bender, Richard Cole, and Rajeev Raman. Exponential Structures for Cache-Oblivious Algorithms. *Proc. 29th International Colloquium on Automata, Languages, and Programming (ICALP)*, 195–207, 2002.

77. Marcelo Sztainberg, Esther M. Arkin, Michael A. Bender, and Joseph S. B. Mitchell. Analysis of Heuristics for the Freeze-Tag Problem. *Proc. 8th Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 270–279, 2002.

78. Michael A. Bender, Richard Cole, Erik D. Demaine, and Martin Farach-Colton. Scanning and Traversing: Maintaining Data for Traversals in a Memory Hierarchy. *Proc. 10th European Symposium on Algorithms (ESA)*, pages 139–151, 2002.

79. Michael A. Bender, Richard Cole, Erik D. Demaine, Martin Farach-Colton, and Jack Zito. Two Simplified Algorithms for Maintaining Order in a List. *Proc. 10th European Symposium on Algorithms (ESA)*, pages 152–164, 2002.

80. Michael A. Bender, Erik D. Demaine, and Marin Farach-Colton. Efficient Tree Layout in a Multilevel Memory Hierarchy. *Proc. 10th European Symposium on Algorithms (ESA)*, pages 165–173, 2002.

81. Vitus J. Leung, Esther M. Arkin, Michael A. Bender, David P. Bunde, Jeanette Johnston, Alok Lal, Joseph S. B. Mitchell, Cynthia A. Phillips, and Steven S. Seiden. Processor Allocation on Cplant: Achieving General Processor Locality Using One-Dimensional Allocation Strategies. *Proc. 4th IEEE International Conference on Cluster Computing (CLUSTER)*, pages 296–304, 2002.

82. Tien-Ruey Hsiang, Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, and Joseph S. B. Mitchell. Algorithms for Rapidly Dispersing Robot Swarms in Unknown Environments. *Proc. 5th Workshop on Algorithmic Foundations of Robotics (WAFR)*, 77–94, 2002.

83. Tien-Ruey Hsiang, Esther M. Arkin, Michael A. Bender, Sándor P. Fekete, and Joseph S. B. Mitchell. Online Dispersion Algorithms for Swarms of Robots. *Proc. 19th Annual ACM Symposium on Computational Geometry (SoCG)*, Video/DVD, pages 382–383, 2003.

84. Esther M. Arkin, Michael A. Bender, Dongdong Ge, Simai He, and Joseph S. B. Mitchell. Improved Approximation Algorithms for the Freeze-Tag Problem. *Proc. 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 295–303, 2003.

85. Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Dongdong Ge, Simai He, Haodong Hu, John Iacono, and Alejandro López-Ortiz. The Cost of Cache-Oblivious Searching. *Proc. 44th Annual Symposium on Foundations of Computer Science (FOCS),* pages 271–280, 2003.

86. Michael A. Bender, Dongdong Ge, Simai He, Haodong Hu, Ron Y. Pinter, Steven Skiena, and Firas Swidan. Improved Bounds on Sorting with Length-Weighted Reversals. *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 912–921, 2004.

87. Michael A. Bender, Bryan Bradley, Geetha Jagannathan, and Krishnan Pillaipakkamnatt. The Robustness of the Sum-of-Squares Algorithm for Bin Packing. *Proc. 6th Workshop on Algorithm Engineering and Experiments (ALENEX)*, 18-30, 2004.

88. Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Charles E. Leiserson. On-the-Fly Maintenance of Series-Parallel Relationships in Fork-Join Multithreaded Programs. *Proc. 16th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 133–144, 2004.

89. Michael A. Bender, Martin Farach-Colton, and Miguel A. Mosteiro. Insertion Sort is $O(n \log n)$. *Proc. 3rd International Conference on Fun with Algorithms (FUN)*, pages 16–23, 2004.

90. Firas Swidan, Michael A. Bender, Dongdong Ge, Simai He, Haodong Hu, and Ron Y. Pinter. Sorting by Length-Weighted Reversals: Dealing with Signs and Circularity. *Proc. 15th Annual Combinatorial Pattern Matching Symposium (CPM)*, Volume 3109 of *Lecture Notes in Computer Science*, pages 32–46, 2004.

91. Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Bradley C. Kuszmaul. Concurrent Cache-Oblivious B-Trees. *Proc. 17th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 228–237, 2005.

92. Michael A. Bender, Martin Farach-Colton, Simai He, Bradley C. Kuszmaul, and Charles E. Leiserson. Adversarial Contention Resolution for Simple Channels. *Proc. 17th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 325–332, 2005.

93. Michael A. Bender, David P. Bunde, Erik D. Demaine, Sándor P. Fekete, Vitus J. Leung, Henk Meijer, and Cynthia A. Phillips. Communication-Aware Processor Allocation for Supercomputers. *Proc. 9th Workshop on Algorithms and Data Structures (WADS)*, Springer LNCS 3608, pages 169–181, 2005.

94. Michael A. Bender and Haodong Hu. An Adaptive Packed-Memory Array. *Proc. 25th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 20–29, 2006. **Best Newcomer Award**.

95. Michael A. Bender, Martin Farach-Colton, and Bradley C. Kuszmaul. Cache-Oblivious String B-Trees. *Proc. 25th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 233–242, 2006.

96. Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Valentin Polishchuk. The Snowblower Problem. *Proc. 7th Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2006.

97. Michael A. Bender, Jeremy T. Fineman, and Seth Gilbert. Contention Resolution with Heterogeneous Job Sizes. *Proc. 14th Annual European Symposium on Algorithms (ESA)*, pages 112–123, 2006.

98. Michael A. Bender and Cynthia A. Phillips. Scheduling DAGs on Asynchronous Processors. *Proc. 19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 35–45, 2007.

99. Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Riko Jacob, and Elias Vicari. Optimal Sparse Matrix Dense Vector Multiplication in the I/O-Model. *Proc. 19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 61–70, 2007.

100. Michael A. Bender, Martin Farach-Colton, Jeremy T. Fineman, Yonatan Fogel, Bradley C. Kuszmaul, and Jelani Nelson. Cache-Oblivious Streaming B-Trees. *Proc. 19th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 81–92, 2007.

101. Kunal Agrawal, Michael A. Bender, and Jeremy T. Fineman. The Worst Page-Replacement Policy. *Proc. 4th International Conference on Fun With Algorithms (FUN)*, pages 135–145, 2007.

102. Michael A. Bender, Jeremy T. Fineman, and Seth Gilbert. A New Approach to Incremental Topological Ordering. *Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1108–1115, 2009.

103. Michael A. Bender, Sándor P. Fekete, Tom Kamphans, and Nils Schweer. Maintaining Arrays of Contiguous Objects. *Proc. 17th International Symposium on the Fundamentals of Computation Theory (FCT)*, LLNCS Volume 6599, pages 14–25, 2009.

104. Michael A. Bender, Haodong Hu, and Bradley C. Kuszmaul. Performance Guarantees for B-trees with Different-Sized Atomic Keys. *Proc. 29th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 305–316, 2010.

105. Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, Dzejla Medjedovic, Pablo Montes, Pradeep Shetty, Richard P. Spillane, and Erez Zadok. Don't Thrash: How to Cache Your Hash on Flash. *Proc. 3rd USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2011.

106. Michael A. Bender and Seth Gilbert. Mutual Exclusion with $O(\log^2 \log n)$ Amortized Work. *Proc. 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 728–737, 2011.

107. Michael A. Bender, Ritwik Bose, Rezaul Chowdhury, and Samuel McCauley. The Kissing Problem: How to End a Gathering When Everyone Kisses Everyone Else Goodbye. *Proc. Sixth International Conference on Fun with Algorithms (FUN)*, pages 28–39, 2012.

108. John Esmet, Michael A. Bender, Martin Farach-Colton, and Bradley C. Kuszmaul. The TokuFS Streaming File System. *Proc. 4th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2012.

109. Dan Alistarh, Michael A. Bender, Seth Gilbert, and Rachid Guerraoui. How to Allocate Tasks Asynchronously. *Proc. 53nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 331–340, 2012.

110. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Reallocation Problems in Scheduling. *Proc. 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 271–279, 2013.

111. Michael A. Bender, David P. Bunde, Vitus J. Leung, Samuel McCauley, and Cynthia A. Phillips. Efficient Scheduling to Minimize Calibrations. *Proc. 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 280–287, 2013.

112. Dan Alistarh, James Aspnes, Michael A. Bender, Rati Gelashvili, and Seth Gilbert. Dynamic Task Allocation. *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 416–435, 2014.

113. Michael A. Bender, Roozbeh Ebrahimi, Jeremy T. Fineman, Golnaz Ghasemiesfeh, Rob Johnson, and Samuel McCauley. Cache-Adaptive Algorithms. *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 958–971, 2014.

114. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Cost-Oblivious Storage Reallocation. *Proc. 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS)*, pages 278-288, 2014.

115. Michael A. Bender, Rezaul Chowdhury, Pramod Ganapathi, Samuel McCauley, and Yuan Tang. The Range 1 Query (R1Q) Problem. *Proc. 20th International Computing and Combinatorics Conference (COCOON)*, pages 116-128, 2014.

116. Michael A. Bender, Martín Farach-Colton, Mayank Goswami, Dzejla Medjedovic, Pablo Montes, and Meng-Tsung Tsai. The Batched Predecessor Problem in External Memory. *Proc. 22nd Annual European Symposium on Algorithms (ESA)*, pages 112-124, 2014.

117. William Jannen, Jun Yuan, Yang Zhan, Amogh Akshintala, John Esmet, Yizheng Jiao, Ankur Mittal, Prashant Pandey, Phaneendra Reddy, Leif Walsh, Michael Bender, Martín Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. BetrFS: A Right-Optimized Write-Optimized File System. *Proc 13th USENIX Conference on File and Storage Technologies (FAST)*, 2015.

118. Michael A. Bender, Jonathan Berry, Simon D. Hammond, K. Scott Hemmert, Samuel McCauley, Branden Moore, Benjamin Moseley, Cynthia A. Phillips, David Resnick, and Arun Rodrigues. Two-Level Main Memory Co-Design: Multi-Threaded Algorithmic Primitives, Analysis, and Simulation. *Proc. 29th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2015. **Best paper award.**

119. Michael A. Bender, Martin Farach-Colton, Sándor P. Fekete, Jeremy T. Fineman, and Seth Gilbert. Cost-Oblivious Reallocation for Scheduling and Planning. *Proc. 27th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 143–154, 2015.

120. Michael A. Bender, Jonathan Berry, Simon D. Hammond, Branden Moore, Benjamin Moseley, and Cynthia A. Phillips. k-Means Clustering on Two-Level Memory Systems. *International Symposium on Memory Systems (MEMSYS)*, 197–205, 2015.

121. Michael A. Bender, Samuel McCauley, Andrew McGregor, Shikha Singh, and Hoa T. Vu. Run Generation Revisited: What Goes Up May or May Not Come Down. *Proc. 26th International Symposium on Algorithms and Computation (ISAAC)*, pages 703–714, 2015.

122. Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, and Maxwell Young. How to Scale Exponential Backoff: Constant Throughput, Polylog Access Attempts, and Robustness. *Proc. 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 636–654, 2016.

123. Jun Yuan, Yang Zhan, William Jannen, Prashant Pandey, Amogh Akshintala, Kanchan Chandnani, Pooja Deo, Zardosht Kasheff, Leif Walsh, Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. Optimizing Every Operation in a Write-Optimized File System. *Proc. 14th USENIX Conference on File and Storage Technologies (FAST)*, pages 1–14, 2016. **Best paper award.**

124. Michael A. Bender, Rezaul Chowdhury, Alex Conway, Martín Farach-Colton, Pramod Ganapathi, Rob Johnson, Samuel McCauley, Bertrand Simon, and Shikha Singh. The I/O Complexity of Computing Prime Tables. *Proc. 12th Latin American Theoretical Informatics Symposium (LATIN)*, pages, 192–206, 2016.

125. Michael A. Bender, Samuel McCauley, Bertrand Simon, Shikha Singh, and Frédéric Vivien. Resource Optimization for Program Committee Members: A Subreview Article. *Proc. Eighth International Conference on Fun with Algorithms (FUN)*, page 7:1–7:20, 2016.

126. Michael A. Bender, Tsvi Kopelowitz, Seth Pettie, and Maxwell Young. Contention Resolution with Log-Logstar Channel Accesses. *Proc. 48th Annual Symposium on the Theory of Computing (STOC)*, pages 499–508, 2016.

127. William Jannen, Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. Lazy Analytics: Let Other Queries Do the Work For You. *Proc. 8th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage)*, 2016.

128. Michael A. Bender, Jon Berry, Rob Johnson, Thomas M. Kroeger, Samuel McCauley, Cynthia A. Phillips, Bertrand Simon, Shikha Singh, and David Zage. Anti-Persistence on Persistent Storage: History-Independent Sparse Tables and Dictionaries. *Proc. 35th ACM Symposium on Principles of Database Systems (PODS)*, pages 289–302, 2016.

129. Michael A. Bender, Erik D. Demaine, Roozbeh Ebrahimi, Jeremy T. Fineman, Rob Johnson, Andrea Lincoln, Jayson Lynch, and Samuel McCauley. Cache-Adaptive Analysis. *Proc. 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 135–144, 2016.

130. Michael A. Bender, Jeremy T. Fineman, Seth Gilbert, Tsvi Kopelowitz, and Pablo Montes. File Maintenance: When in Doubt, Change the Layout! *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1502–1522, January 2017.

131. Peyman Afshani, Michael A. Bender, Martin Farach-Colton, Jeremy T. Fineman, Mayank Goswami, and Meng-Tsung Tsai. Cross-Referenced Dictionaries and the Limits of Write Optimization. *Proc. 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1523–1532, January 2017.

132. Alexander Conway, Ainesh Bakshi, Yizheng Jiao, Yang Zhan, Michael A. Bender, William Jannen, Rob Johnson, Bradley C. Kuszmaul, Donald E. Porter, Jun Yuan, and Martin Farach-Colton. File Systems Fated for Senescence? Nonsense, Says Science! *Proc. 15th USENIX Conference on File and Storage Technologies (FAST)*, pages 45–58, February 2017.

133. Prashant Pandey, Michael A. Bender, Rob Johnson, and Rob Patro. A General-Purpose Counting Filter: Making Every Bit Count. *Proc. 2017 International Conference on Management of Data (SIGMOD)*, pages 775–787, May 2017.

134. Michael A. Bender, Martin Farach-Colton, Rob Johnson, Simon Mauras, Tyler Mayer, Cynthia Phillips, and Helen Xu. Write-Optimized Skip Lists. *Proc. 36th ACM Symposium on Principles of Database Systems (PODS)*, pages 69–78, May 2017.

135. Yang Zhan, Alexander Conway, Yizheng Jiao, Eric Knorr, Michael A. Bender, Martin Farach-Colton, William Jannen, Rob Johnson, Donald E. Porter, and Jun Yuan. The Full Path to Full-Path Indexing. *Proc. 16th USENIX Conference on File and Storage Technologies (FAST)*, pages 123–138, February 2018.

136. Michael A. Bender, Martin Farach-Colton, Mayank Goswami, Rob Johnson, Samuel McCauley, and Shikha Singh. Bloom filters, adaptivity, and the dictionary problem. *Proc 59th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, October 2018.

Technical Magazine Articles

137. Michael A. Bender, Jeremy T. Fineman, Mahnush Movahedi, Jared Saia, Varsha Dani, Seth Gilbert, Seth Pettie, and Maxwell Young. Resource-Competitive Algorithms. *SIGACT News*, 46(3): 57–71, 2015.

138. Michael Bender, Martin Farach-Colton, William Jannen, Rob Johnson, Bradley C. Kuszmaul, Donald E. Porter, Jun Yuan, and Yang Zhan. An Introduction to B^{ε}-Trees and Write-Optimization. *:login; magazine*, 40(5): 22–28, October 2015.

Presentations

1. **Washington University.** St. Louis, MO. Jan 1995.
   The Power of Team Exploration: Two Robots Can Learn Unlabeled Directed Graphs.

2. **SPAA 96.** Padua, Italy. June 1996.
   Efficient Execution of Nondeterministic Parallel Programs on Asynchronous Systems.

3. **ICALP 96.** Paderborn, Germany. July 1996.
   Efficient Asynchronous Consensus with the Value-Oblivious Adversary Scheduler.

4. **Summer Institute on Parallel Discrete Algorithms.** Halifax, Nova Scotia. July 1996.
   Efficient Asynchronous Consensus with the Value-Oblivious Adversary Scheduler.

5. **FOCS 96.** Burlington, VT. Oct 1996.
   Fault-Tolerant Data Structures.

6. **Bell Labs.** Murray Hill, NJ. Sep 1997.
   Flow and Stretch Metrics for Scheduling Continuous Job Streams.

7. **Carnegie Mellon.** Pittsburgh, PA. Nov 1997.
   The Power of a Pebble: Exploring and Mapping Directed Graphs.

8. **DIMACS Workshop on Randomization Methods in Algorithm Design.** Princeton, NJ. Dec 1997.
   Efficient Asynchronous Consensus with the Value-Oblivious Adversary Scheduler.

9. **Michigan State.** East Lansing, MI. Feb 1998.
   New Algorithms and Metrics for Scheduling.

10. **Arizona State.** Phoenix, AZ. Feb 1998.
    New Algorithms and Metrics for Scheduling.

11. **Stony Brook University.** Stony Brook, NY. Mar 1998.
    New Algorithms and Metrics for Scheduling.

12. **Polytechnic University.** Brooklyn, NY. Mar 1998.
    New Algorithms and Metrics for Scheduling.

13. **Bell Labs, Lucent Technologies.** Murray Hill, NJ. Mar 1998.
    New Algorithms and Metrics for Scheduling.

14. **Case Western Reserve University.** Cleveland, OH. Mar 1998.
    New Algorithms and Metrics for Scheduling.

15. **U of Arizona.** Tucson, AZ. Apr 1998.
    New Algorithms and Metrics for Scheduling.

16. **U of Houston.** Houston, TX. Apr 1998.
    New Algorithms and Metrics for Scheduling.

17. **AT&T Research.** Florham Park, NJ. Apr 1998.
    New Algorithms and Metrics for Scheduling.

18. **U of New Mexico.** Albuquerque, NM. Apr 1998.
    New Algorithms and Metrics for Scheduling.

19. **STOC 98.** Dallas, TX. May 1998.
    The Power of a Pebble: Exploring and Mapping Directed Graphs.

20. **U of Maryland.** College Park, MD. Mar 1999.
    New Algorithms and Metrics for Scheduling Job Streams.

21. **Escuela de Ciencias Informáticas 99.** U of Buenos Aires, Argentina. July 1999.
    Estructuras de Datos Avanzados (Advanced Data Structures).
    Week-long Course.

22. **WADS 99.** Vancouver, Canada. Aug 1999.
    General Multiprocessor Task Scheduling: Approximate Solutions in Linear Time.
    Conference Talk for K. Jansen and L. Porkolab.

23. **WADS 99.** Vancouver, Canada. Aug 1999.
    Performance Guarantees for the TSP with a Parameterized Triangle Inequality.

24. **WADS 99.** Vancouver, Canada. Aug 1999.
    The Lazy Bureaucrat Scheduling Problem.

25. **Dagstuhl Seminar: Scheduling in Computer and Manufacturing Systems.** Germany. Oct 1999.
Flow and Stretch Metrics for Scheduling Continuous Job Streams.

26. **Stony Brook University.** Stony Brook, NY. Feb 2000.
Algorithms for Asynchronous Consensus. Operations Research Seminar.

27. **Stony Brook University.** Stony Brook, NY. Mar 2000.
Algorithms for Graph Exploration. Operations Research Seminar.

28. **LATIN 2000.** Punta del Este, Uruguay. Apr 2000.
The LCA Problem Revisited.

29. **SPAA 2000.** Bar Harbor, ME. July 2000.
Scheduling Cilk Multithreaded Computations on Processors of Different Speeds.

30. **Dagstuhl Seminar: Experimental Algorithms.** Germany Sep 2000.
Cache-Oblivious Algorithms and Data Structures.

31. **DIMACS Workshop on Sublinear Algorithms.** Princeton, NJ Oct 2000.
Cache-Oblivious Algorithms and Data Structures.

32. **MARS/DR/SDR DARPA PI meeting.** Knoxville, TN. Oct 2000.
Algorithms in Support of Pheromone Robotics.

33. **FOCS 2000.** Redondo Beach, CA. Nov 2000.
Cache-Oblivious B-Trees.

34. **Yale.** New Haven, CT. Mar 2001.
Cache-Oblivious Data Structures.

35. **Stony Brook Univ.** Stony Brook, NY. Apr 2001.
Cache-Oblivious B-Trees. Operations Research Seminar.

36. **Brookhaven National Labs.** Brookhaven, NY. Mar 2001.
Cache-Oblivious Algorithms.

37. **McGill.** Montréal, Canada. Apr 2001.
The Power of a Pebble: Exploring and Mapping Directed Graphs.

38. **IBM TJ Watson.** Yorktown Heights, NY. May 2001.
Cache-Oblivious Algorithms.

39. **MARS/DR/SDR DARPA PI meeting.** Nashville, TN. July 2001.
Algorithms in Support of Pheromone Robotics.

40. **First Arizona Workshop on Algorithms.** Tempe, AZ. Nov 2001.
Cache-Oblivious Data Structures and Algorithms.

41. **SODA 2002.** San Franciso, CA. Jan 2002.
The Freeze-Tag Problem: How to Wake Up a Swarm of Robots.

42. **Dagstuhl Seminar on Data Structures.** Germany. Feb 2002.
Two Simplified Algorithms for Maintaining Order in a Linked List.

43. **Brown.** Providence, RI. May 2002.
Cache-Oblivious Data Structures.

44. **Dagstuhl Seminar: Scheduling in Computer and Manufacturing Systems.** Germany. June 2002.
Scheduling and Resource Allocation in Commodity Supercomputers.

45. **Sandia Petaflops Systems Workshop.** Albuquerque, NM. June 2002.
Cache-Oblivious Algorithms.

46. **Bell Labs.** Murray Hill, NJ. July 2002.
Cache-Oblivious Data Structures.

47. **AT&T Research.** Florham Park, NJ. July 2002.
Cache-Oblivious Data Structures.

48. **Sun Microsystems Inc.** Burlington, MA. Aug 2002.
Cache-Oblivious Data Structures.

49. **ESA 2002.** Rome, Italy. Sep 2002.
Efficient Tree Layout in a Multilevel Memory Hierarchy.

50. **ESA 2002.** Rome, Italy. Sep 2002.
Two Simplified Algorithms for Maintaining Order in a List.

51. **Polytechnic University.** Brooklyn, NY. Oct 2002.
Cache-Oblivious Data Structures.

52. **MIT.** Cambridge, MA. Oct 2002.
Cache-Oblivious Data Structures.

53. **Workshop on Models and Algorithms for Planning and Scheduling Problems (MAPSP 2003).** Aussois, France. Apr 2003.
The Freeze-Tag Problem: How to Awaken a Swarm of Robots.

54. **SPAA 2003.** San Diego, CA. May 2003.
Improved Approximation Algorithms for the Freeze-Tag Problem.

55. **SoCG 2003.** San Diego, CA. May 2003.
Online Dispersion Algorithms for Swarms of Robots.

56. **SODA 2004.** New Orleans, LA. Jan 2004.
Improved Bounds on Sorting with Length-Weighted Reversals.

57. **SIAM Conference on Parallel Processing for Scientific Computing (PP04).** San Francisco, CA. Feb 2004.
Communication-Aware Processor Allocation for Supercomputers.

58. **SIAM Workshop Combinatorial Scientific Computing.** San Francisco, CA. Feb 2004.
Cache-Oblivious Data Structures and Algorithms.

59. **Boston University.** Boston, Massachusetts. Mar 2004.
Cache-Oblivious B-Trees.

60. **Harvard.** Cambridge, Massachusetts. Mar 2004.
Cache-Oblivious Searching.

61. **NSF Next Generation Software Program – PI Workshop at IPDPS.**
Sante Fe, NM. Apr 2004.
Worst-Case Analysis of Randomized Backoff.

62. **Dagstuhl Seminar: Scheduling in Computer and Manufacturing Systems.**
Germany. Apr 2004.
Scheduling Algorithms for Transactional Memory.

63. **Northeastern.** Boston, Massachusetts. Apr 2004.
Cache-Oblivious B-Trees.

64. **FUN with Algorithms.** Elba, Italy. May 2004.
Insertion Sort is $O(N \log N)$.

65. **Dagstuhl Seminar: Cache-Oblivious and Cache-Aware Algorithms.** Germany. July 2004.
Efficient Tree Layouts.

66. **Sun Microsystems Laboratories.** Burlington, Massachusetts. Aug 2004.
Concurrent Cache-Oblivious Search Trees.

67. **3rd International Conference on Parallel Computing Systems.** Colima, Mexico. Sep 2004.
Communication-Aware Processor Allocation for Supercomputers.
Plenary Talk.

68. **U of Southampton.** Southampton, UK. Oct 2004.
Communication-Aware Processor Allocation for Supercomputers.

69. **U of Durham.** Durham, UK. Oct 2004.
Cache-Oblivious Data Structures.

70. **Kings College London.** London, UK. Oct 2004.
Communication-Aware Processor Allocation for Supercomputers.

71. **U of Nottingham.** Nottingham, UK. Oct 2004.
Communication-Aware Processor Allocation for Supercomputers.

72. **U of Edinburgh.** Edinburgh, UK. Nov 2004.
Cache-Oblivious Data Structures.

73. **U of Braunschweig.** Braunschweig, Germany. Nov 2004.
Cache-Oblivious Data Structures.

74. **U of Hertfordshire.** Hatfield, UK. Nov 2004.
Cache-Oblivious Data Structures.

75. **U of Leeds.** Leeds, UK. Nov 2004.
Communication-Aware Processor Allocation for Supercomputers.

76. **Imperial College.** London, UK. Dec 2004.
Cache-Oblivious Data Structures.

77. **Université de Grenoble.** Grenoble, France. Dec 2004.
Communication-Aware Processor Allocation for Supercomputers (in French).

78. **5th Haifa Workshop on Interdisciplinary Applications of Graph Theory, Combinatorics and Computing.** Haifa, Israel. May 2005.
Adversarial Contention Resolution.
Plenary Talk.

79. **U of Patras.** Patras, Greece. June 2005.
Cache-Oblivious B-trees.

80. **6to. Congreso Int. de las Ciencias Computacionales.** Colima, Mexico. Sep 2005.
Cache-Oblivious Data Structures (in Spanish).
Plenary Talk.

81. **Fall Workshop on Computational Geometry.** Philadelphia, PA. Nov 2005.
What is the Optimal Shape of a Blob?

82. **York U.** Toronto, CA. Nov 2005.
Adversarial Contention Resolution.

83. **U of Toronto.** Toronto, CA. Nov 2005.
Cache-Oblivious B-trees.

84. **SIAM Conference on Parallel Processing for Scientific Computing (PP06).** San Francisco, CA. Feb 2006.
Cache-Oblivious Algorithms for Massive Data.

85. **Dagstuhl Seminar on Data Structures.** Germany. Feb 2006.
An Adaptive Packed-Memory Array.

86. **Workshop on Parallelism in Algorithms and Architectures.** U Maryland, MD. May 2006.
An Adaptive Packed-Memory Array.

87. **CIRM Workshop: Scheduling Algorithms for New Emerging Applications.** Luminy, France. June 2006.
Scheduling Algorithms for Procrastinators.

88. **Université de Grenoble.** Grenoble, France. June 2006.
An Adaptive Packed-Memory Array (in French).

89. **PODS 2006.** Chicago, IL. June 2006.
An Adaptive Packed-Memory Array. **Award Presentation.**

90. **High-End Computing File Systems and I/O (HEC FSIO) 2006.** Arlington, VA. Aug 2006.
Streaming B-trees.

91. **Encuentro Mexicano de Computación (ENC 2006).** San Luis Potosí, MX. Sep 2006.
An Adaptive Packed-Memory Array.
Plenary Talk.

92. **Dagstuhl Seminar on Robot Navigation.** Germany. Oct 2006.
The Snowblower Problem.

93. **Los Alamos National Laboratories.** Los Alamos, NM. Jan 2007.
An Adaptive Packed-Memory Array.

94. **Columbia.** New York, NY. Apr 2007.
An Adaptive Packed-Memory Array.

95. **Oberwolfach Workshop on Algorithm Engineering.** Germany. June 2007.
Engineering B-trees and Cache-Oblivious B-trees on Real Memory Hierarchies (Ignorance is Bliss)

96. **SPAA 2007.** San Diego, CA. June 2007.
Scheduling DAGs on Asynchronous Processors.

97. **U of Buenos Aires.** Buenos Aires, Argentina. Oct 2007.
An Adaptive Packed-Memory Array (in Spanish).

98. **Washington Univ.** St. Louis, MO. Dec 2007.
An Adaptive Packed-Memory Array.

99. **Dagstuhl Seminar on Data Structures.** Germany. Feb 2008.
Cache-Oblivious Streaming B-Trees.

100. **High-End Computing File Systems and I/O (HEC FSIO).** Arlington, VA. Aug 2008.
Techniques for Streaming File Systems and Databases.

101. **IBM TJ Watson.** Yorktown Heights, NY. Nov 2008.
Cache-Oblivious Streaming B-trees.

102. **Quantum Mechanics in the Complex Domain.** St. Louis, MO. Mar 2009.
An Adaptive Packed-Memory Array.

103. **NSF Workshop on the Science of Power Management.** Arlington, VA. Apr 2009.
An Adaptive Packed-Memory Array. **Spotlight Talk.**

104. **The Graduate Center, CUNY.** New York, NY. May 2009.
An Adaptive Packed-Memory Array.

105. **Symposium on Experimental Algorithms (SEA).** Dortmund, Germany. June 2009.
From Streaming B-trees to Tokutek: How a Theoretician Learned to be VP of Engineering.
Keynote Talk.

106. **IBM Supercomputing Professional Interest Community.** IBM TJ Watson. Aug 2009.
Performance of Fractal Tree Databases.

107. **Scalable Approaches to High Performance and High Productivity Computing (ScalPerf).** Bertinoro, Italy. Sep 2009.
Performance of Fractal Tree Databases.

108. **Brookhaven National Labs.** Brookhaven, NY. Feb 2010.
Performance of Fractal Tree Databases.

109. **Dagstuhl Seminar on Data Structures.** Germany. Mar 2010.
Fractal Tree Databases: Data Structures for Fun and Profit.

110. **New Topics in Distributed Algorithms.** EPFL, Switzerland. Apr 2010.
Fractal Tree Databases: Concurrency Challenges.

111. **St. Louis University.** St. Louis, MO. Apr 2010.
Performance of Fractal Tree Databases.

112. **Annual Case Lecture at St. Louis University.** St. Louis, MO. Apr 2010
Gaps in My Education: Mailboxes, Libraries, and How to Insert into an Array.
Graduation and Award-Ceremony Talk.

113. **PODS 2010.** Indianapolis, IN. June 2010.
Performance Guarantees for B-trees with Different-Size Atomic Keys.

114. **High-End Computing File Systems and I/O (HEC FSIO).** Arlington, VA. Aug 2010.
Multidimensional and String Indexes for Streaming Data.

115. **CNRS Workshop on New Challenges in Scheduling Theory.** Fréjus, France. Sep 2010.
Scheduling DAGs on Asynchronous Processors.

116. **Hofstra University.** Hempstead, NY. Sep 2010.
How to Index Massive Data Sets Quickly.

117. **Morrelly Homeland Security Center.** Bethpage, NY. Jan 2011.
How to Index Massive Data Sets Quickly.

118. **Sustainable Energy-Efficient Data Management (SEEDM).** Arlington, VA. May 2011.
How Fast Indexing Makes Databases Greener.

119. **Bertinoro Workshop on Algorithms and Data Structures (ADS).** Bertinoro, Italy. June 2011.
Don't Thrash: How to Cache Your Hash on Flash.

120. **High-End Computing File Systems and I/O (HEC FSIO).** Arlington, VA. Aug 2011.
Better Metadata Management Through Data Structures.

121. **BIRS Workshop on Probabilistic versus Deterministic Techniques for Shared Memory Computation.** Banff Centre, Canada. Feb 2012.
Asynchronous Shared-Memory Mutual Exclusion in $O(\log^2 \log n)$ RMRs.

122. **STOC Tutorial on Algorithms for Memory-Sensitive Computing.**
New York, NY. May 2012.
Tutorial. *Databases and External Memory: Indexes, Write Optimization, and Cryptosearches.*

123. **NSF Workshop on Research Directions in the Principles of Parallel Computation.**
Pittsburgh, PA. June 2012
Indexing Big Data.

124. **Dagstuhl Seminar on Database Workload Management.** Germany. July 2012.
Write-Optimized Data Structures.

125. **Dagstuhl Seminar on Database Workload Management.** Germany. July 2012.
The Procrastination Scheduling Problem. "Gong Show" Presentation.

126. **Sandia National Laboratories.** Albuquerque, NM. Aug 2012.
Indexing Massive Data Sets.

127. **6th Extremely Large Databases Conference, Workshop, and Tutorials (XLDB).**
Stanford, CA. Sep 2012.
Tutorial. *Data Structures and Algorithms for Big Databases.*

128. **NSF Center for Dynamic Data Analytics (CDDA).** Stony Brook, NY. Mar 2013.
Eliminating the Data Ingestion Bottleneck in Big Data Applications.

129. **Waterloo.** Waterloo, CA. Mar 2013.
Data Structures for Indexing Massive Data Sets. Database Seminar.

130. **White House Big Data Workshop.** White House Conference Center, Washington, DC. May 2013.
Write Optimization: the Future of Databases/File Systems. One-minute talk.

131. **CEWIT Industrial Advisory Board.** Stony Brook, NY. May 2013.
Overview of Big Data in Computer Science Department.

132. **Workshop on Algorithms and Data Structures (ADS).** Bertinoro, Italy. June 2013.
Cache-Adaptive Algorithms.

133. **7th Extremely Large Databases Conference, Workshop, and Tutorials (XLDB).**
Stanford, CA. Sep 2013.
Tutorial. *Data Structures and Algorithms for Big Databases.*

134. **Rutgers.** New Brunswick, NJ. Oct 2013.
Write Optimization.

135. **SUNY Conference (SUNYCON): Building a Smarter University.** Oct 2013.
Research and Education on Big Data.

136. **Two Sigma.** New York, NY. Nov 2013.
Read-Write Optimization.

137. **CEWIT.** Stony Brook, NY. Nov 2013.
Overview of Big Data in Computer Science Department.

138. **New Challenges in Scheduling Theory.** Aussois, France. Mar 2014.
Reallocation Problems in Scheduling.

139. **XLDB Healthcare Workshop.** Stony Brook, NY. May 2014.
Indexing Big Data.

140. **Symposium on Principles of Database Systems (PODS).** Snowbird, UT. June 2014.
Cost-Oblivious Storage Reallocation.

141. **Workshop on Randomized Algorithms for Distributed Computing and Networks (RADICON)** Rennes, France TBD: **T**hree **B**ackoff **D**ilemmas. July 2014.

142. **10th Workshop on the Foundations of Mobile Computing (FOMC).** Philadelphia, PA. Aug 2014.
Keynote Talk.

143. **Workshop on Streaming Graph Algorithms (WSGA).** Albuquerque, NM. Oct 2014.
Indexing Big Data

144. **BigData Techcon.** San Francisco, CA. Oct 2014.
Tutorial. *Data Structures and Algorithms for Big Databases.*

145. **Lawrence Berkeley Laboratories.** Berkeley, CA. Feb 2015.
Indexing Big Data

146. **Workshop on Algorithms and Data Structures (ADS).** Bertinoro, Italy. June 2015.
Three Backoff Dilemmas.

147. **NoSQL Now.** San Jose, CA. Aug 2015.
Tutorial. *Write-Optimized Data Structures for Databases.*

148. **Georgia Tech.** Atlanta, GA. Sep 2015.
Write-Optimization.

149. **U Pennsylvania.** Philadelphia, PA. Nov 2015.
TBA: Three Backoff Dilemmas.

150. **U Connecticut.** Storrs, CT March 2016.
Write-Optimization.

151. **New Challenges in Scheduling Theory.** Centre CNRS "Paul Langevin," Aussois, France March 2016.
TBD: Three Backoff Dilemmas.

152. **Escuela de Ciencias Inform'aticas.** U Buenos Aires July 2016.
Week-long course. *Estructuras de Datos para los Sistemas de Almacenaje (Data structures for Storage Systems).*

153. **Complexity & Analysis of Distributed Algorithms.** CMO-BIRS Oaxaca, MX Nov 2016.
Three Backoff Dilemmas.

154. **Bertinoro Workshop on Algorithms and Data Structures (ADS).** Bertinoro, Italy. June 2017.
Anti-Persistence on Persistent Storage:? History-Independent Sparse Tables and Dictionaries.

155. **Faculty seminar and lunch.** Stony Brook, NY January 2017.
Write Optimization.

156. **New Challenges in Scheduling Theory.** Centre CNRS "Paul Langevin," Aussois, France April 2018.
Scheduling Algorithms for Program-Committee Members.

157. **33rd IEEE IPDPS.** Vancouver, CA May 2018.
The Algorithmics of Write Optimization.
Keynote talk.

Selected Department and University Service

1. Member, Graduate Admissions Committee, 1998-2002, 2006-2007
2. Member, School of Engineering Scholarship Committee, 2001
3. Organizer, Theory Qualifier, 1999, 2003
4. Host, Distinguished Lecturers (Bentley, Farach-Colton, Hart, Karger, Leiserson, Valiant, Goldwasser, Muthukrishnan), 2001, 2002, 2003, 2008, 2013, 2013, 2014.
5. CS Representative, Strategic Plan Advisory and Coordinating Committee (SPACC), 2002
6. Chair, Visibility/PR Committee, 2003-2004
7. Judge, Graduate Student Research Conference, 2003
8. Member, Quality of Life Committee, 2005-2008
9. CS Representative, Women in Science and Engineering (WISE), Course Counseling, 2002, 2005-present
10. Member, Undergraduate committee, 2005-2008
11. Director, Undergraduate honors program, 2005-2008
12. Member, Seawulf Cluster Management Committee, 2006-2008
13. Member, Recruiting Committee, 2010-2011
14. Member, Ad-hoc Awards Committee, 2012.
15. Co-director, Undergraduate CS Honors Program, 2010-2016.
16. Director, Undergraduate CS Honors Program, 2016-present.

Teaching at Stony Brook

1. **CSE 150 Honors Foundations of Computer Science**, Fall 2005, 2006, 2010, 2012, 2014, 2016.
   Approximately 10-40 students per semester.

2. **CSE 215 Foundations of Computer Science**, Spring 2015.
   Approximately 130 students per semester.

3. **CSE 303 Introduction to the Theory of Computation**, Spring 2003.
   Approximately 100 students per semester.

4. **CSE 350 Theory of Computation: Honors**, Spring 2006, 2007, 2014, 2016, 2018.
   Approximately 10-20 students per semester.

5. **CSE 373 Undergraduate Analysis of Algorithms**, Spring 1999, 2001, 2002, 2005, 2017.
   Approximately 80-110 students per semester.

6. **CSE 548 Graduate Analysis of Algorithms**, Fall 1998–2002, 2010, 2011, 2013, 2018. Spring 2012, 2013.
   Approximately 60–150 students per semester.

7. **CSE 638 Advanced Algorithms and Data Structures**, Spring 2005, 2007, 2011, 2012.
   Approximately 25 students per semester.

8. **CSE 642 Seminar in the Analysis of Algorithms**, Spring 2005–2007, Fall 2010–present.
   Approximately 20-40 students and faculty per semester.

9. **CSE 648 Randomized Algorithms and Advanced Data Structures**, Spring 2000, 2002, 2003.
   Approximately 30 students per semester.

10. **CSE 652 Seminar in the Analysis of Algorithms**, Fall 1999–Fall 2004.
    Approximately 20-30 students and faculty per semester.

Other Teaching

1. **Advanced Topics in Data Structures** (in Spanish), July 1999.
   18 hour course. Escuela de Ciencias Informáticas, U Buenos Aires, Argentina.
   Approximately 110 students.

2. **MIT 6.895 Theory of Parallel Systems**, Fall 2003.
   Approximately 20 students.

3. **MIT 6.896 Theory of Parallel Hardware**, Spring 2004.
   Approximately 20 students.

4. **Algorithms for External Memory** (in Spanish), July 2005.
   18 hour course. Escuela de Ciencias Informáticas, U Buenos Aires, Argentina.
   Approximately 40 students.

5. **Theoretical Foundations of Computer Science** (in Spanish), July 2005.
   8 hour workshop. Instituto Tecnológico de Colima, Mexico.
   Approximately 10 students.

6. **Algorithms for External Memory** (in Spanish), Sept 2006.
   5 hour tutorial. 7imo Encuentro Internacional de Ciencias de la Computación (ENC '06), San Luis Potosí, MX.
   Approximately 30 students.

7. **Databases and External Memory: Indexes, Write Optimization, and Cryptosearches**, May 2012.
   STOC Tutorial on Algorithms for Memory-Sensitive Computing. New York, NY. Approximately 40 attendees.

8. **Data Structures and Algorithms for Big Databases**, Sept 2012.
   4 hour tutorial. 6th Extremely Large Databases Conference, Workshop, and Tutorials (XLDB). Stanford, CA.
   Approximately 50 attendees.

9. **Data Structures and Algorithms for Big Databases**, Sept 2013.
   4 hour tutorial. 7th Extremely Large Databases Conference, Workshop, and Tutorials (XLDB). Stanford, CA.
   Approximately 50 attendees.

10. **Data Structures for Big Data**, Jan 2015.
    Week-long course. Ecole Normale Supériere de Lyon, Lyon France.

11. **Estructuras de Datos para los Sistemas de Almacenaje (Data structures for Storage Systems)** Jul 2016.
    Week-long course. Escuela de Ciencias Inform'aticas. University of Buenos Aires. Approximately 70 students.

PhD Students Advised

1. Roozbeh Ebrahimi (Google)
2. Yonatan Fogel (Oracle)
3. Haodong Hu (Shanghai University of Finance and Economics), Fall 2007
4. Tien-Ruey Hsiang (National Taiwan University of Science and Technology), Fall 2002
   (co-advised with E. Arkin and J.S.B. Mitchell)
5. Nenad Jovanovic (First Spring Corp), Winter 2002
   (co-advised with E. Arkin and J.S.B. Mitchell)
6. Tyler Mayer (Charles River Analytics), Spring 2018. (co-advised J.S.B. Mitchell)
7. Samuel McCauley (IT University of Copenhagen), Summer 2016
8. Dzejla Medjedovic (Sarajevo School of Science and Technology), Summer 2014
9. Pablo Montes (Google), Summer 2014
10. Prashant Pandey
11. Shikha Singh (Wellesley College), Summer 2018. (co-advised with J. Chen)

12. Marcelo Sztainberg (Northeastern Illinois University), Fall 2003
    (co-advised with Arkin and Mitchell)