

Martín Farach-Colton

Employment History

- 7/05 - **Professor**, Rutgers University Computer Science Department.
Graduate Director, 7/03 - 6/06, 7/12 - 6/14
- 12/06 - 4/14 **Founder & CTO**, Tokutek
Acquired by Percona, Inc.
- 7/98 - 6/05 **Associate Professor**, Rutgers University Computer Science Department.
- 7/00 - 7/02 **Senior Research Scientist**, Google, Inc.
- 7/97 - 8/98 **Member of Technical Staff**, Bell Laboratories, Lucent Technologies.
- 9/94 - 6/98 **Assistant Professor**, Rutgers University Computer Science Department.
- 9/91 - 8/94 **Postdoctoral Fellow**, DIMACS: Center for Discrete Mathematics and Theoretical Computer Science.

Education

- Ph.D.** 1991, Computer Science, University of Maryland, College Park.
- M.D.** 1988, The Johns Hopkins University School of Medicine, Baltimore, Maryland.
- B.S.** 1984, Mathematics and Chemistry, *magna cum lauda*, with Honors from South Carolina College, University of South Carolina, Columbia.

Selected Awards and Honors

- 2013 Distinguished Alumnus, 40th Anniversary of the University of Maryland Computer Science Department.
- 2012 Rutgers Computer Science Graduate Student Society *Award for Excellence in Teaching*.
- 2012 Imre Simon “Test of Time” Award for LATIN 2000 paper “The LCA Problem Revisited.”
- 2005 Plenary Speaker on the Future of Algorithms at the *New Horizons in Computing* Conference in Kyoto, Japan.
- 1999 Rutgers Computer Science Graduate Student Society *Award for Excellence in Teaching*.
- 1996 Alfred P. Sloan Research Fellowship.
- 1994, 96 National Academy of Sciences “Frontiers of Science” Symposium invited participant, Irvine CA.

Editorial Boards

- 1/04 – **Editorial Board Member**, *ACM Transactions on Algorithms*
- 7/94 – 1/04 **Editorial Board Member**, *Journal of Algorithms*.
- 9/92 - 3/94 **Associate Editor** for Computation and Engineering in Genomics, *IEEE Engineering in Medicine and Biology*.

Non-academic Activities

- 2015 Bronze Medal, IBJJF World Master Brazillian Jiu Jitsu Championship, Master 5, Purple, Middle-heavyweight.
- 2013 – Ali Forney Center, Board of Directors.
- 2008 – 14 Lambda Legal, Board of Directors.
Development Committee Co-Chair 08-12.
Board Governance Committee Co-Chair 13-14.

Selected Professional Activities

- 8/99 – 9/03 **Chair**, DIMACS Special Activities on Bioinformatics, Including the DIMACS Special Year held in '00-'03.
- 10/97 – 10/99 **Commmittee Member**, National Academy of Science, National Research Council committee on *Strengthening the Linkages between the Sciences and Mathematics*.
- 1/97 – 8/99 **Member**, Steering Committee for DIMACS Special Year in Massive Data Sets.
- 2/96 – 6/06 **Member**, DIMACS Council, advisory board of directors.
- 9/93 – 8/96 **Assistant Chair**, DIMACS Special Year on “Mathematical Support for Molecular Biology.”

Recent Program Committees

- 6/16 Combinatorial Pattern Matching (CPM '16).
- 4/16 12th Latin American Theoretical INformatics Conference (LATIN '16)
- 6/14 Combinatorial Pattern Matching (CPM '14).
- 5/14 IEEE/ASE International Conf on Big Data Science and Computing (BigData '14)
- 5/13 IEEE/ASE International Conf on Big Data Science and Computing (BigData '13)
- 9/12 Workshop on Massive Data Algorithmics (MASSIVE '12).
- 6/12 Combinatorial Pattern Matching (CPM '12).
- 4/12 10th Latin American Theoretical INformatics Conference (LATIN '12)
- 10/10 17th String Processing and Information Retrieval (SPIRE '10).
- 9/09 International Conference on Parallel Processing (ICPP-2009).
- 6/09 Combinatorial Pattern Matching (CPM '09).
- 10/07 14th String Processing and Information Retrieval (SPIRE '07).
- 9/07 Conference Chair for Encuentro Nacional de la Informatica (ENC '07) IR Track.
- 10/06 13th String Processing and Information Retrieval (SPIRE '06).
- 12/05 12th Annual International Conference on High Performance Computing (HiPC '05).
- 8/05 11th Annual International Computing and Combinatorics Conference (Cocoon '05).
- 6/05 13th Annual International conference on Intelligent Systems for Molecular Biology (ISMB 2005).

- 6/05 Combinatorial Pattern Matching (CPM '05).
- 5/05 37th Annual ACM Symposium on the Theory of Computing (STOC '05).
- 6/04 Third International Conference on FUN with Algorithms (FUN '04).
- 4/04 Program Committee Chair for Sixth Latin American Theoretical INformatics Conference (LATIN '04)
- 1/03 Program Committee Chair for Symposium on Discrete Algorithms (SODA '03)

Supervisory Activities

- Postdoctoral Fellows: Richa Agarwala, 94–96, now at NIH; S. Muthukrishnan, 94–95, now on the faculty of Rutgers University and at Microsoft Labs; R. Ravi, 94–95, now on the faculty of Carnegie Mellon University; Vineet Bafna, 94–96, now at Celera; Scott Decatur, 96–97, now at GMO, Inc.; Mona Singh, 96–97, now a faculty member at Princeton University; Shibu Yooseph, 97–98, now at Celera; Elizabeth Sweedyk, 97–99, now a faculty member at Harvey Mudd Collge.
- Ph.D. Students: Current: Meng-Tsung Tsai, Alex Conway. Completed Ph.D.: Rohan Fernades, May 2007, now a analyst at J.P. Morgan, New York; Yang Huang, May 2007, now a postdoc at NIH; Miguel Mosteiro, June 2006, now a postdoctoral researcher at Universidad Rey Juan Carlos, Madrid; Vicky Choi, November 2001, now a faculty member at Virginia Tech; Rahul Shah, November 2001, now a faculty member at Louisiana State University; Gabriela Hristescu, August 2000, now a faculty member at Rowan College; Rick Desper, June 1998, now a researcher at the National Human Genome Research Institute (NIH); Vincenzo Liberatore, June 1998, now a faculty member at Case Western Reserve University.

Grants

- 2016 PI for ENC Grant *High-Performance File Systems*.
- 2015-2018 co-PI for HHH Grant *Meaningful Data Compression and Reduction of High-Throughput Sequencing Data*, PI Alexander Schliep.
- 2014-2017 PI for NSF Grant *CSR: FTFS: A Read/Write-optimized Fractal Tree File System*.
- 2013-2016 PI for NSF Grant IIS-1247750 *BIGDATA: Eliminating the Data Ingestion Bottleneck in Big-Data Applications*.
- 2011-2014 PI for NSF Grant CCF-1114930 *Algorithms for Reallocation Problems*.
- 2011-2013 Co-PI for NSF SBIR Grant *A Multithreaded Storage Engine using Highly-Concurrent Fractal Trees*.
- 2009-2012 PI for NSF Grant CCF-0937829 *HECURA: Multidimensional and String Indexes for Streaming Data*.
- 2008-2011 PI for DOE Small Business Grant *Efficient Databases for HPC Clusters*.
- 2008 PI for NSF SBIR Grant *A Storage Engine for High-Volumn Data*.
- 2006-09 PI for NSF Grant CCF-0632838 *Adversarial Contention Resolution*.

- 2006-09 PI for NSF Grant CCF-0621425 *Techniques for Streaming File Systems and Databases*.
- 2006-08 PI for NSF Grant CCF-0541097 *High-Performance Data Access through Memory Abstraction*.
- 1999-03 PI for NSF Grant CCR-9820879 *Algorithms for Embedding Metric Spaces*.
- 1995-98 PI for NSF Career Development Award CCR-9501942.
- 1996 PI for NATO Grant CRG 960215.
- 1995 PI for NSF Research Experience for Undergraduates Award.
- 1994-95 Co-PI with Michael Noordewier for NSF Grant BIR-9403650, *Computational Methods for Analyzing Genetic Regulatory Sequences*.
- 1900-03 PI for DIMACS Activities in Bioinformatics, Fred Roberts and Craig Nevill-Manning Co-PIs.
 - Sloan Foundation.
 - Celera Genomics.
- 1995-98 *Mathematical Support for Molecular Biology*.
 - NSF BIR 94-12594.
 - National Center for Human Genome Research/NIH (through NSF) BIR 94-12594-03-CONF.
 - National Security Agency Grant ID#95I-191.
 - Misc. small grants from Roche Molecular, Centers for Disease Control, National Institute for Allergies and Infectious Diseases, SmithKline Beecham, Merck Pharmaceutical, BioSym Tech, and Ciba Geigy, totalling approximately.
- 1994-96 PI for NSF Grant DMS-9407039, *Improved Algorithms for Functional Analysis of Genomic Information*.
- 1992-94 PI for NSF Grant CCR-9207931, *Applications of String Matching to Molecular Biology*.

Journal Articles (with Conference Versions)

- (1) Martin Farach-Colton and Meng-Tsung Tsai. Exact Sublinear Binomial Sampling. Accepted for publication to *Algorithmica*, Special Issue for ISAAC 2013.
- (2) 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: A Right-Optimized Write-Optimized File System. Accepted for publication to *ACM Transaction on Storage*.
- (3) M. Bender, M. Farach-Colton, S. Fekete, J. Fineman, and S. Gilbert. Reallocation Problems in Scheduling. *Algorithmica*, Vol 73:389–409, 2015.
- (4) M. Farach-Colton, A. Fernández-Anta and M. A. Mosteiro. Optimal Memory-aware Sensor Network Gossiping. Accepted for publication to *Theoretical Computer Science*.

- (5) P. Bille and M. Farach-Colton. Fast and Compact Regular Expression Matching. Accepted for publication to *Theoretical Computer Science*.
- (6) M. Farach-Colton, R. Fernandes and M. Mosteiro. Bootstrapping a Hop-optimal Network in the Weak Sensor Model. Accepted for publication to *TALG*.
- (7) M. Farach-Colton and M. Mosteiro. Initializing Sensor Networks of Non-uniform Density in the Weak Sensor Model. *Algorithmica*: Volume 73, Issue 1 (2015), Page 87-114.
- (8) N. Alon, M. Badoiu, E. Demaine, M. Farach-Colton, M. Hajiaghayi, and A. Sidiropoulos. Ordinal Embeddings of Minimum Relaxation: General Properties, Trees, and Ultrametrics. *ACM Transactions on Algorithms*, 4(4): 2008.
- (9) M. Farach-Colton, G. Landau, C. Sahinalp and D. Tsur. Optimal spaced seeds for faster approximate string matching. *Journal of Computer and System Science*, 73(7): 1035-1044, 2007.
- (10) M. Bender, M. Farach-Colton, and M. Mosteiro. Insertion Sort is $O(n \log n)$. *Theory of Computing Systems*, 39(3): 391-397, 2006. Special Issue on LATIN '04.
- (11) M. Bender, M. Farach-Colton, G. Pemmasani, S. Skiena, and P. Sumazin. All-Pairs Lowest Common Ancestors in Trees and Directed Acyclic Graphs. *Journal of Algorithms*, 57(2): 75-94, 2005.
- (12) M. Farach-Colton and R. Shah. On the Complexity of Ordinal Clustering. *Journal of Classification*, accepted for publication.
- (13) V. Choi and M. Farach-Colton. BARNACLE: An Assembly Algorithm for Clone-based Sequences of Whole Genomes. *Gene* 320 (2003) 165–176.
- (14) M. Bender and M. Farach-Colton. The Level Ancestor Problem Simplified. *Theoretical Computer Science*, in press (2003).
- (15) M. Bender, E. Demaine and M. Farach-Colton. Cache-Oblivious B-Trees. *SICOMP*, in press (2003).
- (16) Y. Bartal, M. Farach-Colton, S. Yooseph, and L. Zhang. Fast, Fair and Frugal Bandwidth Allocation in ATM Networks. *Algorithmica*, 33(3):272–286 (2002).
- (17) R. Cole, M. Farach-Colton, R. Hariharan, T. Przytycka and M. Throup. An $O(n \log n)$ Algorithm for the Maximum Agreement Subtree Problem for Binary Trees. *SIAM Journal of Computing*, 30(5):1385–1404 (2000).
- (18) M. Farach-Colton and V. Liberatore. On Local Register Allocation. *J. Algorithms* 37(1):37–65 (2000).
- (19) K. Chen, D. Durand, and M. Farach-Colton. Notung: Dating gene duplications using gene family trees. *Journal of Computational Biology*, 7(3-4):429–447 (2000).
- (20) M. Farach-Colton, P. Ferragina and S. Muthukrishnan. On the Sorting-Complexity of Suffix Tree Construction. *Journal of the Association for Computing Machinery*, 47(6):987–1011 (2000).
- (21) M. Farach and S. Kannan. Efficient Algorithms for Inverting Evolution. *Journal of the Association for Computing Machinery*, 46(4):437–449 (1999).
- (22) R. Agarwala, V. Bafna, M. Farach, M. Paterson and M. Thorup. On the Approximability of Numerical Taxonomy (Fitting Distances by Tree Metrics). *SIAM Journal on Computing*, 28(3):1073–1085 (1999).

- (23) A. Amir, G. Benson and M. Farach. Optimal Parallel Two Dimensional Text Searching on a CREW PRAM. *Information and Computation*, 144(1):1–17 (1998).
- (24) M. Farach and M. Thorup. String Matching in Lempel-Ziv Compressed Strings. *Algorithmica* 20(4):388–404 (1998).
- (25) M. Farach and S. Muthukrishnan. Optimal Parallel Randomized Renaming. *Information Processing Letters*, Vol 61, pp. 7–10, 1997.
- (26) M. Farach and M. Thorup. Sparse Dynamic Programming for Evolutionary Tree Comparison. *SIAM Journal on Computing*, Vol 26, No 1, pp. 210–230, 1997.
- (27) A. Amir, G. Benson and M. Farach. Optimal Two Dimensional Compressed Matching. *Journal of Algorithms*, Vol 24, pp. 354–379, 1997.
- (28) M. Farach. Recognizing Circular Decomposable Metrics. *Journal of Computational Biology*, Vol 4, pp. 157–162, 1997.
- (29) Richa Agarwala, Serafim Batzoglou, Vlado Dančik, Scott Decatur, Martin Farach, Sridhar Hannenhalli, S. Muthukrishnan and Steven Skiena. Local Rules for Protein Folding on the Triangular Lattice and Generalized Hydrophobicity in the HP Model. *Journal of Computational Biology*, 4(3):275–296 (1997).
- (30) Jaime Cohen and Martin Farach. Numerical Taxonomy on Data: Experimental Results. *Journal of Computational Biology*, 1997, pp 547-558.
- (31) A. Amir, G. Benson and M. Farach. Let Sleeping Files Lie: Pattern Matching in Z-compressed Files. *Journal of Computer and System Science*, Vol 52, No 2, pp. 299–307, 1996.
- (32) M. Farach, T. Przytycka and M. Thorup. On the Agreement of Many Trees. *Information Processing Letters*, Vol 55, pp. 297–301, 1995.
- (33) M. Farach and M. Thorup. Fast Comparison of Evolutionary Trees. *Information and Computation*, Vol 123, No. 1, pp.29–37, 1995.
- (34) A. Amir, M. Farach, R. Idury, H. La Poutré and A. Schäffer. Improved Dynamic Dictionary Matching. *Information and Computation*, Vol 119, No. 2, pp. 258–282, 1995.
- (35) A. Amir, M. Farach and S. Muthukrishnan. Alphabet Dependence in Parameterized Matching. *Information Processing Letters*, Vol 49, No. 3, pp. 111-120, 1995.
- (36) A. Amir and M. Farach. Efficient 2-dimensional Approximate Matching of Half-rectangular Figures. *Information and Computation*, Vol 118, No. 1, pp. 1-11, 1995.
- (37) M. Farach, S. Kannan and T. Warnow. A Robust Model for Finding Optimal Evolutionary Trees. *Algorithmica, Special Issue on Computational Biology*, Vol 13, pp. 155-179, 1995.
- (38) A. Amir, M. Farach, R. Giancarlo, Z. Galil and K. Park. Dynamic Dictionary Matching. *Journal of Computer and Systems Science*, Vol 49, No. 2, pp. 208-222, 1994.
- (39) A. Amir, G. Benson and M. Farach. An Alphabet Independent Approach to Two Dimensional Pattern Matching. *SIAM Journal on Computing*, Vol 23, No. 2, pp. 313-323, 1994.
- (40) A. Amir and M. Farach. Two Dimensional Dictionary Matching. *Information Processing Letters*, Vol 44, pp. 233-239, 1992.

- (41) A. Apostolico, M. Farach and C. Iliopoulos. Optimal Superprimitivity Testing for Strings. *Information Processing Letters*, Vol 39, pp. 17–20, 1991.
- (42) A. Amir and M. Farach. Efficient Matching of Nonrectangular Shapes. *Annals of Mathematics and Artificial Intelligence*, Vol 4, 1991.
- (43) J. Wald, M. Farach, M. Tagamets and J. Reggia. Generating Plausible Diagnostic Hypotheses with Self-processing Causal Networks. *Journal of Experimental and Theoretical Artificial Intelligence*, Vol 1. #2, 1989.

Conference Papers

- (44) Michael Bender, Rezaul Chowdhury, Alex Conway, Martin Farach-Colton, Pramod Ganapathi, Rob Johnson, Samuel McCauley, Bertrand Simon and Shikha Singh. The I/O Complexity of Computing Prime Tables. *Proceedings of LATIN '16*.
- (45) Martin Farach-Colton and Meng-Tsung Tsai. Tight Approximations of Graph Degeneracy in Large Graphs. *Proceedings of LATIN '16*.
- (46) Martin Farach-Colton and Meng-Tsung Tsai. On the Complexity of Computing Prime Tables. *Proceedings of ISAAC '15*.
- (47) Martin Farach-Colton, Tsan-sheng Hsu, Meng Li and Meng-Tsung Tsai. Finding Articulation Points of Large Graphs in Linear Time. *Proceedings of WADS '15*.
- (48) Michael A. Bender, Martin Farach-Colton, Sandor Fekete, Jeremy Fineman and Seth Gilbert. Cost-Oblivious Reallocation for Scheduling and Planning. *Proceedings of SPAA '15*.
- (49) 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: A Right-Optimized Write-Optimized File System. *13th USENIX Conference on File and Storage Technologies, (FAST '15)*.
- (50) Michael A. Bender, Martin Farach-Colton, Mayank Goswami, Dzejla Medjedovic, Pablo Montes, and Meng-Tsung Tsai. The Batched Predecessor Problem in External Memory. *European Symposium on Algorithms (ESA '14)*.
- (51) Martin Farach-Colton, Katia Leal, Miguel A. Mosteiro and Christopher Thraves. Dynamic Windows Scheduling with Reallocations. *Proceedings of SEA '14*.
- (52) Michael A. Bender, Martin Farach-Colton, Sandor Fekete, Jeremy Fineman, and Seth Gilbert. Cost-Oblivious Storage Reallocation. *Proceedings of PODS '14*.
- (53) Martin Farach-Colton and Meng-Tsung Tsai. Computing the Degeneracy of Large Graphs. *Proceedings of LATIN '14*.
- (54) Martin Farach-Colton and Meng-Tsung Tsai. Exact Sublinear Binomial Sampling. *Proceedings of ISAAC '13*.
- (55) Michael A. Bender, Martin Farach-Colton, Sandor Fekete, Jeremy Fineman, and Seth Gilbert. Reallocation Problems in Scheduling. *Proceedings of SPAA '13*.
- (56) 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. *Proceedings of VLDB '12*.

- (57) John Esmet, Michael A. Bender, Martin Farach-Colton, and Bradley C. Kuszmaul. The TokuFS File System. *Proceedings of HotStorage '12*.
- (58) M. Farach-Colton, A. Fernandez-Anta, A. Milani, M. A. Mosteiro and S. Zaks. Opportunistic Information Dissemination in Mobile Ad-hoc Networks: adaptiveness vs. obliviousness and randomization vs. determinism. *Proceedings of LATIN '12*.
- (59) Michael A. Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, Dzejla Medjedovic, Pablo Montes, Pradeep Shetty, Richard P. Spillane, Erez Zadok. Don't Thrash: How to Cache Your Hash on Flash. *Proceedings of HotStorage '11*.
- (60) P. S. Almeida, C. Baquero, M. Farach-Colton, P. Jesus and M. A. Mosteiro. Fault-Tolerant Aggregation: Flow Update Meets Mass Distribution. *Proceedings of the 15th International Conference On Principles Of Distributed Systems (OPODIS '11)*.
- (61) A. Mitrofanova, M. Farach-Colton, and B. Mishra. Efficient and Robust Prediction Algorithms for Protein Complexes using Gomory-Hu Tree. *Pacific Symposium on Biocomputing (PSB '09)*.
- (62) M. Farach-Colton and Y. Huang. A Linear Delay Algorithm for Building Concept Lattices. *Symposium on Combinatorial Pattern Matching (CPM '08)*.
- (63) M. Farach-Colton and M. Mosteiro. Sensor Network Gossiping or How to Break the Broadcast Lower Bound. *18th International Symposium on Algorithms and Computation (ISAAC '07)*.
- (64) M. Farach-Colton and M. Mosteiro. Initializing Sensor Networks of Non-uniform Density in the Weak Sensor Model. *Workshop on Data structures and Algorithms (WADS '07)*.
- (65) M. Bender, M. Farach-Colton, J. Fineman, Y. Fogel, B. Kuszmaul and J. Nelson. Cache-Oblivious Streaming B-Trees. *Symposium on Parallelism in Algorithms and Architectures (SPAA '07)*.
- (66) M. Farach-Colton and Y. Huang. Lattice based clustering of temporal gene-expression matrices. *7th SIAM International Conference on Data Mining (SDM'07)*.
- (67) M. Farach-Colton, R. J. Fernandes and M. A. Mosteiro. Lower Bounds for Clear Transmissions in Radio Networks. *Latin American Theoretical Computer Science (LATIN '06)*.
- (68) M. Bender, M. Farach-Colton and B. Kuszmaul. Cache-Oblivious String B-Trees. *Symposium on Principles of Database Systems (PODS '06)*.
- (69) M. Farach-Colton, R. J. Fernandes and M. A. Mosteiro. Bootstrapping a Hop-optimal Network in the Weak Sensor Model. *European Symposium on Algorithms (ESA '05)*.
- (70) M. Bender, M. Farach-Colton, S. He, B. Kuszmaul, and C. Leiserson. Adversarial Contention Resolution for Simple Channels. *Symposium on Parallelism in Algorithms and Architectures (SPAA '05)*.
- (71) N. Alon, M. Badoiu, E. Demaine, M. Farach-Colton, M. Hajiaghayi, and A. Sidiropoulos. Ordinal Embeddings of Minimum Relaxation: General Properties, Trees, and Ultrametrics. *Symposium on Discrete Algorithms (SODA '05)*.
- (72) M. Farach-Colton, Y. Huang and J. Woolford. Discovering Temporal Relations in Molecular Pathways Using Protein-Protein Interactions. *RECOMB '04*.

- (73) M. Bender, M. Farach-Colton and M. Mosteiro. Insertion Sort is $O(n \log n)$. *FUN '04*.
- (74) M. Bender and M. Farach-Colton. The Level Ancestor Problem Simplified. *LATIN '02*. See Paper (14) for journal version.
- (75) M. Bender, R. Cole, E. Demaine, M. Farach-Colton and J. Zito. Two Simplified Algorithms for Maintaining Order in a List. *European Symposium on Algorithms (ESA '02)* 152–164.
- (76) M. Bender, R. Cole, E. Demaine and M. Farach-Colton. Scanning and Traversing: Maintaining Data for Traversals in a Memory Hierarchy. *European Symposium on Algorithms (ESA '02)* 139–151.
- (77) M. Bender, E. Demaine and M. Farach-Colton. Efficient Tree Layout in a Multilevel Memory Hierarchy. *European Symposium on Algorithms (ESA '02)* 165–173.
- (78) M. Charikar, K. Chen and M. Farach-Colton. Finding Frequent Items in Data Streams. *International Colloquium on Automata, Languages, and Programming (ICALP '02)* 508–515.
- (79) R. Shah and M. Farach-Colton. Undiscretized Dynamic Programming: Faster Algorithms for Facility Location and Related Problems on Trees. *Symposium on Discrete Algorithms (SODA '02)* 108–115.
- (80) M. Farach-Colton and G. Hristescu. Structure-based Feature Extraction from Protein Databases. 2001 International Conference on Mathematics and Engineering Techniques in Medicine and Biological Sciences *METMBS '01*.
- (81) M. Farach-Colton and R. Shah. On the Midpath Tree Conjecture: A Counter-Example. *Symposium on Discrete Algorithms (SODA '01)* 208–209.
- (82) M. Farach-Colton and G. Hristescu. COFE: A Scalable Method for Feature Extraction from Complex Objects. 2nd International Conference on Data Warehousing and Knowledge Discovery *DaWaK '00* 358–371.
- (83) M. Bender and M. Farach-Colton. The LCA Problem Revisited. *LATIN '00* 88–94.
- (84) K. Chen, D. Durand, and M. Farach-Colton. Notung: Dating gene duplications using gene family trees. *RECOMB '00* 96–106. See Paper (19) for journal version.
- (85) M. Bender, E. Demaine and M. Farach-Colton. Cache-Oblivious B-Trees. *FOCS '00* 399–409. See Paper (15) for journal version.
- (86) M. Farach-Colton and P. Indyk. Approximate Nearest Neighbor Algorithms for Hausdorff Metrics via Embeddings. *Foundations of Computer Science (FOCS '99)* 171–180.
- (87) M. Farach-Colton, U. Kremer and V. Liberatore. Evaluation of Algorithms for Local Register Allocation. *International Conference on Compiler Construction (CC '99)* 137–152.
- (88) Y. Bartal, M. Farach-Colton, S. Yooseph, and L. Zhang. Fast, Fair and Frugal Bandwidth Allocation in ATM Networks. *Symposium on Discrete Algorithms (SODA '99)* 92–101. See Paper (16) for journal version.
- (89) M. Farach, P. Ferragina and S. Muthukrishnan. Overcoming Memory Bottlenecks in Suffix Tree Construction. *Foundations of Computer Science (FOCS '98)* 174–185. See Paper (20) for journal version.

- (90) M. Farach-Colton and V. Liberatore. On Local Register Allocation. *Symposium on Discrete Algorithms (SODA '98)* 564–573. See Paper (18) for journal version.
- (91) A. Ambainis, R. Desper, M. Farach and S. Kannan. Nearly Tight Bounds on the Learnability of Evolution. *Foundations of Computer Science (FOCS '97)* 524–533.
- (92) C. Womack and M. Farach. Randomization, Rigor and Persuasiveness in Proofs. *Logic and Mathematical Reasoning*, Mexico City, Mexico, October 1997.
- (93) G. Hristescu, C. Benham, and M. Farach. DNA Strand Separation Prediction: A Parallel Implementation. *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '97)*.
- (94) M. Farach, S. Kannan, E. Knill and S. Muthukrishnan. Group Testing Problems with Sequences in Experimental Molecular Biology. *SEQUENCES '97*, Possitano, Italy, June 1997.
- (95) M. Farach. Optimal Suffix Tree Construction with Large Alphabets. *Foundations of Computer Science (FOCS '97)* 137–143. See Paper (20) for journal version.
- (96) Jaime Cohen and Martin Farach. Numerical Taxonomy on Data: Experimental Results. *Symposium on Discrete Algorithms (SODA '97)*. See Paper (30) for journal version.
- (97) Richa Agarwala, Serafim Batzoglou, Vlado Dančik, Scott Decatur, Martin Farach, Sridhar Hannenhalli, S. Muthukrishnan and Steven Skiena. Local Rules for Protein Folding on the Triangular Lattice and Generalized Hydrophobicity in the HP Model. *Symposium on Discrete Algorithms (SODA '97)* 390–399. See Paper (29) for journal version.
- (98) M. Farach and S. Muthukrishnan. Perfect Hashing for Strings: Formalization, Algorithms and Open Problems. *Symposium on Combinatorial Pattern Matching (CPM '96)*, Laguna Beach, California, July 1996.
- (99) G. Christopher, M. Farach and M. Trick. The Structure of Circular Decomposable Metrics. *European Symposium on Algorithms (ESA '96)*. See Paper (28) for journal version.
- (100) M. Farach and S. Kannan. Efficient Algorithms for Inverting Evolution. *Symposium on the Theory of Computing (STOC '96)*. See Paper (21) for journal version.
- (101) M. Farach and S. Muthukrishnan. Optimal Logarithmic Time Randomized Suffix Tree Construction. *International Colloquium on Automata Languages and Programming (ICALP '96)*. See Paper (20) for journal version.
- (102) M. Farach, T. Przytycka and M. Thorup. Computing the Agreement of Trees with Bounded Degrees. *European Symposium on Algorithms (ESA '95)* 381–393. See Paper (32) for journal version.
- (103) M. Farach and S. Muthukrishnan. Optimal Parallel Dictionary Matching and Compression. *Symposium on Parallel Algorithms and Architectures (SPAA '95)* pp. 244–253, Santa Barbara, California, July 1995.
- (104) M. Farach, M. Noordewier, S. Savari, L. Shepp, A. Weiner and J. Ziv. On the Entropy of DNA: Algorithms and Measurements based on Memory and Rapid Convergence. *Symposium on Discrete Algorithms (SODA '95)* pp. 48–57, San Francisco, January 1995.

- (105) M. Farach and M. Thorup. String Matching in Lempel-Ziv Compressed Strings. *Symposium on the Theory of Computing (STOC '95)*, 703–712. See Paper (24) for journal version.
- (106) M. Gu., M. Farach and R. Beigel. Dynamic On-line Text Indexing. *Symposium on Discrete Algorithms (SODA '94)* pp. 697–704, Washington, DC, January 1994.
- (107) M. Farach and M. Thorup. Optimal Evolutionary Tree Comparison by Sparse Dynamic Programming. *Foundations of Computer Science (FOCS '94)* 770–779. See Paper (26) for journal version.
- (108) A. Amir, G. Benson and M. Farach. Let Sleeping Files Lie: Pattern Matching in Z-compressed Files. *Symposium on Discrete Algorithms (SODA '94)* 705–714. See Paper (31) for journal version.
- (109) M. Farach and M. Thorup. Fast Comparison of Evolutionary Trees. *Symposium on Discrete Algorithms (SODA '94)* 481–488. See Paper (33) for journal version.
- (110) A. Amir, G. Benson and M. Farach. Optimal Two Dimensional Compressed Matching. *International Colloquium on Automata Languages and Programming (ICALP '94)*. See Paper (27) for journal version.
- (111) M. Farach, S. Kannan and T. Warnow. A Robust Model for Finding Optimal Evolutionary Trees. *Symposium on the Theory of Computing (STOC '93)* 137–145. See Paper (37) for journal version.
- (112) A. Amir, G. Benson and M. Farach. Parallel Two Dimensional Pattern Matching. *Symposium on Parallel Algorithms and Architectures (SPAA '93)*, 79–85. See Paper (23) for journal version.
- (113) A. Amir, M. Farach, R. Idury, H. La Poutré and A. Schäffer. Improved Dynamic Dictionary Matching. *Symposium on Discrete Algorithms (SODA '93)* 392–401. See Paper (34) for journal version.
- (114) A. Amir, M. Farach and Y. Matias. Efficient Randomized Dictionary Matching Algorithms. *Symposium on Combinatorial Pattern Matching (CPM '92)*, Tucson, Arizona, April 1992.
- (115) A. Amir, G. Benson and M. Farach. Alphabet Independent Two Dimensional Matching. *Symposium on the Theory of Computing (STOC '92)* 59–68. See Paper (39) for journal version.
- (116) A. Amir and M. Farach. Adaptive Dictionary Matching. *Foundations of Computer Science (FOCS '91)* 760–761. See Paper (38) for journal version.
- (117) A. Amir and M. Farach. Efficient 2-dimensional Approximate Matching of Non-rectangular Figures. *Symposium on Discrete Algorithms (SODA '91)* 212–223. See Paper (36) for journal version.
- (118) A. Amir and M. Farach. The Whole is Faster than its Parts. *Conference on the Foundations of Artificial Intelligence (FAI '89)*. See Paper (42) for journal version.

Books and Chapters

- (1) D. Durand and M. Farach-Colton, On the Design of Optimization Criteria for Multiple Sequence Alignment. In *Biological Evolution and Statistical Physics*, Springer Verlag. Michael Lässig and Angelo Valleriani, editors.

- (2) M. Farach, F. Roberts, M. Vingron and M. Waterman, Editors. *Mathematical Support for Molecular Biology*, DIMACS Book Series.
- (3) A. Amir and M. Farach, Extensions of String Matching and Array Pattern Matching. In *Pattern Matching Algorithms and Applications*, A. Apostolico and Z. Galil, Editors.