

Scott F. Kaplan

Department of Mathematics and Computer Science 155 Farmington Rd.
Amherst College Amherst, MA 01002
Amherst, MA 01002-5000 (413) 241-7345
Phone: (413) 542-2377
Email: sfkaplan@amherst.edu
URI: <https://sfkaplan.people.amherst.edu>

Research Interests

Virtual memory management and file system caching, CPU/memory/disk scheduling and allocation, garbage collection and manual memory allocation, reference trace collection, reference trace reduction, prefetching, page replacement policies, and in-memory data compression

Degrees

Doctor of Philosophy in Computer Sciences, 1999
The University of Texas at Austin
Advisors: Paul R. Wilson and Donald Fussell
Dissertation: *Compressed Caching and Modern Virtual Memory Simulation*

Bachelor of Arts, 1995
Amherst College
Magna cum laude in Computer Science

Professional Experience

Amherst College, Amherst, MA. Professor of Computer Science, July 2012 to present. Chair, Department of Computer Science, July 2014 to June 2015, October 2019 to present.

University of Massachusetts, Amherst, MA. Adjunct Assistant Professor of Computer Science, September 2004 to present.

Hampshire College, Amherst, MA. Adjunct Associate Professor of Computer Science, January 2012 to May 2012.

Amherst College, Amherst, MA. Associate Professor of Computer Science, July 2006 to June 2012. Chair, Department of Computer Science, Fall 2008 to Spring 2009, Fall 2011 to Spring 2012.

Microsoft Research, Redmond, WA. Visiting Researcher, January 2011 to March 2011.

University of Washington, Seattle, WA. Visiting Scholar, Department of Computer Science and Engineering, March 2011 to June 2011.

Universidad Técnica Federico Santa María, Valparaíso, Chile. Fulbright Scholar, Spring 2007.

Amherst College, Amherst, MA. Assistant Professor of Computer Science, Fall 1999 to Spring 2006.

IBM TJ Watson Research Center, Yorktown Heights, NY. Visiting Researcher, July 2002 to July 2003.

IBM TJ Watson Research Center, Hawthorne, NY. Visiting Researcher, Summer 2000.

The University of Texas at Austin, Department of Computer Science, Austin, TX. Graduate Research Assistant, Spring 1996 to Fall 1998.

Have OOP, Will Travel, Austin, TX. Teaching Assistant, Fall 1996 to Fall 1998.

The University of Texas at Austin, Department of Computer Science, Austin, TX. Assistant Instructor, CS 105 (Introduction to C++), Spring 1997.

The Santa Fe Institute, Santa Fe, NM. Intern Researcher, Summer 1994.

IBM TJ Watson Research Center, Yorktown Heights, NY. Intern Developer, Summer 1993.

Awards and Honors

National Science Foundation, *Interdisciplinary Cluster Computing at a Liberal Arts College*. (Proposal #0521169; \$298,839; September 2005 through August 2007)

National Science Foundation, *Acquisition of a Laboratory Testbed for Networked Embedded Systems and Sensor Research*. (Proposal #0520729; \$300,000; September 2005 through August 2007)

National Science Foundation, *Collaborative Research: A Five-College Partnership for Information Assurance Education*. (Proposal #0416851; \$299,986; September 2004 through August 2006)

Faculty Research Award Program: Amherst College, January 2004

The Miner D. Crary Sabbatical Fellowship: Amherst College, July 2002

Best Paper: The USENIX Annual Technical Conference, June 1999, for, "The Case for Compressed Caching in Virtual Memory Systems".

The Schlumberger Fellowship: The University of Texas, Department of Computer Sciences and The Schlumberger Corporation, Fall 1998

The MCD Fellowship: The University of Texas, Department of Computer Sciences, Fall 1995 to Spring 1997

Conference Papers

"Redline: First Class Support for Interactivity in Commodity Operating Systems," Ting Yang, Tongping Liu, Emery D. Berger, Scott F. Kaplan, and J. Eliot B. Moss. *Proceedings of The Eighth USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, December 2008, USENIX Association.

- “CRAMM: Virtual Memory Support for Garbage-Collected Applications,” Ting Yang, Emery D. Berger, Scott F. Kaplan, and J. Eliot B. Moss. *Proceedings of The Seventh USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, pp. 103–116, November 2006, USENIX Association.
- “Autonomic Heap Sizing: Taking Real Memory Into Account,” Ting Yang, Matthew Hertz, Emery D. Berger, Scott F. Kaplan, and J. Eliot B. Moss. *Proceedings of The Fourth International Symposium on Memory Management*, pp. 61–72, October 2004, ACM Press.
- “Adaptive Caching for Demand Prepaging,” Scott F. Kaplan, Lyle A. McGeoch, and Megan F. Cole. *Proceedings of The Third International Symposium on Memory Management*, pp. 114–126, June 2002, ACM Press.
- “The Case for Compressed Caching in Virtual Memory,” Paul R. Wilson, Scott F. Kaplan, and Yannis Smaragdakis. *Proceedings of The USENIX Annual Technical Conference*, pp. 101–116, June 1999, USENIX Association.
- “Trace Reduction for Virtual Memory Simulation,” Scott F. Kaplan, Yannis Smaragdakis, and Paul R. Wilson. *Proceedings of The ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pp. 47–58, May 1999, ACM Press.
- “EELRU: Simple and Effective Adaptive Page Replacement,” Yannis Smaragdakis, Scott F. Kaplan, and Paul R. Wilson. *Proceedings of The ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pp. 122–133, May 1999, ACM Press.

Journal Papers

- “Flexible Reference Trace Reduction for VM Simulations,” Scott F. Kaplan, Yannis Smaragdakis, and Paul R. Wilson. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, pp. 1–38, January 2003, ACM Press.
- “The EELRU Adaptive Replacement Algorithm,” Yannis Smaragdakis, Scott F. Kaplan, and Paul R. Wilson. *Performance Evaluation*, pp. 93–123, May 2003, Elsevier Science Publishers.

Books and book chapters

- “Adaptive Replacement Algorithm Templates and EELRU,” in *The Handbook of Research on Advanced Operating Systems and Kernel Applications: Techniques and Technologies*, 2009.

Workshop Papers (refereed)

- “Collecting Whole-System Reference Traces of Multiprogrammed and Multithreaded Workloads,” Scott F. Kaplan, *Proceedings of The ACM Workshop on Software Performance (WOSP)*, pp. 228–237, January 2004, ACM Press.
- “Class-based Prioritized Resource Control in Linux,” Shailabh Nagar, Hubertus Franke, Jonghyuk Choi, Chandra Seetharaman, Scott F. Kaplan, Nivedita Singhvi, Vivek Kashyap, and Mike

Kravetz. *Proceedings of The Ottawa Linux Symposium (OLS)*, pp. 161–180, July 2003, Linux Symposium.

Posters (refereed)

“Scheduler-Aware Virtual Memory Management,” Emery Berger, Scott Kaplan, Bhuvan Urgaonkar, Pritesh Sharma, Abhishek Chandra, Prashant Shenoy. Presented at *The Symposium on Operating System Principles (SOSP)*, October 2003, Bolton Landing, New York.

Extended Abstracts (refereed)

“Two Methods for Collecting Whole-System Reference Traces of Multiprogrammed and Multi-threaded Workloads,” Scott F. Kaplan. *Proceedings of The ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pp. 420–421, June 2004, ACM Press.

Technical Reports

“NoFuS: Automatically Detecting” + `String.fromCharCode(32)` + `ObFuSCateD.toLowerCase()` + `JavaScript Code`,” Scott Kaplan, Benjamin Livshits, Benjamin Zorn, Christian Siefert, Charlie Curstinger. *Microsoft Research, Technical Report MSR-TR-2011-57*, April 2011.

“kVMTrace: An Efficient Reference Trace Collector of Full Workloads for Main Memory Simulations,” Scott F. Kaplan. *Amherst College, Department of Mathematics and Computer Science, Technical Report TR-2004-02*, November 2004.

“Re-evaluating Page Replacement Policies: Why LRU Remains King,” Scott F. Kaplan, Thomas B. Jablin. *Amherst College, Department of Mathematics and Computer Science, Technical Report TR-2004-01*, October 2004.

“Autonomic Heap Sizing: Taking Real Memory into Account,” Ting Yang, Matthew Hertz, Emery Berger, Scott F. Kaplan, and J. Eliot B. Moss. *The University of Massachusetts Amherst, Department of Computer Science, Technical Report TR-04-14*, March 2004.

“EELRU: Simple and Effective Adaptive Page Replacement,” Yannis Smaragdakis, Scott F. Kaplan, and Paul R. Wilson. *The University of Texas at Austin, Department of Computer Sciences, Technical Report TR-98-23*, December 1998.

“Trace Reduction for Virtual Memory Simulations,” Scott F. Kaplan, Yannis Smaragdakis, and Paul R. Wilson. *The University of Texas at Austin, Department of Computer Sciences, Technical Report TR-98-24*, December 1998.

Invited Talks

“Adaptive Prepaging Using Recency Information,” International Business Machines, TJ Watson Research Center, February 2002.

“How Virtual Memory Systems Can Respond to Reference Behavior,” The University of Massachusetts, Amherst, Computer Science Department, January 2001.

“Recency and Compressed Caching: How to Respond to Reference Behavior,” Lucent Technologies, Bell Labs, January 2000.

“Recency and Compressed Caching: How to Respond to Reference Behavior,” International Business Machines, TJ Watson Research Laboratories, January 2000.

Professional Activities

Referee, *Journal of Systems and Software*, Summer 2011.

Referee, *Journal of Systems and Software*, Spring 2011.

Referee, *File And Storage Technologies (FAST)*, 2009.

Panelist, *National Science Foundation, Science and Technology Centers program*, January 2009.

Panelist, *National Science Foundation, Computer Systems Research*, June 2008.

Referee, *IEEE Transactions on Computers*, 2007.

Panelist, *National Science Foundation, Computer Systems Research*, March 2006.

Panelist, *National Science Foundation, Foundations of Computing Processes and Artifacts*, November 2005.

Referee, *Performance Evaluation Review*, Fall 2004.

Program Committee, *USENIX Annual Technical Conference*, 2004.

Referee, *USENIX Annual Technical Conference*, 2002.

Referee, *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 2001.

Publicity co-chair (North America), *ACM SIGMETRICS/PERFORMANCE*, 2001.

Referee, *Software: Practice and Experience*, 1999.

Member, Association for Computing Machinery.

Member, ACM Special Interest Group on Operating Systems.

Member, ACM Special Interest Group on Measurement and Evaluation.

Member, USENIX Advanced Computing Systems Association.

Classes Taught at Amherst College

Fall 2021:	CS 163: Introduction to Computer Architecture
	CS 171 : Computer Systems
Spring 2021:	CS 211: Data Structures
	CS 365: Performance Evaluationa and Optimization
Fall 2020:	=CS 171: Computer Systems
Spring 2020:	CS 112: Introduction to Computer Science II
	CS 283: Networks
Fall 2019:	CS 163: Introduction to Computer Architecture
	CS 171: Computer Systems
	CS 211: Data Structures

Spring 2019: CS 111: Introduction to Computer Science I
CS 365: Performance Evaluation and Optimization

Fall 2018: CS 171: Computer Systems
CS 283: Networks

Fall 2017: CS 111: Introduction to Computer Science I
CS 171: Computer Systems

Spring 2017: CS 112: Introduction to Computer Science II

Fall 2016: CS 112: Introduction to Computer Science II
CS 461: Advanced Operating Systems

Spring 2016: CS 111: Introduction to Computer Science I
CS 261: Computer Systems II

Spring 2015: CS 261: Computer Systems II
CS 281: Networks and Cryptography

Fall 2014: CS 111: Introduction to Computer Science I
CS 161: Computer Systems I

Spring 2014: CS 111: Introduction to Computer Science I
CS 261: Computer Systems II

Fall 2013: CS 111: Introduction to Computer Science I
CS 161: Computer Systems I

Spring 2013: CS 111: Introduction to Computer Science I
CS 281: Networks and Cryptography

Fall 2012: CS 111: Introduction to Computer Science I
CS 161: Computer Systems I

Spring 2012: CS 261: Computer Systems II
CS 461: Advanced Operating Systems
Hampshire college CS 156: Computer Systems

Fall 2011: First Year Seminar: Science and Religion
CS 161: Computer Systems I

Fall 2010: CS 16: Computer Systems I
CS 28: Networks and Cryptography

Spring 2010: CS 11: Introduction to Computer Science
CS 26: Computer Systems II

Fall 2009: CS 16: Computer Systems I
CS 39: Advanced Operating Systems (*new course*)

Spring 2009: CS 26: Computer Systems II (*new course*)
CS 28: Networks and Cryptography (*new course*)

Fall 2008: First Year Seminar: Science and Religion
CS 16: Computer Systems I (*new course*)

Spring 2008: CS 11: Introduction to Computer Science I
CS 29: Networks

Fall 2007: First Year Seminar: Science and Religion (*new course*)
CS 12: Introduction to Computer Science II

Spring 2006:	CS 14, Introduction to Computer Systems (lectures and labs)
Fall 2005:	CS 11, Introduction to Computer Science I CS 39, Principles of Operating System Design
Spring 2005:	CS 14, Introduction to Computer Systems (lectures and labs)
Fall 2004:	CS 11, Introduction to Computer Science I CS 29, Networks (<i>new course</i>)
Spring 2004:	CS 12, Introduction to Computer Science II CS 40, Seminar in Computer Science: Runtime Resource Management
Fall 2003:	CS 12, Introduction to Computer Science II (<i>new course</i>) CS 39, Principles of Operating System Design
Spring 2002:	CS 14, Introduction to Computer Systems (lectures and labs)
Fall 2001:	CS 29, Networks and Concurrent Computing (<i>new course</i>) CS 39, Principles of Operating System Design
Spring 2001:	CS 14, Introduction to Computer Systems (lectures and labs)
Fall 2000:	CS 11, Introduction to Computer Science (lectures and labs) CS 39, Principles of Operating System Design
Spring 2000:	CS 11, Introduction to Computer Science (lectures and labs)
Fall 1999:	CS 11, Introduction to Computer Science (lectures and labs) (co-taught with C. McGeoch) CS 40, Seminar in Computer Science: Virtual Memory and Data Representations

Advising of Honors Theses at Amherst College

Luka Duranovic 2022. *Untitled*.

Enes Kristo 2022. *Untitled*.

Christopher Rabasa 2020. *Whole-system reference trace gathering and modeling of elastic process memory allocation*.

Jeffrey Ewing 2018. *Redesigned Compressed Caching for Modern Systems*.

Matthew Macoy 2017. *Compressed caching in the age of multicore processors*.

Allan Li 2015. *Examination of the behavior of compressed caching and the optimal use thereof*.

Michael Millian 2015. *Reliability and capabilities of browser-based botnets for distributed computing*.

Gregory Cohan 2014. *Incorporating Twitter data into machine learning models to predict crime*.

Benjamin S. Pullman 2013. *Automatic detection of source code obfuscation*.

Sulaiman Alkhaled 2013. *Machine learning for page replacement*.

Mark Santolucito 2013. *In-memory data compression for modern systems*.

Nicholaus G. Mollé 2010. *Neural Networks and Page Replacement Policies*.

Kevin J. Nattinger 2008. *Evolving a Better Page Replacement Algorithm*.

Owen S. Hofmann 2006. *Reference Trace Reduction via Reference Distribution Sampling*.

Thomas B. Jablin 2006. *The Effects of Hard Disk Caches on Prefetching and Clustering*.

Blair A. Trump 2005. *A Modular Memory Manager in Simulation*.

Deborah S. Katz 2004. *Controlling The Overhead and Accuracy of Online Recency Information*

Gathering.

Isuru M. Seneviratne 2004. *Page-level Reference Locality and User-level Memory Allocators.*

Stephen B. Thomas 2002. *Auction: A New Approach to Cache Partitioning.*

Special Topics Courses

Adam R. Lerner, Fall 2009. *An Exploration of Radio Frequency Identification Protocols and Security*.

Elenor Jessop, Daniel Lees, and Adam Lerner, Fall 2007. *Machine learning for page replacement*.

Thomas B. Jablin, Fall 2004. *Prefetching and Clustering*.

Jared R. White, Spring 2002. *Strategies of Computational Biology*.

George C. Shaw, Fall 2001. *Lossless Compression of Floating Point Data*.

Matthew B. Bevers, Spring 2001. *Auctioning in Variable Space Allocation*.

Rachel L. Rubin, Spring 2000 and Fall 2000. *Variable Space Allocation*.

Advising of Summer Interns

David Dang, Luka Duranovic, Jared Kim, Summer 2021. Measuring and modeling reference distributions for elastic processes.

Luka Duranovic, Nicholas Govus, Tian Xia, Summer 2020 (SURF). Towards a comprehensive re-evaluation of page replacement policies.

Christopher Rabasa, Summer 2019. Efficient gathering of whole-system reference trace collection.

Eden Seyoum and Gabriel Young, Summer 2016. Re-evaluating compressed caching on modern systems.

James Buchanan and Peter Le, Summer 2008. Compressed caching vs. disk access latency hiding in multicore systems.

Elisabeth Baseman, Nicholas Mollel, and Benjamin Zalinger, Summer 2008. The PsuedoOpt page replacement policy with proper evaluation and comparison to recently proposed policies.

Laura M. Strickman and Xin Zheng, Summer 2005. Achieving high utilization under memory pressure for CPU-bound workloads with working set scheduling.

Thomas B. Jablin, Blair A. Trump, and Lisa M. Wallmark, Summer 2004. Experimental evaluation and comparison of page replacement policies using realistic main memories and workloads.

Ana M. Mocanu, Summer 2003. Development and evaluation of whole-system reference trace collection for memory management simulation.

Megan F. Cole, Albree Fellow, Summer 2001. Development and experimentation on demand prepaging.

Stephen B. Thomas, Hughes Fellow, Summer 2000. Development of dynamic memory management policy simulators.

Advising of Graduate Students

Tongping Liu, Department of Computer Science, University of Massachusetts Amherst, Ph.D., Fall 2012. Dissertation committee.

Mel Gorman, Department of Computer Science and Information Systems, University of Limerick, Ph.D. Summer 2010. External examiner.

Gene Novark, Department of Computer Science, University of Massachusetts Amherst, Ph.D., Summer 2010. Dissertation committee member.

Ting Yang, Department of Computer Science, University of Massachusetts Amherst, Ph.D. Spring 2009. Dissertation committee member.

Matthew Hertz, Department of Computer Science, University of Massachusetts Amherst, Ph.D. Fall 2005. Dissertation committee member.

Mel Gorman, Department of Computer Science and Information Systems, University of Limerick, M.S. Spring 2003. External thesis committee member.

Committees and College Service

Science Faculty Steering Committee, Fall 2019 to present.

Orientation Committee, Fall 2018 to Spring 2019.

First Year Seminar Committee, Fall 2013 to Spring 2014.

Student Fellowships Committee, Fall 2011 to Spring 2012.

Science Planning Committee, Fall 2009 to Fall 2010; Fall 2011 to 2015.

Faculty Library Committee, Fall 2007 to Spring 2009.

Library Renovation Committee, Fall 2007 to Spring 2009.

Director of Athletics Search Committee, Spring 2006.

Faculty Liaison to Men's Squash, Fall 2004 to Spring 2006.

Athletics Committee, Chair, Fall 2003 to Spring 2005.

Orientation Committee, Fall 2000 to Spring 2002.

Dean of Student Activities Search Committee, Spring 2000.

Network Administrator Search Committee, Spring 2001 to Fall 2001.

Assistant Coach, Men's Squash, Fall 2000 to Spring 2005.

Activities, Department of Computer Science

Chair of Computer Science, July 2008 to June 2009; Fall 2011 to Spring 2012; Fall 2014 to Spring 2015; October 2019 to present.

Library representative for Computer Science, Spring 2008.

Comprehensive examination organizer, Spring 2005 to Spring 2006, Fall 2007 to Spring 2008.

Co-manager of departmental systems, Fall 2000 to present.

ACM Chapter advisor, Fall 2003 to Spring 2004.

ACM Chapter advisor, Fall 2001 to Spring 2002.

Library representative for Computer Science, Fall 1999 to Spring 2000.