

Fred Dougliis

EMC Corp.
100 Canal Pointe Blvd #205
Princeton, NJ 08540

@FredDougliis

<http://www.linkedin.com/in/freddougliis>

f.dougliis@computer.org

Professional Experience

- 3/15-present, Consultant Research Scientist, [EMC Corp.](#), Princeton, NJ.
1/10-3/15, Consultant Software Engineer, [EMC Corp.](#), Princeton, NJ.
6/09-1/10, Senior Member Technical Staff, Data Domain, Princeton NJ (acquired).
 - Member of the CTO office of EMC Core Technologies Division, [researching and developing](#) technologies for improved storage, especially data protection.
 - Lead strategy for research publications within the division (3/15-).
 - Analysis of disk reliability and proactive prediction of disk failure.
 - Caching algorithms for flash storage.
 - Analysis of deduplication performance and effectiveness in production environments.
 - Scalable routing of content within a deduplicating backup cluster.
 - Automated content-aware assignment of backup clients to a cluster of independent deduplicating backup servers.
- 4/02-2/09, Research Staff Member, [IBM T.J.Watson Research Center](#), Hawthorne, NY.
 - Program codirector, System S streaming data analysis project, 1/08-2/09. Managed technical interactions between IBM and U.S. Department of Defense with respect to project planning, status, technical exchanges and other project operations. Responsible for overall project documentation for use within IBM and Government. On-site consultant for application of System S in classified environment.
 - Team lead for site interoperation component of networked stream processing system., 6/05-1/08. The CLASP system provided a framework for interoperating multiple, autonomous System S sites. It applied techniques from grid computing, such as negotiation for shared resources, to the stream processing environment. I was technical lead for the component (a team of four researchers full-time, collaborating with several others). I worked to apply this component within a classified environment.
 - Team co-lead for storage system supporting a highly scalable data streaming environment. 6/04-6/05. I collaborated with several researchers at the IBM Almaden Research Center to design and prototype a storage

system for System S. Novel features included a “value-based retention” model, which allows different types of data to be assigned values that vary over time, and permits the storage system to delete the lowest-valued data automatically.

- Researcher, System S project, 3/03-2/09. I participated on the project since its inception. Early contributions include the initial evaluation of the requirements for its storage system and investigation into supporting provenance of streaming data, i.e. the relationships between data as they flow through the system.
- Improving performance in a variety of environments, through the use of techniques to eliminate redundant data in communication and storage, 4/02-12/03. The “Data Redundancy Avoidance Toolkit” project explored novel deduplication techniques and compared alternative techniques. I formed the project with a colleague after performing earlier work on using similarity detection to identify files suitable for delta-encoding. Further exploration determined that doing similarity detection on pieces of files provided substantial improvements in size reduction with lower overhead.
- 9/98-3/02, Department Head, Distributed Systems Research, [AT&T Labs - Research](#), Florham Park NJ.
7/96-8/98, Principal Technical Staff Member, [AT&T Labs - Research](#), Florham Park NJ.
1/96-7/96 Senior Technical Staff Member, [AT&T Labs - Research](#),
6/94-1/96 Member of Technical Staff, AT&T Bell Laboratories.
 - Liaison between AT&T Labs--Research and the part of AT&T that offered the Intelligent Content Distribution service. Architecture, standards, project management, vendor evaluation.
 - Research areas included:
 - Internet performance, such as content distribution, web proxy caching, dynamic content generation, and delta-encoding of web pages.
 - Web applications, such as tracking changes to web pages, updating pages dynamically based on user activity, and mobile web access.
 - Adaptive techniques for disconnecting users of modem pools based on inactivity.
- 9/96-1/97, Visiting Lecturer, [Computer Science Department](#), Princeton University, Princeton NJ.
 - Taught [distributed computing](#) (enrollment: approximately 35 undergraduate and graduate students). Advised two-person undergraduate independent project building a Java-based FTP client.
- 10/91-6/94, Scientist, Matsushita Information Technology Laboratory, Princeton NJ.
 - Researched system support for mobile computing, and especially storage systems for mobile computers. Investigated memory-based file systems using DRAM and flash memory.
 - Investigated methods to reduce power consumption of magnetic disks using on-line and off-line heuristics, and using flash memory.

- Built and evaluated mechanism to integrate compression and virtual memory, to improve paging performance on machines with limited memory.
- 10/90-9/91, Visiting Professor, [Vrije Universiteit](#), Amsterdam.
 - Redesigned and implemented process management and virtual memory in the Amoeba system. Investigated the usefulness and feasibility of client file caching in a message-based environment. Compared Amoeba with the Sprite network operating system in order to evaluate kernel architectures, resource allocation models, and communication protocols.
 - Taught seminar on distributed systems (enrollment: approximately 30 students).
- 5/86-9/90, Designing and Implementing Transparent Process Migration, U.C. Berkeley.
 - Evaluated tradeoffs between transparency, performance, reliability, and simplicity in a process migration facility, which permits processes to move transparently during execution and to execute independently of their location. Allowed idle workstations to be used for parallel compilations and other processor-intensive applications, while respecting the autonomy of individual workstations. Investigated alternative methods for managing access to idle hosts. Examined design decisions in the context of practical experience in a community of 20-30 users over a period of over a year.
- 10/84-9/90, Implementing Sprite, U.C. Berkeley.
 - Implemented the Sprite network operating system with other graduate students. In addition to process migration, my contributions included much of the original low-level kernel scheduling, synchronization, and process management code; a UNIX compatibility library to permit porting of UNIX source programs to Sprite; and support for many user-level programs such as emacs and TeX.
- 9/87-6/89, Investigating High-Performance File Systems, U.C. Berkeley.
 - Worked on the preliminary design of a log-structured file system, which would provide high performance by reducing the disk seeks necessary to write small files in arbitrary locations on a disk. Instrumented Silicon Graphics workstations to trace file system operations for the purpose of trace-driven analysis of file system performance.
- Summer 1987, Building an Archive Facility for Optical Disks, Digital Western Research Laboratory.
 - Built a network-wide archive facility for use with write-once optical disks. Users on any host could archive a directory hierarchy and retrieve the archive using a unique identifier.
- 1/86-5/86, Teaching Assistant, U.C. Berkeley.
 - Assisted in teaching the undergraduate operating systems course, working with Professor John Ousterhout.

Education

- Ph.D. (Computer Science) 1990, [University of California, Berkeley](#). Thesis: *Transparent Process Migration in the Sprite Operating System*. Advisor: John Ousterhout.
- M.S. (Computer Science) 1987, [University of California, Berkeley](#). Master's Report: *Process Migration in the Sprite Operating System*. Advisor: John Ousterhout.
- B.S. (Computer Science) *magna cum laude* with Distinction, 1984, [Yale College](#), New Haven, Connecticut.

Honors and Awards

- IEEE Computer Society, Golden Core Award, January 2012
- IEEE Computer Society, Meritorious Service Award, January 2011
- IBM Outstanding Technical Achievement Award for *System S*, July 2008
- IBM Research Division award, 2005 Pat Goldberg Memorial Best Paper Award in Computer Science, Electrical Engineering and Math
- Best paper, International World Wide Web Conference (WWW2004)
- Netherlands Organization for Scientific Research Post-Doctoral Fellowship, 1990-1991.
- California MICRO Fellowship, 1985-1986.
- General Electric Foundation Fellowship, 1984-1985.

Professional Activities

- Editorial Boards
 - [Internet Computing](#), 1/99-present. **Editor in chief**, 1/07-12/10. Associate editor in chief, 10/05-12/06. Guest-edited special issues: Cloud Storage, May/June 2016 (with Robin Chen, forthcoming); [Web-Scale Data Centers](#), July-August 2014; [Virtualization](#) (with Orran Krieger), March-April 2013; [Grid Computing](#) (with Ian Foster), July-August, 2003; [Scalable Internet Services](#) (with Frans Kaashoek), July-August, 2001.
 - [World Wide Web: Internet and Web Information Systems](#) (formerly *World Wide Web*), 1996-2014. Guest-edited a special issue on World Wide Web Characterization and Performance Evaluation, 1999. Guest-edited a special issue of best papers from WWW2005, Dec. 2006.
 - [International Journal of Communication Systems](#), 1999-2004.
- Committees
 - Chair, [IEEE Computer Society Technical Committee on the Internet](#), 11/97-9/00. (Formerly the Task Force on Internetworking). Founder of the committee.
 - Chair, [IEEE Computer Society Technical Committee on Operating Systems and Application Environments](#), 7/96-12/98. Vice-chair, 9/93-7/96.
- Conference Leadership

- Member, steering committee, [USENIX Workshop on Hot Topics in Cloud Computing](#), 2015-.
- Member, steering committee, [ACM/IFIP/USENIX Annual Middleware Conference](#), 2010-. Liaison to USENIX Association, 2010-12.
- Member, steering committee, [IEEE Workshop on Hot Topics on the Web](#), 2006- (sponsored by IEEE-CS TCI).
- **Program Co-chair**, [WWW2005](#)
- **Program chair**, [8th International Web Content Caching and Distribution Workshop](#)
- Vice-chair for applications, [WWW2003](#) program committee.
- Vice-chair for performance, [WWW2002](#) program committee.
- **Program Co-chair**, [2001 Symposium on Applications and the Internet](#). Steering Committee Chair, 6/99-8/00.
- **Program Chair**, [USENIX 2nd Symposium on Internet Technologies and Systems](#).
- **Program Chair**, [1998 USENIX Annual Technical Conference](#).
- Vice-chair for performance, [WWW-8](#) program committee.
- **General Chair** and member of program committee, 4th Workshop on Workstation Operating Systems (1993), subsequently renamed the Workshop on Hot Topics in Operating Systems (HOTOS). Member, steering committee, HOTOS, 1993-1999.
- Member, program committee
 - [INFLOW 2015](#)
 - [ACM/IFIP/USENIX 16th Annual Middleware Conference \(2015\)](#)
 - [2015 ACM Symposium on Cloud Computing](#)
 - [2015 USENIX Annual Technical Conference](#)
 - [FAST 2015](#)
 - [ACM/IFIP/USENIX 15th Annual Middleware Conference \(2014\)](#)
 - [2014 ACM Symposium on Cloud Computing](#)
 - [FAST 2013](#)
 - [ICAC 2012](#)
 - [ACM/IFIP/USENIX 12th Annual Middleware Conference \(2011\)](#)
 - [ICDCS 2011](#)
 - [ACM/IFIP/USENIX 11th Annual Middleware Conference \(2010\)](#)
 - [ICDCS 2010](#)
 - [ACM/IFIP/USENIX 10th Annual Middleware Conference \(2009\)](#)
 - [WISE 2009](#)
 - [ICDCS 2009](#)
 - [ACM/IFIP/USENIX 9th Annual Middleware Conference \(2008\)](#). Also Industrial chair.
 - [8th USENIX Symposium on Operating Systems Design and Implementation](#) (2008)
 - [ACM/IFIP/USENIX 8th Annual Middleware Conference \(2007\)](#)
 - [ICDCS 2007](#)
 - [ACM/IFIP/USENIX 7th Annual Middleware Conference \(2006\)](#)
 - [13th](#), [10th](#), [9th](#), [7th](#) International World Wide Web Conferences

- [2004, 2000, 1997 USENIX Annual Technical Conferences](#)
- [Global Internet 2002](#)
- [4th Symposium on Operating Systems Design & Implementation](#)
- [WISP99, WISP98](#)
- 14th International Conference on Distributed Computing Systems
- 4th Symposium on Experiences with Distributed and Multiprocessor Systems
- Panelist, [ACM/IFIP/USENIX 7th Annual Middleware Conference \(2006\)](#)
- Panelist, National Science Foundation, CSR Small 2014; Parallel and Distributed Systems, 2007, 2006, 2005, 2004, 1999. Member, Committee of Visitors, 2009.
- Member, Ph.D. dissertation committee
 - Xing Lin, University of Utah, *Using Similarity in Content and Access Pattern to Improve Space Efficiency And Performance in Storage Systems*, 2015 (expected)
 - Florin Sultan, Rutgers University, *System Support for Service Availability, Remote Healing and Fault Tolerance Using Lazy State Propagation*, 2004.
 - Arno Bakker, Vrije Universiteit Amsterdam, *An Object-Based Software Distribution Network*, 2002.
- Referee for ACM Transactions on Storage, ACM Transactions on Computer Systems, ACM/IEEE Transactions on Networking, ACM Computing Surveys, IEEE Computer, IEEE Internet Computing, IEEE Transactions on Parallel and Distributed Systems, Computing Systems, Journal of Internetworking, Software - Practice & Experience.
- Senior Member, IEEE.
Member, ACM, USENIX.

Refereed Publications

- Xing Lin, Fred Douglass, Jim Li, Xudong Li, Robert Ricci, Stephen Smaldone, and Grant Wallace, “Metadata Considered Harmful...to Deduplication,” *Proceedings of the 7th USENIX Workshop on Hot Topics in Storage and File Systems*, July 2015.
- Ao Ma, Fred Douglass, Guanlin Lu, Darren Sawyer, Surendar Chandra, and Windsor Hsu, “RAIDShield: Characterizing, Monitoring, and Proactively Protecting Against Disk Failures,” *Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST'15)*, Feb. 2015. Selected for fast-track publication to *ACM Transactions on Storage*.
- Cheng Li, Philip Shilane, Fred Douglass, Hyong Shim, and Stephen Smaldone, “Nitro: A Capacity-Optimized SSD Cache for Primary Storage.” *Proceedings of the 2014 Annual Technical Conference*, June 2014.
- Cheng Li, Philip Shilane, Fred Douglass, Darren Sawyer, and Hyong Shim, “Assert(!Defined(Sequential I/O)),” *Proceedings of HotStorage'14*, June 2014.
- Xing Lin, Guanlin Lu, Fred Douglass, Philip Shilane, and Grant Wallace, “Migratory Compression: Coarse-grained Data Reordering to Improve

- Compressibility,” *Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST'14)*, Feb. 2014.
- Grant Wallace, Fred Dougkis, Hangwei Qian, Philip Shilane, Stephen Smaldone, Mark Chamness, and Windsor Hsu, “Characteristics of Backup Workloads in Production Systems,” *Proceedings of the 10th USENIX Conference on File and Storage Technologies (FAST'12)*, Feb. 2012.
 - Fred Dougkis, Deepti Bhardwaj, Hangwei Qian, and Philip Shilane, “Content-aware Load Balancing for Distributed Backup”, USENIX 25th Large Installation System Administration Conference. Dec. 2011.
 - Wei Dong, Fred Dougkis, Kai Li, Hugo Patterson, Sazzala Reddy, and Philip Shilane. “Tradeoffs in scalable data routing for deduplication clusters,” *Proceedings of the 9th USENIX Conference on File and Storage Technologies (FAST'11)*, Feb. 2011.
 - Bo Yang, Hao Wang, and Fred Dougkis, “Applying Knowledge Sharing for Business Intelligence Collaboration,” IEEE International Conference on Web Services (ICWS) 2009.
 - Bo An, Fred Dougkis, and Fan Ye, “Heuristics for Negotiation Schedules in Multi-plan Optimization,” in *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems*, May 2008.
 - Fred Dougkis, “Best Practices for the Care and Feeding of a Program Committee, and Other Thoughts on Conference Organization,” Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems (WOWCS '08), April 2008.
 - Kirsten Hildrum, Fred Dougkis, Joel L. Wolf, Philip Yu, Lisa Fleischer, and Akshay Katta, "Storage Optimization for Large-Scale Distributed Stream Processing Systems," *ACM Transactions on Storage*, Feb. 2008. An earlier version appeared in the Third International Workshop on System Management Techniques, Processes and Services (SMTPS 2007), March 2007.
 - Michael Branson, Fred Dougkis, Brad Fawcett, Zhen Liu, Anton Riabov and Fan Ye, “CLASP: Collaborating, Autonomous Stream Processing Systems,” 8th International Middleware Conference, November 2007.
 - Michael Branson, Fred Dougkis, Brad Fawcett, Zhen Liu, Anton Riabov and Fan Ye, "Autonomic Operations in Cooperative Stream Processing Systems," Second Workshop on Hot Topics in Autonomic Computing, June 2007.
 - Bin Rong, Fred Dougkis, Zhen Liu, Cathy Xia, "Failure Recovery in Cooperative Data Stream Analysis," Second International Conference on Availability, Reliability and Security (ARES 2007), April 2007.
 - Ranjita Bhagwan, Fred Dougkis, Kirsten Hildrum, Jeffrey O. Kephart, and William E. Walsh, "Time-varying management of data storage," Workshop on Hot Topics in System Dependability, June 2005.
 - Lakshmish Ramaswamy, Arun Iyengar, Ling Liu, and Fred Dougkis, "Automatic Fragment Detection in Dynamic Web Pages and its Impact on Caching" , *IEEE Transactions on Knowledge and Data Engineering*, Vol.17 No.6, June 2005. *IBM Research best paper award*, 2005.
 - Fred Dougkis, John Palmer, Elizabeth S. Richards, David Tao, William H. Tetzlaff, John M. Tracey, and Jian Yin, "Position: Short Object Lifetimes Require

- a Delete-Optimized Storage System," 11th ACM SIGOPS European Workshop, September 2004.
- Purushottam Kulkarni, Fred Douglass, Jason LaVoie, and John M. Tracey, Redundancy Elimination Within Large Collections of Files, USENIX 2004, Boston, MA, June 2004.
 - Lakshmin Ramaswamy, Arun Iyengar, Ling Liu, and Fred Douglass, Automatic Detection of Fragments in Dynamically Generated Web Pages, 13th International World Wide Web Conference (WWW2004), New York City, May 2004. *Best paper award recipient.*
 - Fred Douglass and Arun Iyengar, Application-Specific Delta-Encoding via Resemblance Detection, 2003 USENIX Annual Technical Conference, June 2003.
 - Michael Rabinovich, Zhen Xiao, Fred Douglass, and Chuck Kalmanek, Moving Edge Side Includes to the Real Edge—the Clients, 4th USENIX Symposium on Internet Technologies and Systems, March 2003.
 - Zhuoqing Morley Mao, Charles D. Cranor, Fred Douglass, Michael Rabinovich, Oliver Spatscheck, and Jia Wang, A Precise and Efficient Evaluation of the Proximity between Web Clients and their Local DNS Servers, 2002 USENIX Annual Technical Conference, June 2002.
 - Fred Douglass, Sonia Jain, John Klensin, and Michael Rabinovich, "Click-once Hypertext: Now You See It, Now You Don't," WIAPP'01, July 2001.
 - Alex Biliris, Chuck Cranor, Fred Douglass, Michael Rabinovich, Sandeep Sibal, Oliver Spatscheck, and Walter Sturm, "CDN Brokering," WCW'01, June 2001.
 - Yih-Farn Chen, Fred Douglass, Huale Huang, and Kiem-Phong Vo, "TopBlend: An Efficient Implementation of HtmlDiff in Java", WebNet'00, November 2000.
 - Dejan Milojicic, Fred Douglass, Yves Panedeine, Ric Wheeler, and Songnian Zhou, "Process Migration," *ACM Computing Surveys*, 32(3), pp. 241-299, September 2000.
 - Fred Douglass, Server-side tracking of new documents, *Proceedings of the 1st International Workshop on Web Site Evolution*, Atlanta, GA, October 1999.
 - Fred Douglass and Thomas Killian, "Adaptive Modem Connection Lifetimes." 1999 USENIX Annual Technical Conference, pp. 27--41, June, 1999.
 - Anja Feldmann, Ramón Cáceres, Fred Douglass, Gideon Glass, and Michael Rabinovich, "Performance of Web Proxy Caching in Heterogeneous Bandwidth Environments." *Proceedings of IEEE Infocom'99*, pp. 106-116, March 1999.
 - Ramón Cáceres, Fred Douglass, Anja Feldmann, Gideon Glass, and Michael Rabinovich. "Web Proxy Caching: The Devil is in the Details ." SIGMETRICS Workshop on Internet Server Performance, June 1998.
 - Fred Douglass, Thomas Ball, Yi h-Farn Chen, and Eleftherios Koutsofios, "The AT&T Internet Difference Engine: Tracking and Viewing Changes on the Web", *World Wide Web*, January 1998. Also published as AT&T Labs--Research TR 97.23.1, April, 1997.
 - Fred Douglass, Anja Feldmann, Balachander Krishnamurthy, and Jeffrey Mogul, "Rate of Change and other Metrics: a Live Study of the World Wide Web," USENIX Symposium on Internet Technologies and Systems, December 1997.

- Fred Douglass, Antonio Haro, and Michael Rabinovich, "HPP: HTML Macro-Preprocessing to Support Dynamic Document Caching." USENIX Symposium on Internet Technologies and Systems, December 1997.
- Jeffrey Mogul, Fred Douglass, Anja Feldmann, and Balachander Krishnamurthy, "Potential Benefits of Delta-encoding and Data Compression for HTTP." Proceedings of Sigcomm'97, September, 1997, pp. 181--194. An extended version appears as Research Report 97/4, Digital Equipment Corporation Western Research Laboratory, July, 1997.
- Gaurav Banga, Fred Douglass, and Michael Rabinovich. Optimistic Deltas for WWW Latency Reduction. In *Proceedings of 1997 USENIX Technical Conference*, pages 289--303, January 1997.
- Fred Douglass, Tom Ball, Robin Chen, and Eleftherios Koutsofios. Webguide: Querying and navigating changes in web repositories. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.
- Bill N. Schilit, Fred Douglass, David M. Kristol, Paul Krzyzanowski, James Sienicki, and John A. Trotter. Teleweb: Loosely connected access to the world wide web. In *Proceedings of the Fifth International World Wide Web Conference*, Paris, France, May 1996.
- Fred Douglass and Thomas Ball. Tracking and viewing changes on the web. In *Proceedings of 1996 USENIX Technical Conference*, pages 165--176, San Diego, CA, January 1996.
- Fred Douglass, P. Krishnan, and Brian Bershad. Adaptive disk spin-down policies for mobile computers. *Computing Systems*, 8(4):381--413, Fall 1995. An earlier version appeared in *Proceedings of the Second Symposium on Mobile and Location-independent Computing*, pp. 121--137, April 1995.
- Fred Douglass, Ramón Cáceres, Brian Marsh, Frans Kaashoek, Kai Li, and Joshua Tauber. "Storage Alternatives for Mobile Computers," in *Mobile Computing*, chapter 18, pages 473--505. Kluwer Academic Publishers, 1996. An earlier version appeared in *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pp. 25--37, November, 1994.
- Fred Douglass, P. Krishnan, and Brian Marsh. "Thwarting the Power-Hungry Disk." *Proceedings of 1994 Winter USENIX Conference*, pp. 293--306, January 1994.
- Brian Marsh, Fred Douglass, and P. Krishnan. "Flash Memory File Caching for Mobile Computers." *Proceedings of the 27th Hawaii Conference on Systems Sciences*, pp. 451--460, January 1994.
- Ramón Cáceres, Fred Douglass, Kai Li, and Brian Marsh. "Operating Systems Implications of Solid-State Mobile Computers," *Proceedings of the Fourth Workshop on Workstation Operating Systems (WWOS-IV)*, pp. 21--27, October 1993.
- Daniel Barbara, Chris Clifton, Fred Douglass, Hector Garcia-Molina, Stephen Johnson, Ben Kao, Sharad Mehrotra, Jens Tellefsen, and Rosemary Walsh. "The {G}old mailer" *9th International Conference on Data Engineering*, Vienna, April 1993.

- Fred Douglass, "The Compression Cache: Using On-line Compression to Extend Physical Memory," *Proceedings of 1993 Winter USENIX Conference*, pp. 519--529, January 1993.
- Fred Douglass, "On the Role of Compression in Distributed Systems," *Proceedings of the Fifth ACM SIGOPS European Workshop*, September 1992. Also appears in *ACM Operating Systems Review*, 27(2):88--93, April 1993.
- Fred Douglass and Brian Marsh, "The Workstation as a Waystation: Integrating Mobility into Computing Environments," *Proceedings of the Third Workshop on Workstation Operating Systems (WWOS-III)*, pp. 17--21, April 1992.
- Fred Douglass, John K. Ousterhout, M. Frans Kaashoek, and Andrew. S. Tanenbaum, "A Comparison of Two Distributed Systems: Amoeba and Sprite," *Computing Systems*, 4(4):353--384, Fall 1991.
- Fred Douglass and John Ousterhout, "Transparent Process Migration: Design Alternatives and the Sprite Implementation," *Software--Practice & Experience*, 21(8):757--785, August 1991.
- Fred Douglass, "Experience with Process Migration in Sprite," *Workshop on Experiences with Building Distributed and Multiprocessor Systems*, pp. 59--72, October 1989.
- Brent Welch, Mary Baker, Fred Douglass, John Hartman, Mendel Rosenblum, and John Ousterhout, "Sprite Position Statement: Use Distributed State for Failure Recovery," *Proceedings of the Second Workshop on Workstation Operating Systems (WWOS-II)*, pp. 130--133, September 1989.
- John Ousterhout, Andrew Cherenon, Fred Douglass, Michael Nelson, and Brent Welch, "The Sprite Network Operating System," *Computer*, 21(2):23--36, February 1988.
- Fred Douglass and John Ousterhout, "Process Migration in the Sprite Operating System," *Proceedings of the 7th International Conference on Distributed Computing Systems*, pp. 18--25, September 1987.

Books

- Dejan Milojicic, Frederick Douglass, and Richard Wheeler. eds., *Mobility: Processes, Computers, and Agents*. Addison-Wesley Longman, and ACM Press, February 1999.

Unrefereed Publications

- *IEEE Internet Computing* editor-in-chief columns
 - Fred Douglass, "From the Editor in Chief: Thanks for the Fish ??? but I'm Drowning!," *IEEE Internet Computing*, vol. 14, no. 6, Nov/Dec, 2010.
 - Fred Douglass, "From the Editor in Chief: Closing the Open (Face) Book," *IEEE Internet Computing*, vol. 14, no. 5, Sep/Oct, 2010.
 - Fred Douglass, "From the Editor in Chief: As I Emerge from the Mobile Phone Dark Ages, I Look Around in Fear and Wonder," *IEEE Internet Computing*, vol. 14, no. 4, Jul/Aug, 2010.

- Fred Douglass, "From the Editor in Chief: Sinking or Swimming in the Sea of Data," *IEEE Internet Computing*, vol. 14, no. 3, May/June, 2010.
- Fred Douglass, "From the Editor in Chief: Turnpikes or Freeways?," *IEEE Internet Computing*, vol. 14, no. 2, Mar/Apr, 2010.
- Fred Douglass, "From the Editor in Chief: It's All About the (Social) Network," *IEEE Internet Computing*, vol. 14, no. 1, Jan/Feb, 2010.
- Fred Douglass, "From the Editor in Chief: Phone + Internet Café = Secure Banking? You Betcha," *IEEE Internet Computing*, vol. 13, no. 6, Nov/Dec, 2009.
- Fred Douglass, "From the Editor in Chief: Information Overload, 140 Characters at a Time," *IEEE Internet Computing*, vol. 13, no. 4, Jul/Aug, 2009.
- Fred Douglass, "From the Editor in Chief: Staring at Clouds," *IEEE Internet Computing*, vol. 13, no. 3, May/June, 2009.
- Fred Douglass, "From the Editor in Chief: Unease of Use," *IEEE Internet Computing*, vol. 13, no. 2, Mar/Apr, 2009.
- Fred Douglass, "From the Editor in Chief: Protect Us from Ourselves," *IEEE Internet Computing*, vol. 13, no. 1, Jan/Feb, 2009.
- Fred Douglass, "From the Editor in Chief: Full Stream Ahead," *IEEE Internet Computing*, vol. 12, no. 6, Nov/Dec, 2008.
- Fred Douglass, "From the Editor in Chief: Ideas Ahead of Their Time," *IEEE Internet Computing*, vol. 12, no. 5, Sep/Oct, 2008.
- Fred Douglass, "From the Editor in Chief: The Eye of the Beholder: Whence IC?" *IEEE Internet Computing*, vol. 12, no. 3, May/June, 2008.
- Fred Douglass, "From the Editor in Chief: Not-so-secret Identities," *IEEE Internet Computing*, vol. 12, no. 2, Mar/Apr, 2008.
- Fred Douglass, "From the Editor in Chief: On Social Networking and Communication Paradigms," *IEEE Internet Computing*, vol. 12, no. 1, Jan/Feb, 2008.
- Fred Douglass, "From the Editor in Chief: Collective Wisdom: A Modest Proposal to Improve Peer Review, Part 2," *IEEE Internet Computing*, vol. 11, no. 6, Nov/Dec, 2007.
- Fred Douglass, "From the Editor in Chief: Collective Wisdom: A Modest Proposal to Improve Peer Review, Part 1," *IEEE Internet Computing*, vol. 11, no. 5, Sep/Oct, 2007.
- Fred Douglass, "From the Editor in Chief: What's Your PageRank?," *IEEE Internet Computing*, vol. 11, no. 4, pp. 3-4, Jul/Aug, 2007
- Fred Douglass, "From the Editor in Chief: The Search for Jim, and the Search for Altruism," *IEEE Internet Computing*, vol. 11, no. 3, pp. 4-5, May/June, 2007
- Fred Douglass, "From the Editor in Chief: Pet Peeves," *IEEE Internet Computing*, vol. 11, no. 2, pp. 4-5, Mar/Apr, 2007
- Fred Douglass, "From the Editor in Chief: Looking Forward, Looking Back," *IEEE Internet Computing*, vol. 11, no. 1, pp. 4-5, Jan/Feb, 2007
- RFCs

- RFC 3229, "Delta encoding in HTTP", J Mogul, Y Goland, Arthur van Hoff, Fred Dougliis, Anja Feldmann, Balachander Krishnamurthy, Daniel Hellerstein, January 2002.
- Internet-Drafts
 - "Known Mechanisms for Content Internetworking", Fred Dougliis, Imran Chaudhri, Phillip Rzewski, 11/09/2001.
 - "Known CDN Request-Routing Mechanisms", R. Nair, Mark Green, Brad Cain, Oliver Spatscheck, Fred Dougliis, Markus Hofmann, Doug Potter, A Barbir, 11/01/2001.
- Fred Dougliis, Michael Branson, Kirsten W. Hildrum, Bin Rong, and Fan Ye, "Multisite Cooperative Data Stream Analysis," *ACM SIGOPS Operating Systems Review*, July 2006.
- Fred Dougliis, Steve Chapin, and Jim Isaak, "Internet Research on Internet Time," Technical Activities Forum, *Computer*, 31(11), November, 1998, pages 76--78.
- Fred Dougliis. Experiences with the AT&T Internet Difference Engine. In *Proceedings of the 22nd International Conference for the Resource Management & Performance Evaluation of Enterprise Computing System (CMG96)*, December 1996.
- Thomas Ball and Fred Dougliis. An internet difference engine and its applications. In *Digest of Paper, COMPCON '96*, pages 71--76, February 1996.
- Fred Dougliis and Brian Marsh, *Low-power Disk Management for Mobile Computers*, Matsushita Information Technology Laboratory, Technical Report MITL-TR-53-93, April 1993.
- Fred Dougliis, Daniel Barbara, Brian Marsh, and Sharad Mehrotra, *Memory Management to Support Large UNIX Applications*, Matsushita Information Technology Laboratory, Technical Report MITL-TR-37-92, December 1992.
- Frederick Dougliis, *Transparent Process Migration in the Sprite Operating System*, Computer Science Division (EECS), U.C. Berkeley, Technical Report UCB/CSD 90/598, September 1990. PhD. thesis.
- Fred Dougliis and John Ousterhout, "Log-Structured File Systems," *Digest of Papers, Spring Compcon '89*, pp. 124--129, February 1989.
- Fred Dougliis and John Ousterhout, "Process migration in Sprite: A status report," *IEEE Computer Society Technical Committee on Operating Systems Newsletter*, 3(1):8--10, Winter 1989.
- Brian Marsh and Fred Dougliis and P. Krishnan, *Flash Memory File Caching for Mobile Computers*, Matsushita Information Technology Laboratory, Technical Report MITL-TR-59-93, June 1993.
- John Ousterhout and Fred Dougliis, "Beating the I/O Bottleneck: A Case for Log-Structured File Systems", *ACM Operating Systems Review*, 23(1):11--28, January 1989.

Patents

- "Methods, systems and products for generating web pages," (with S.Jain, J. Klensin, and M. Rabinovich), [U.S. patent 9,020,904](#), April 28, 2015.

- “Unified web hosting and content distribution system and method for assuring predetermined performance levels ,” (with D.H. Shur, J.M. Sommer, and J.E. Van Der Merwe), [U.S. Patent 8,943,193](#), January 27, 2015.
- “Indexing architecture for deduplicated cache system of a storage system,” (with P. N. Shilane, H. Shim, S. Smaldone, G. Wallace, and C. Li), [U.S. Patent 8,935,446](#), January 13, 2015.
- “Out-of-core similarity matching,” (with G. Wallace and P. N. Shilane), [U.S. Patent 8,914,338](#), December 16, 2014.
- “Bulk deletion through segmented files,” (with E. Chron and S. Morgan), [U.S. Patent 8,914,330](#), December 16, 2014.
- “Method for the interoperation of virtual organizations,” (with M. Branson, B. Fawcett, Z. Liu, W. Waller, and F. Ye), [U.S. Patent 8,889,624](#), November 18, 2014.
- “Method and system for detecting unwanted content of files,” (with G. Wallace and P. N. Shilane), [U.S. Patent 8,825,626](#), September 2, 2014.
- “Method and apparatus for efficiently searching data in a storage system,” (with G. Wallace and P. N. Shilane), [U.S. Patent 8,756,249](#), June 17, 2014.
- “State-based directing of segments in a multinode deduplicated storage system,” (with P. Shilane and R. H. Patterson), [U.S. Patent 8,751,448](#), June 10, 2014.
- “Method and apparatus for content-aware resizing of data chunks for replication,” (with P. N. Shilane and G. Wallace), [U.S. Patent 8,712,963](#), April 29, 2014.
- “Method for inter-site data stream transfer in cooperative data stream processing,” (with M. J. Branson, Z. Liu, and F. Ye), [U.S. Patent 8,688,850](#), April 1, 2014.
- “System and method for optimizing data in value-based storage system,” (with N. Bansal, L. K. Fleischer, K. W. Hildrum, A. K. R. Katta, J. D. Palmer, E. S. Richards, D. Tao, W. H. Tetzlaff, and J. L. Wolf), [U.S. Patent 8,683,150](#), March 25, 2014.
- “Efficient content meta-data collection and trace generation from deduplicated storage,” (with P. Shilane and G. Wallace), [U.S. Patent 8,667,032](#), March 4, 2014.
- “Negotiation management in data processing systems,” (with B. An, B. W. Fawcett, A. V. Riabov, and F. Ye), [U.S. Patent 8,656,403](#), February 18, 2014.
- “Method and apparatus for determining optimal chunk sizes of a deduplicated storage system,” (with P. N. Shilane and G. Wallace), [U.S. Patent 8,639,669](#), January 28, 2014.
- “Efficient content meta-data collection and trace generation from deduplicated storage,” (with P. Shilane and G. Wallace), [U.S. Patent 8,631,052](#), January 14, 2014.
- “Data migration with load balancing and optimization,” [U.S. Patent 8,554,918](#), October 8, 2013.
- “System and method for allocating a storage unit for backup in a storage system with load balancing,” (with R. H. Patterson and P. N. Shilane), [U.S. Patent 8,484,356](#), July 9, 2013.
- “System and method for composition of stream processing service environments,” (with M. J. Branson, P. Dube, and Z. A. Garbow), [U.S. Patent 8,478,874](#), July 2, 2013.

- “Mechanism for execution of multi-site jobs in a data stream processing system,” (with M. J. Branson and F. Ye), [U.S. Patent 8,417,762](#), April 9, 2013.
- “Apparatus for reducing accesses to levels of a storage hierarchy in a computing system,” (with M. C. Olsen, M. Rosu, and T. R. Hildner), [U.S. Patent 8,392,655](#), March 5, 2013.
- “Method and apparatus for cooperative data stream processing,” (with M. J. Branson, B. W. Fawcett, Z. Liu, W. Waller, and F. Ye), [U.S. Patent 8,359,347](#), January 22, 2013.
- “Systems and methods for securely processing sensitive streams in a mixed infrastructure,” (with M. J. Branson, P. Dube, and Z. A. Garbow), [U.S. Patent 8,276,200](#), September 25, 2012.
- “System and method for optimizing data in value-based storage system,” (with N. Bansal, L. K. Fleischer, K. W. Hildrum, A. K. R. Katta, J. D. Palmer, E. S. Richards, D. Tao, W. H. Tetzlaff, and J. L. Wolf), [U.S. Patent 8,275,957](#), September 25, 2012.
- “Methods and apparatus for effective on-line backup selection for failure recovery in distributed stream processing systems,” (with Z. Liu, H. Xia, and B. Rong), [U.S. Patent 8,225,129](#), July 17, 2012.
- • “Mechanism for recovery from site failure in a stream processing system,” (with M. Branson, B. Fawcett, Z. Liu, B. Rong, and F. Ye), [U.S. Patent 8,219,848](#), July 10, 2012.
- • “Unified web hosting and content distribution system and method for assuring predetermined performance levels,” (with D. Shur, J. Sommer, and J. Van der Merwe), [U.S. Patent 8,176,172](#), May 8, 2012.
- “Method and apparatus for data redundancy elimination at the block level,” (with P. Kulkarni, J. LaVoie, and J. M. Tracey), [U.S. Patent 8,135,683](#), March 13, 2012.
- “Methods and systems for detecting fragments in electronic documents,” (with A. Ivengar and L. Ramaswamy), [U.S. Patent 8,037,081](#), October 11, 2011.
- “Optimizing a storage system to support short data lifetimes,” (with K.S. Anderson, N. Halim, J.D. Palmer, J. Davis, E.S. Richards, D. Tao, W.H. Tetzlaff, J.M. Tracey, and J.L. Wolf), [U.S. Patent 7,958,093](#), Jun 7, 2011.
- “Methods and systems for publishing electronic documents with automatic fragment detection,” (with A.K. Iyengar and L.M. Ramaswamy), [U.S. Patent 7,870,100](#) January 11, 2011.
- “Method and apparatus for limiting reuse of domain name system information,” (with M. Rabinovich and O. Spatscheck), [U.S. Patent 7,725,536](#), May 25, 2010.
- “Unified web hosting and content distribution system and method for assuring predetermined performance levels,” (with D.H. Shur, J.M. Sommer and J.E. Van Der Merwe), [U.S. Patent 7,657,622](#), February 2, 2010.
- “Eliding web page content.” (with S.Jain, J. Klensin, and M. Rabinovich), [U.S. patent 7,620,898](#), November 17, 2009.
- “Method and apparatus for content distribution network brokering and peering,” (with A. Biliris, C. D. Cranor, C. Nelson, M. Rabinovich, S. Sibal, O. Spatscheck, and W. Sturm), [U.S. Patent 7,562,153](#), Jul 14, 2009.

- "Method and apparatus for limiting reuse of domain name system response information," (with M. Rabinovich and O. Spatscheck), [U.S. Patent 7,444,371](#), October 28, 2008.
- "Method and apparatus for scalable peer-to-peer inquiries in a network of untrusted parties," (with F. Jania, J. LaVoie, and J. Tracey), [U.S. Patent 7,325,034](#), January 29, 2008.
- "Method and apparatus for scalable peer-to-peer inquiries in a network of untrusted parties," (with F.L. Jania, J.D. LaVoie, and J.M. Tracey), [U.S. patent 7,216,297](#), May 8, 2007.
- "Method for content distribution in a network supporting a security protocol," (with M. Rabinovich, A. Rubin, and O. Spatscheck), [U.S. patent 7,149,803](#), Dec. 12, 2006.
- "Method for transferring and displaying data pages on a data network," (with G. Banga, H.V. Jagadish, M. Rabinovich, and K.-P. Vo), [U.S. patent 6,910,073](#), June 21, 2005.
- "Adaptive modem connection lifetimes," (with T. Killian), [U.S. patent 6,598,082](#), July 22, 2003.
- "Identifying changes in on-line data repositories," (with T. Ball), [U.S. patent 6,596,030](#), July 22, 2003
- "Management of time and expense when communicating between a host and a communication network," (with D.M. Kristol, P. Krzyzanowski, J.A. Trotter, J.P. Sienicki, and W.N. Schilit), [U.S. patent 6,587,877](#), July 1, 2003.
- "Adaptive modem connection lifetimes," (with T. Killian), [U.S. patent 6,487,596](#), November 26, 2002.
- "Method and apparatus for tracking and viewing changes on the web," (with T. Ball), [U.S. patent 6,366,933](#), April 2, 2002.
- "Personalizing the display of changes to records in an on-line repository," [U.S. patent 6,249,795](#), June 19, 2001.
- "Method for reducing perceived delay between a time data is requested and a time data is available for display," (with G. Banga, H.V. Jagadish, and M. Rabinovich), [U.S. patent 6,240,447](#), May 29, 2001.
- "Method and apparatus for dynamic data transfer on a web page," (with A. Haro and M. Rabinovich), [U.S. patent 6,021,426](#), February 2, 2000.
- "Method for reducing the delay between the time a data page is requested and the time the data page is displayed," (with G. Banga, H.V. Jagadish, M. Rabinovich, and K.-P. Vo), [U.S. patent 5,931,904](#), August 3, 1999.
- "Querying and navigating changes in web repositories" (with T. Ball, Y. Chen, and E. Koutsofios), [U.S. patent 5,860,071](#), January 12, 1999.
- "A method and system for reducing access latency by providing fine grain direct access to flash memory concurrent with a block transfer therefrom" (with R. Cáceres, B. Marsh, and B. Bershad), [U.S. patent 5,802,554](#), September 1, 1998.
- "Adaptive Disk Spin-down Method for Managing the Power Distributed to a Disk Drive in a Laptop Computer," (with B. Marsh, B. Bershad, and P. Krishnan), [U.S. patent 5,493,670](#), February 20, 1996.

- "Method for Managing the Power Distributed to a Disk Drive in a Laptop Computer," (with B. Marsh and P. Krishnan), [U.S. patent 5,481,733](#), January 2, 1996.

Approximately 40 patent applications under review.