

CS 535 Advanced Operating Systems  
Craig E. Wills  
Monday, January 14, 2008

WPI, Spring 2008  
Course Syllabus

## Instructor

Craig E. Wills, FL-236, cew@cs.wpi.edu. Office hours: as needed by students. Any time for short questions. Electronic mail is an effective method to contact me.

## Course Web Page

Copies of all handouts, assignments and notes will be posted as appropriate on the course Web page. The address for it is <http://www.cs.wpi.edu/~cs535/s08/>.

## Purpose

This is a graduate-level course in the design of advanced systems. It focuses on the issues of advanced operating systems and distributed systems, which have evolved from general-purpose multiprogramming systems covered in previous courses. The goals are 1) to familiarize the student with current literature in the area, 2) for students to be able to place work in its context in terms of its relative importance and relationship with other work, 3) to give students experience in making public presentations of technical work, and 4) to develop students' ability for critical thinking and discussions concerning design choices, tradeoffs, and their consequences.

## Prerequisites

A first course in operating systems, such as CS 502. Background in architecture, networking, compilers and programming languages would also be helpful for issues related to these topics. An interest in reading, thinking about and discussing issues in advanced system design.

## Text Book

There is no required text book for the course. Relevant text books are:

1. *Distributed Systems: Principles and Paradigms*. Andrew S. Tanenbaum and Maarten van Steen, Prentice Hall, 2002.

2. *Distributed Systems: Concepts and Designs*, Second Edition. George Coulouris, Jean Dollimore and Tim Kindberg, Addison Wesley 1994.
3. *Operating Systems Concepts*, Seventh Edition. Abraham Silberschatz, Peter Baer Galvin and Greg Gagne, John Wiley & Sons, 2005.
4. *Distributed Systems*. Sape Mullender, Addison Wesley 1993.

They provide reasonably up-to-date discussions on issues in advanced and distributed systems. Some material from these sources will be introduced in class.

Much of the course will also be based on seminal and current literature in the field. Literature will be selected both by the instructor and the students. Access to this material will be discussed in class.

## Grading Policy

Final grades will be computed as follows:

Homework: 35%;

Final Project and Presentation: 30%;

Final Exam: 25%;

Class participation: 10%.

Final grades will reflect the extent to which you have demonstrated understanding and insight for the material. No incomplete grades will be assigned unless there exist exceptional, extenuating circumstances. There will be one final exam. There will be one final project that may or may not involve programming.

Class meetings will consist of discussions. Each student should be prepared to pose and answer questions on the day a paper or topic is discussed. When reading papers you are expected to take notes on the main points of each article, questions concerning its contents and suggestions on how the described research can be followed up. You should try to discern the strengths and weaknesses of the paper. After reading the paper, you should write a summary of the paper along with questions and suggestions for followup (roughly a page or so). It is important to include your personal thoughts on the paper concerning what you liked or disliked about the paper. An excellent review will include a summary containing the key points of the paper along with personal observations of the paper.

For your own benefit you should maintain an online bibliography of the papers you read. The bibliography should contain papers you read for the class and any other references of interest that you come across as you read or peruse the literature. These entries will provide ideas for further reading for the course.

## Late Policy

You are expected to keep up with the pace of papers for the course. Entries turned in late will be penalized 5% of total assignment value per day or partial day. All entries are due at the *start* of class on the due date. Those turned in after the start of class will be counted late. Exceptions to these rules can be made only *a priori*. Finally, no entries will be accepted after the last day of class to allow sufficient time for grading.

## Cheating

Unless explicitly noted, all work is to be done on an individual basis. You are encouraged to talk with others about ideas and material in the course, particularly in preparing for exams. However all work, in the form of reviews, code or answers to problems, you submit for grading must be your work. Misrepresentation of the work of another as one's own submitted work is a violation of academic honesty. Aiding someone else to commit an act of academic dishonesty is also a violation. Submitting individually-assigned work that was jointly done with another person is a violation of academic honesty.

Any violation of the WPI's guidelines for academic honesty will result in no credit for the course and referral to the Student Affairs Office. More information on definitions, responsibilities and procedures regarding the WPI academic honesty policy can be found at <http://www.wpi.edu/Pubs/Policies/Judicial/sect5.html>.

## Schedule

The following is a tentative outline of the material that will be covered in this course.

**week 1:** 1-14. Course outline, contents, grading policy, introduction, and overview.

Papers to read: [109, 66].

**week 2:** 1-21. System Structure.

Papers to read: [11, 110].

Other papers of interest: [83, 12, 114, 49, 7, 45, 53, 39, 18, 69, 31, 65, 123].

**week 3:** 1-28. Performance evaluation/Benchmarking

Papers to read: [24, 104].

Other papers of interest: [44, 94, 42, 73, 34, 25, 59, 98].

**week 4:** 2-4. Communication/RPC/multicast/network protocols.

Papers to read: [95, 17].

Other papers of interest: [80, 20, 10, 70, 77, 32, 15, 48, 106, 52, 78, 82].

**week 5:** 2-11. Synchronization—clocks, mutual exclusion, deadlocks.

Papers to read: [60].

Other papers of interest: [108].

**week 6:** 2-18. Processes and multiprocessors, threads, memory management, load sharing, scheduling.

Papers to read: [33, 121].

Other papers of interest: [35, 64, 21, 22, 51, 3, 96, 14, 107, 47, 5, 71, 117, 63, 62, 29, 85].

**week 7:** 2-25. Distributed File Systems and File Sharing.

Papers to read: [41, 84].

Other papers of interest: [6, 2, 75, 81, 57, 91, 100, 43, 103, 4, 112, 68, 99].

**weeks 8-14:** 3-3 – 4-21. Tentatively no class: March 3.

Other issues:

Name servers/resource location [111, 46, 16, 105, 19].

Object-oriented systems [28].

Protection and security [124, 54, 36, 30, 116, 120, 122].

Experience with systems: [27, 97, 61].

Distributed operating systems [74, 119, 115, 93, 79, 26].

Distributed computing environments [87, 40, 58, 37, 50].

Web as a Distributed System [86, 101, 13, 72, 23, 8, 55, 56].

Caching/Prefetching [76]

Peer-to-Peer Systems [90, 102]

Wireless/Mobile Systems [113, 92, 38, 1]

Network Servers and Services [118, 9, 88]

Sensor Networks [67]

Grid Computing [89], Network Management, Web Services

## Systems Literature

We will be reading articles from journals and conference proceedings. Three primary publishers of journals and conference proceedings are ACM ([www.acm.org](http://www.acm.org)), IEEE Computer Society ([www.computer.org](http://www.computer.org)) and USENIX ([www.usenix.org](http://www.usenix.org)). Each of these has their own digital library, which is a prominent link from each of these organization's home page. WPI is a member of each of these organizations and access from a machine on the WPI campus will allow download of all articles from ACM and IEEE. If you are off-campus, then you can use the WPI library proxy (<http://www.wpi.edu/Academics/CCC/Help/Software/proxy.html>) to gain access to these repositories. USENIX allows access to most of their collection without a password. If you need to access a collection that is less than a year old then a member password may be needed. Contact the instructor to obtain the WPI password.

Articles are also available in the journals and proceedings themselves. The following lists sources that are usually, if not always, systems related.

- ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems (published as part of Performance Evaluation Review)

- ACM SIGOPS Conference

- ACM Symposium on Operating Systems Principles (published as part of Operating System Review)

- ACM Transactions on Computer Systems

- HotNets: Workshop on Hot Topics in Networking (ACM)

- IEEE Proceedings of the International Conference on Distributed Computing Systems
- Operating Systems Review (ACM)

- Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ACM)

- USENIX Technical Conference

- USENIX Symposium on Operating Systems Design and Implementation

- USENIX Workshop on Hot Topics in Operating Systems

- USENIX Symposium on Networked Systems Design and Implementation

Articles in these journals and conference proceedings are sometimes systems related.

- ACM Computing Surveys

- Communications of the ACM

- Computer Communications

- Computer Networks

- IEEE Computer

- IEEE International Symposium on Computer Architecture\*

International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)  
ACM Transactions on Internet Technology  
IEEE Internet Computing  
IEEE Network  
IEEE/ACM Transactions on Networking  
IEEE Transactions on Software Engineering  
IEEE Transactions on Computers  
IEEE Transactions on Parallel and Distributed Systems  
International World Wide Web Conference  
Journal of Parallel and Distributed Computing  
Performance Evaluation  
Performance Evaluation Review (ACM)  
SIGCOMM Symposium on Communications, Architectures and Protocols (ACM) (published as part of Computer Communication Review)  
Software Practice and Experience

Other journals and conference proceedings not listed here may also contain relevant material.

## References

- [1] Atul Adya, Paramvir Bahl, and Lili Qiu. Characterizing and analyzing alert and browse services of mobile clients. In *Proceedings of the 2002 USENIX Annual Technical Conference*, Monterey, CA USA, June 2002. USENIX Association.
- [2] Atul Adya, William J. Bolosky, Miguel Castro, Gerald Cermak, Ronnie Chaiken, John R. Douceur, Jon Howell, Jacob R. Lorch, Marvin Theimer, and Roger P. Wattenhofer. Farsite: Federated, available, and reliable storage for an incompletely trusted environment. In *Symposium on Operating Systems Design and Implementation*, Boston, MA, December 2002.
- [3] Thomas E. Anderson, Brian N. Bershad, Edward D. Lazowska, and Henry M. Levy. Scheduler activations: Effective kernel support for the user-level management of parallelism. *ACM Transactions on Computer Systems*, 10(1):53–79, February 1992.
- [4] Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roeselli, and Randolph Y. Wang. Serverless network file systems. *ACM Transactions on Computer Systems*, 14(1):41–79, February 1996.

- [5] Thomas E. Anderson, Henry M. Levy, Brian N. Bershad, and Edward D. Lazowska. The interaction of architecture and operating system design. In *Proceedings of the Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 108–120, April 1991.
- [6] Siddhartha Annapureddy, Michael Freedman, and David Mazieres. Shark: Scaling file servers via cooperative caching. In *USENIX Symposium on Networked Systems Design and Implementation*, 2005.
- [7] G. Back, P. Tullmann, L. Stoller, W.C. Hsieh, and J. Lepreau. Techniques for the design of java operating systems. In *Proceedings of 2000 USENIX Technical Conference*, San Diego, California, June 2000.
- [8] Gaurav Banga and Peter Druschel. Measuring the capacity of a web server. In *USENIX Symposium on Internet Technology and Systems*, Monterey, California, USA, December 1997. USENIX Association.
- [9] Gaurav Banga, Peter Druschel, and Jeffrey C. Mogul. Better operating system features for faster network servers. In *Workshop on Internet Server Performance*, Madison, Wisconsin USA, June 1998.
- [10] Gaurav Banga, Jeffrey C. Mogul, and Peter Druschel. A scalable and explicit event delivery mechanism for UNIX. In *Proceedings of 1999 USENIX Technical Conference*, Monterey, California, June 1999.
- [11] Paul Barham, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer, Ian Pratt, and Andrew Warfield. Xen and the art of virtualization. In *Proceedings of the Symposium on Operating Systems Principles*, Bolton Landing, NY USA, October 2003.
- [12] Andy Bavier, Larry Peterson, Mike Wawrzoniak, Scott Karlin, Tammo Spalink, Timothy Roscoe, David Culler, Brent Chun, and Mic Bowman. Operating system support for planetary-scale network services. In *USENIX Symposium on Networked Systems Design and Implementation*, San Francisco, CA, USA, March 2004.
- [13] Tim Berners-Lee, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. The World-Wide Web. *Communications of the ACM*, 37(8):76–82, August 1994.
- [14] Brian N. Bershad, Thomas E. Anderson, Edward D. Lazowska, and Henry M. Levy. Lightweight remote procedure call. *ACM Transactions on Computer Systems*, 8(1):37–55, February 1990.
- [15] Kenneth P. Birman. The process group approach to reliable distributed computing. *Communications of the ACM*, 36(12):37–53,103, December 1993.
- [16] Andrew D. Birrell, Roy Levin, Roger M. Needham, and Michael D. Schroeder. Grapevine: An exercise in distributed computing. *Communications of the ACM*, 25(4):260–274, April 1982.

- [17] Andrew D. Birrell and Bruce Jay Nelson. Implementing remote procedure calls. *ACM Transactions on Computer Systems*, 2(1):39–59, February 1984.
- [18] David L. Black, David B. Golub, Daniel P. Julin, Richard F. Rashid, Richard P. Draves, Randall W. Dean, Alessandro Forin, Joseph Barrera, Hideyuki Tokuda, Gerald Malan, and David Bohman. Microkernel operating system architecture and Mach. In *Proceedings of the Micro-kernels and Other Kernel Architectures Workshop*, pages 11–30. USENIX Association, April 1992.
- [19] Mic Bowman, Larry L. Peterson, and Andrey Yeatts. Univers: An attribute-based name server. *Software–Practice and Experience*, 20(4):403–424, April 1990.
- [20] Tim Brecht, David Pariag, and Louay Gammo. accept()able strategies for improving web server performance. In *Proceedings of the USENIX Annual Technical Conference*, Boston, MA USA, June 2004.
- [21] B. Brock, G. Carpenter, E. Chiprout, E. Elnozahy, M. Dean, D. Glasco, J. Peterson, R. Rajamony, F. Rawson, R. Rockhold, and A. Zimmerman. Windows NT in a cc-NUMA system. In *Proceedings of the 3rd USENIX Windows NT Symposium*, Seattle, Washington, July 1999.
- [22] Edouard Bugnion, Scott Devine, and Mendel Rosenblum. Disco: Running commodity operating systems on scalable multiprocessors. In *Proceedings of the 16th ACM Symposium on Operating System Principles*, Saint-Malo, France, October 1997. SIGOPS.
- [23] Pei Cao and Sandy Irani. Cost-aware WWW proxy caching algorithms. In *USENIX Symposium on Internet Technology and Systems*, Monterey, California, USA, December 1997. USENIX Association.
- [24] Abhishek Chandra and David Mosberger. Scalability of Linux event-dispatch mechanisms. In *Proceedings of the 2001 USENIX Annual Technical Conference*, Boston, MA USA, June 2001.
- [25] J. Bradley Chen, Yasuhiro Endo, Kee Chan, David Maxières, Antonio Dias, Margo Seltzer, and Michael D. Smith. The measured performance of personal computer operating systems. *ACM Transactions on Computer Systems*, 14(1):3–40, February 1996.
- [26] David R. Cheriton. The V distributed system. *Communications of the ACM*, 31(3):314–333, March 1988.
- [27] David D. Clark. The structuring of systems using upcalls. In *Proceedings of the Tenth ACM Symposium on Operating Systems Principles*, pages 171–180. ACM, December 1985. Available in ACM Operating Systems Review, Vol. 19, No. 5.
- [28] Partha Dasgupta, Richard J. LeBlanc Jr., and William F. Appelbe. The Clouds distributed operating system: Functional description, implementation details and related work. In *Proceedings of the 8th International Conference on Distributed Computing Systems*, pages 2–9, June 1988.

- [29] Peter J. Denning. The working set model for program behavior. *Communications of the ACM*, 11(5):323–333, May 1968.
- [30] Peter J. Denning. Computer viruses. *American Scientist*, 76(3), May-June 1988. Reprinted in *Computers Under Attack Intruders, Worms, and Viruses*, Peter J. Denning, Editor.
- [31] 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.
- [32] Peter Druschel. Operating system support for high-speed communication. *Communications of the ACM*, 39(9):41–51, September 1996.
- [33] Derek L. Eager, Edward D. Lazowska, and John Zahorjan. Adaptive load sharing in homogeneous distributed systems. *IEEE Transactions on Software Engineering*, SE-12(5):662–675, May 1986.
- [34] Yasuhiro Endo, Zheng Wang, J. Bradley Chen, and Margo Seltzer. Using latency to evaluate interactive system performance. In *Second Symposium on Operating Systems Design and Implementation (OSDI '96) Proceedings*, pages 185–199, October 1996.
- [35] Yoav Etsion, Dan Tsafir, and Dror G. Feitelson. Effects of clock resolution on the scheduling of interactive and soft real-time processes. In *Proceedings of the ACM SIGMETRICS Conference*, San Diego, CA, June 2003.
- [36] Robert S. Fabry. Capability-based addressing. *Communications of the ACM*, 17(7):403–412, July 1974.
- [37] James FitzGibbon and Tim Strike. Distributed computing: Moving from CGI to CORBA. In *Proceedings of 2000 USENIX Technical Conference*, San Diego, California, June 2000.
- [38] Jason Flinn and M. Satyanarayanan. Managing battery lifetime with energy-aware adaptation. *ACM Transactions on Computer Systems*, 22(2):137–179, May 2004.
- [39] Eran Gabber, Christopher Small, John Bruno, Jose Brustoloni, and Avi Silberschatz. The pebble component-based operating system. In *Proceedings of 1999 USENIX Technical Conference*, Monterey, California, June 1999.
- [40] James Gettys. The future is coming: Where the X window system should go. In *Proceedings of the Freenix Track 2002 USENIX Annual Technical Conference*, Monterey, CA USA, June 2002.
- [41] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The Google file system. In *Proceedings of the Symposium on Operating Systems Principles*, Bolton Landing, NY USA, October 2003.

- [42] Ashvin Goel, Luca Abeni, Charles Krasic, Jim Snow, and Jonathan Walpole. Supporting time-sensitive applications on a commodity OS. In *Symposium on Operating Systems Design and Implementation*, Boston, MA, December 2002.
- [43] Burra Gopal and Udi Manber. Integrating content-based access mechanisms with hierarchical file systems. In *Symposium on Operating Systems Design and Implementation*, New Orleans, LA, USA, February 1999.
- [44] Krishna Gummadi, Richard Dunn, Stefan Saroiu, Steven Gribble, Henry Levy, and John Zahorjan. Measurement, modeling and analysis of a peer-to-peer file-sharing workload. In *Proceedings of the Symposium on Operating Systems Principles*, Bolton Landing, NY USA, October 2003.
- [45] Hermann Hartig and Michael Hohmuth. The performance of micro-kernel-based systems. In *Proceedings of the 16th ACM Symposium on Operating System Principles*, Saint-Malo, France, October 1997. SIGOPS.
- [46] John Heidemann and Dhaval Shah. Location-aware scheduling with minimal infrastructure. In *Proceedings of 2000 USENIX Technical Conference*, San Diego, California, June 2000.
- [47] Mark D. Hill, James R. Larus, Steven K. Reinhardt, and David A. Wood. Cooperative shared memory: Software and hardware for scalable multiprocessors. *ACM Transactions on Computer Systems*, 11(4):300–318, November 1993.
- [48] Norman C. Hutchinson and Larry L. Peterson. The x-kernel: An architecture for implementing network protocols. *IEEE Transactions on Software Engineering*, 17(1):64–76, January 1991.
- [49] S. Ioannidis, S. Bellovin, and J. Smith. Sub-operating systems: A new approach to application security. In *Proceedings of the 10th ACM SIGOPS European Workshop*, Saint-Emilion, France, September 2002. ACM.
- [50] Brad Curtis Johnson. A distributed computing environment framework: An OSF perspective. Technical report, Open Software Foundation, January 1992.
- [51] Michael B. Jones and John Regehr. The problems you’re having may not be the problems you think you’re having: Results from a latency study of Windows NT. In *Symposium on Operating Systems Design and Implementation*, New Orleans, LA, USA, February 1999.
- [52] M. F. Kaashoek, R. van Renesse, H. van Staveren, and A. Tanenbaum. FLIP: An inter-network protocol for supporting distributed systems. *ACM Transactions on Computer Systems*, 11(1):73–106, 1993.
- [53] M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector Briceno, Russell Hunt, David Mazieres, and Tom Pinckney. Application performance and flexibility on exokernel systems. In *Proceedings of the 16th ACM Symposium on Operating System Principles*, Saint-Malo, France, October 1997. SIGOPS.

- [54] Charles Killian, James Anderson, Ranjit Jhala, and Amin Vahdat. Life, death, and the critical transition: Finding liveness bugs in systems code. In *USENIX Symposium on Networked Systems Design and Implementation*, Cambridge, MA USA, 2007.
- [55] Balachander Krishnamurthy and Craig E. Wills. Study of piggyback cache validation for proxy caches in the World Wide Web. In *USENIX Symposium on Internet Technology and Systems*, Monterey, California, USA, December 1997. USENIX Association.
- [56] Balachander Krishnamurthy and Craig E. Wills. Piggyback server invalidation for proxy cache coherency. In *Seventh International World Wide Web Conference*, Brisbane, Australia, April 1998.
- [57] Thomas M. Kroeger and Darrell E. Long. Design and implementation of a predictive file prefetching algorithm. In *Proceedings of the 2001 USENIX Annual Technical Conference*, Boston, MA USA, June 2001.
- [58] Albert Lai and Jason Nieh. Limits of wide-area thin-client computing. In *Proceedings of the ACM SIGMETRICS 2002 Conference*, Marina Del Ray, CA, June 2002. ACM.
- [59] Kevin Lai and Mary Baker. A performance comparison of unix operating systems on the pentium. In *Proceedings of the USENIX Technical Conference*, pages 265–277. USENIX Association, January 1996.
- [60] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, pages 558–565, July 1978.
- [61] Butler W. Lampson. Hints for computer system design. In *Proceedings of the 9th Symposium on Operating Systems Principles*, pages 33–48, October 1983. Reprinted in *Operating System Review* 19(5).
- [62] Richard P. LaRowe, Jr. and Carla Schlatter Ellis. Page placement policies for NUMA multiprocessors. *Journal of Parallel and Distributed Computing*, (11):112–129, 1991.
- [63] Richard P. LaRowe Jr., Carla Schlatter Ellis, and Laurence S. Kaplan. The robustness of NUMA memory management. In *Proceedings of the 13th ACM Symposium on Operating Systems Principles*, pages 137–151, October 1991.
- [64] James R. Larus and Michael Parkes. Using cohort scheduling to enhance server performance. In *Proceedings of the 2002 USENIX Annual Technical Conference*, Monterey, CA USA, June 2002. USENIX Association.
- [65] Hugh C. Lauer and Roger M. Needham. On the duality of operating system structures. In *Proceedings Second International Symposium on Operating Systems*. IRIA, October 1978. Reprinted in *Operating Systems Review*, 13(2) April 1979, pp. 3–19.
- [66] Roy Levin and David D. Redell. An evaluation of the Ninth SOSP Submissions. In *Proceedings of the Ninth ACM Symposium On Operating Systems Principles*, pages 35–40. ACM, 1983. Reprinted ACM SIGGRAPH Computer Graphics, 22(5) October 1988.

- [67] Philip Levis, Sam Madden, David Gay, Joseph Polastre, Robert Szewczyk, Alec Woo, Eric Brewer, and David Culler. The emergence of networking abstractions and techniques in TinyOS. In *USENIX Symposium on Networked Systems Design and Implementation*, San Francisco, CA, USA, March 2004.
- [68] Eliezer Levy and Abraham Silberschatz. Distributed file systems: Concepts and examples. *ACM Computing Surveys*, 22(4):321–374, December 1990.
- [69] Jochen Liedtke. Toward real microkernels. *Communications of the ACM*, 39(9):70–77, September 1996.
- [70] X. Liu, C. Kreitz, R. van Renesse, J. Hickey, M. Hayden, K. Birman, and R. Constable. Building reliable, high-performance communication systems from components. In *Proceedings of the 17th Symposium on Operating Systems Principles*, Kiawah Island Resort, SC USA, December 1999.
- [71] Cathy McCann, Raj Vaswani, and John Zahorjan. A dynamic processor allocation policy for multiprogrammed shared-memory multiprocessors. *ACM Transactions on Computer Systems*, 11(2):146–178, May 1993.
- [72] Jeffrey C. Mogul. The case for persistent-connection HTTP. In *Proceedings of the ACM SIGCOMM '95 Conference*. ACM, August 1995.
- [73] Jeffrey C. Mogul. Brittle metrics in operating systems research. In *Workshop on Hot Topics in Operating Systems*, Rio Rico, AZ, USA, March 1999.
- [74] Sape J. Mullender, Guido van Rossum, Andrew S. Tanenbaum, Robbert van Renesse, and Hans van Staveren. Amoeba a distributed operating system for the 1990s. *IEEE Computer*, pages 44–53, May 1990.
- [75] Athicha Muthitacharoen, Robert Morris, Thomer M. Gil, and Benjie Chen. Ivy: A read/write peer-to-peer file system. In *Symposium on Operating Systems Design and Implementation*, Boston, MA, December 2002.
- [76] Michael N. Nelson, Brent B. Welch, and John K. Ousterhout. Caching in the Sprite network file system. *ACM Transactions on Computer Systems*, 6(1):134–154, February 1988.
- [77] Henrik Frystyk Nielsen, Jim Gettys, Anselm Baird-Smith, Eric Prud'hommeaux, Hikon Lie, and Chris Lilley. Network performance effects of HTTP/1.1, CSS1, and PNG. In *Proceedings of the ACM SIGCOMM '97 Conference*. ACM, September 1997.
- [78] H. Orman, E. Menze, S. O'Malley, and L. Peterson. A fast and general implementation of Mach IPC in a network. In *Proceedings of the Third USENIX Mach Conference*. USENIX Association, April 1993.
- [79] John K. Ousterhaut, Andrew R. Chersonson, Frederick Douglass, Michael N. Nelson, and Brent B. Welch. The Sprite network operating system. *IEEE Computer*, pages 23–36, February 1988.

- [80] David Pariag, Tim Brecht, Ashif Harji, Peter Buhr, Amol Shukla, and David R. Cheriton. Comparing the performance of web server architectures. In *EuroSys '07: Proceedings of the ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, pages 231–243, New York, NY, USA, 2007. ACM.
- [81] KyoungSoo Park and Vivek Pai. Deploying large file transfer on an HTTP content distributed network. In *Proceedings of the USENIX First Workshop on Real, Large Distributed Systems*, San Francisco, CA USA, December 2004.
- [82] Craig Partridge. Protocols for high-speed networks: Some questions and a few answers. *Computer Networks and ISDN Systems*, pages 1018–1028, April 1993.
- [83] Larry Peterson, Andy Bavier, Marc E. Fiuczynski, and Steve Muir. Experiences building planetlab. In *OSDI '06: Proceedings of the 7th symposium on Operating systems design and implementation*, pages 351–366, Berkeley, CA, USA, 2006. USENIX Association.
- [84] Michael Piatek, Tomas Isdal, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. Do incentives build robustness in bittorrent? In *USENIX Symposium on Networked Systems Design and Implementation*, Cambridge, MA USA, 2007.
- [85] Jelica Protić, Milo Tomašević, and Veljko Milutinović. Distributed shared memory: Concepts and systems. *IEEE Parallel & Distributed Technology*, 4(2):63–79, Summer 1996.
- [86] Venugopalan Ramasubramanian, Ryan Peterson, and Emin Gun Sirer. Corona: A high performance publish-subscribe system for the world wide web. In *USENIX Symposium on Networked Systems Design and Implementation*, San Jose, CA USA, 2006.
- [87] Robbert Van Renesse, Kenneth Birman, and Werner Vogels. Astrolabe: A robust and scalable technology for distributed system monitoring, management and data mining. *ACM Transactions on Computer Systems*, 21(2):164–206, May 2003.
- [88] John Reumann, Ashish Mehra, Kang G. Shin, and Dilip Kandlur. Virtual services: A new abstraction for server consolidation. In *Proceedings of 2000 USENIX Technical Conference*, San Diego, California, June 2000.
- [89] Matei Ripeanu, Mic Bowman, Jeffrey S. Chase, Ian Foster, and Milan Milenkovic. Globus and planetlab resource management solutions compared. In *HPDC '04: Proceedings of the 13th IEEE International Symposium on High Performance Distributed Computing*, pages 246–255, Washington, DC, USA, 2004. IEEE Computer Society.
- [90] Rodrigo Rodrigues, Barbara Liskov, and Liuba Shrira. The design of a robust peer-to-peer system. In *Proceedings of the 10th ACM SIGOPS European Workshop*, Saint-Emilion, France, September 2002. ACM.
- [91] Drew Roselli, Jacob R. Lorch, and Thomas E. Anderson. A comparison of file system workloads. In *Proceedings of 2000 USENIX Technical Conference*, San Diego, California, June 2000.

- [92] Marcell Rosu, C. Michael Olsen, Chandra Narayanaswami, and Lu Luo. PAWP: A power aware web proxy for wireless LAN clients. In *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, 2004.
- [93] M. Rozier, V. Abrossimov, F. Armand, I. Boule, M. Gien, M. Guillemont, F. Herrmann, C. Kaiser, S. Langlois, P. Léonard, and W. Neuhauser. Overview of the CHORUS distributed operating systems. Technical report, Chorus systèmes, April 1990. This paper is a revised update version of an article published in “*Computing Systems*,” *The Journal of the Usenix Association*, 1(4).
- [94] Yaoping Ruan and Vivek Pai. Making the ”box” transparent: System call performance as a first-class result. In *Proceedings of the USENIX Annual Technical Conference*, Boston, MA USA, June 2004.
- [95] Yaoping Ruan and Vivek Pai. Understanding and addressing blocking-induced network server latency. In *Proceedings of the USENIX Annual Technical Conference*, Boston, MA USA, May/June 2006.
- [96] James D. Salehi, James F. Kurose, and Don Towsley. The effectiveness of affinity-based scheduling on multiprocessor network protocol processing. *IEEE/ACM Transactions on Networking*, 4(4):516–530, August 1996.
- [97] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in systems design. *ACM Transactions on Computer Systems*, 2(4):277–288, November 1984.
- [98] Jerome H. Saltzer and John W. Gintell. The instrumentation of Multics. *Communications of the ACM*, 13(8):495–500, August 1970.
- [99] R. Sandberg, D. Goldberg, S. Kleiman, Dan Walsh, and Bob Lyon. Design and implementation of the Sun Network File System. In *Proceedings of the Summer USENIX Conference*, pages 119–130. USENIX Association, June 1985.
- [100] Douglas J. Santry, Michael J. Feeley, Norman C. Hutchinson, and Alistair C. Veitch. Elephant: The file system that never forgets. In *Symposium on Operating Systems Design and Implementation*, New Orleans, LA, USA, February 1999.
- [101] Stefan Saroiu, Krishna P. Gummadi, Henry M. Levy, and Steven D. Gribble. An analysis of internet content delivery mechanisms. In *Symposium on Operating Systems Design and Implementation*, Boston, MA, December 2002.
- [102] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. Exploring the design space of distributed and peer-to-peer systems: Comparing the web, triad, and chord/cfs. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, Cambridge, MA USA, March 2002.
- [103] Mahadev Satyanarayanan. Scalable, secure, and highly available distributed file access. *IEEE Computer*, 23(5):9–22, May 1990.

- [104] Bianca Schroeder, Adam Wierman, and Mor Harchol-Balter. Open versus closed: A cautionary tale. In *USENIX Symposium on Networked Systems Design and Implementation*, San Jose, CA USA, 2006.
- [105] M. Schroeder, A. Birrell, and R. Needham. Experience with Grapevine: The growth of a distributed system. *ACM Transactions on Computer Systems*, 2(1):3–23, February 1984.
- [106] John F. Shoch. Inter-network naming, addressing, and routing. In *Proceedings IEEE COMPCON*, pages 72–79, Fall 1978.
- [107] John F. Shoch and Jon A. Hupp. The “worm” programs - early experience with a distributed computation. *Communications of the ACM*, 25(3):172–180, March 1982.
- [108] Mukesh Singhal. A heuristically-aided algorithm for mutual exclusion in distributed systems. *IEEE Transactions on Computers*, pages 651–662, May 1989.
- [109] Alan Jay Smith. The task of the referee. *IEEE Computer*, 23(4):65–73, April 1990.
- [110] Stephen Soltész, Herbert Pötzl, Marc E. Fiuczynski, Andy Bavier, and Larry Peterson. Container-based operating system virtualization: a scalable, high-performance alternative to hypervisors. In *EuroSys '07: Proceedings of the ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, pages 275–287, New York, NY, USA, 2007. ACM.
- [111] Craig Soules and Gregory Ganger. Why can't i find my files? new methods for automating attribute assignment. In *USENIX Workshop on Hot Topics in Operating Systems*, Lihue, HI, USA, May 2003.
- [112] Mirjana Spasojevic and M. Satyanarayanan. An emperical study of a wide-area distributed file system. *ACM Transactions on Computer Systems*, 14(2):200–222, May 1996.
- [113] Jing Su, Alvin Chin, Anna Popivanova, Ashvin Goel, and Eyal de Lara. User mobility for opportunistic ad-hoc networking. In *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, 2004.
- [114] Michael M. Swift, Brian N. Bershad, and Henry M. Levy. Improving the reliability of commodity operating systems. In *Proceedings of the Symposium on Operating Systems Principles*, Bolton Landing, NY USA, October 2003.
- [115] Andrew S. Tanenbaum and Robbert van Renesse. Distributed operating systems. *ACM Computing Surveys*, 17(4):419–470, December 1985.
- [116] Ken Thompson. Reflections on trusting trust. *Communications of the ACM*, 27(8):761–763, August 1984.

- [117] Raj Vaswani and John Zahorjan. The implications of cache affinity on processor scheduling for multiprogrammed, shared memory multiprocessors. In *Proceedings of the 13th ACM Symposium on Operating Systems Principles*, pages 26–40, October 1991.
- [118] Rob von Behren, Eric Brewer, Nikita Borisov, Mike Chen, Matt Welsh, David Culler, Josh MacDonald, Jeremy Lau, and Steven D. Gribble. Ninja: A framework for network services. In *Proceedings of the 2002 USENIX Annual Technical Conference*, Monterey, CA USA, June 2002. USENIX Association.
- [119] B. Walker, G. Popek, R. English, C. Kline, and G. Thiel. The Locus distributed operating system. In *Proceedings of the 9th ACM Symposium on Operating Systems Principles*, pages 49–70, October 1983.
- [120] Ian H. Witten. Computer (in)security: Infiltrating open systems. In Peter J. Denning, editor, *Computers Under Attack Intruders, Worms, and Viruses*, chapter 5, pages 105–142. Addison Wesley, 1990. Originally published in *ABACUS*, Vol. 4, No. 4, Summer 1987.
- [121] Timothy Wood, Prashant Shenoy, Arun Venkataramani, and Mazin Yousif. Black-box and gray-box strategies for virtual machine migration. In *USENIX Symposium on Networked Systems Design and Implementation*, Cambridge, MA USA, 2007.
- [122] W. Wulf, E. Cohen, W. Corwin, A. Jones, R. Levin, C. Pierson, and F. Pollack. HYDRA: The kernel of a multiprocessor operating system. *Communications of the ACM*, 17(6):337–345, June 1974.
- [123] Michael Young, Avadis Tevanian, Richard Rashid, David Golub, Jeffrey Eppinger, Jonathan Chew, William Bolosky, David Black, and Rober Baron. The duality of memory and communications in the implementation of a multiprocessor operating system. In *Proceedings of the 11th ACM Symposium on Operating Systems Principles*, pages 63–76, 1987. Reprinted in *Operating System Review* 21(5).
- [124] Aydan Yumerefendi, Benjamin Mickle, and Landon Cox. Tightlip: Keeping applications from spilling the beans. In *USENIX Symposium on Networked Systems Design and Implementation*, Cambridge, MA USA, 2007.