

CS555 Syllabus—Spring 1999

John Heidemann

March 5, 1999

Class begins January 13 and ends May 3, with January 18, February 15, and March 15 and 17th off. May 5 is the stop day. The final is May 10th, 8am-10am.

This syllabus may be updated over the semester. The most recent version can always be found at <http://www.isi.edu/~johnh/WORK/CS555/SP1999/SYLLABUS>.

Optional text: [Coulouris94a] .

- [1. **Coulouris94a**] George Coulouris, Jean Dollimore, and Tim Kindberg. *Distributed Systems: Concepts and Design*. Addison-Wesley, second edition, 1994.

1 Introduction

Class 0 (Jan. 13): Diagnostic exam

Class 1 (Jan. 20): Introduction Overview/Reading: [Hanson89a, Levin83a, BernersLee94a] . Homework 1 given out.

- [2. **Hanson89a**] Michael J. Hanson. Efficient reading of papers in science. Brochure of unknown origin, 1989.

- [3. **Levin83a**] Roy Levin and David D. Redell. An evaluation of the ninth SOSP submissions, or how (and how not) to write a good systems paper. *ACM Operating Systems Review*, 17(3):35–40, July 1983.

- [4. **BernersLee94a**] Tim Berners-Lee, Robert Cailliae, Ari Luotonen, Henrik Frystyk Nielsen, and Arthur Secret. The World-Wide Web. *Communications of the ACM*, 37(8):76–82, August 1994.

Class 2 (Jan. 25): design principles [Saltzer81a, Lampson83a, Clark80a] . (Optional: Text Chapter 2 (Design Goals).)

- [5. **Saltzer81a**] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *Proceedings of the 2nd International Conference on Distributed Computing Systems*, pages 509–512, April 1981.

- [6. **Lampson83a**] Butler Lampson. Hints for computer system design. In *Proceedings of the 9th Symposium on Operating Systems Principles*, pages 33–48, Bretton Woods, New Hampshire, October 1983. ACM.

- [7. **Clark80a**] David Clark and Liba Svobodov. Design of distributed systems supporting local autonomy. In *Proceedings of the 20th IEEE COMPCON*, pages 438–444. IEEE, February 1980.

2 Concurrency

Class 3 (Jan. 27): Monitors and messages [Lampson80a, Hoare78a, Lauer78a] . Homework 1 due, homework 2 given out.

- [8. **Lampson80a**] B.W. Lampson and D.D. Ridell. Experiences with processors and monitors in Mesa. *Communications of the ACM*, 23(2):105–117, February 1980.
- [9. **Hoare78a**] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, August 1978.
- [10. **Lauer78a**] Hugh C. Lauer and Roger M. Needham. On the duality of operating system structures. In *Proceedings of the Second International Symposium on Operating Systems*, pages 408–423. INRIA, October 1978. reprinted in *Operating Systems Review* 13(2), April 1979, pp. 3–19.

Class 4 (Feb. 1): RPC [Birrell84a, Hauser93a] . Optional: Text Chapter 5 (RPC).

- [11. **Birrell84a**] A. Birrell and B. Nelson. Implementing remote procedure calls. *ACM Transactions on Computer Systems*, 2(1):39–59, February 1984.
- [12. **Hauser93a**] Carl Hauser, Christian Jacobi, Marvin Theimer, Brent Welch, and Mark Weiser. Using threads in interactive systems: A case study. In *Proceedings of the 14th Symposium on Operating Systems Principles*, pages 94–105, December 1993.

Class 5 (Feb. 3): DSM [Li86a, Carriero85a] .

- [13. **Li86a**] Kai Li and Paul Hudak. Memory coherence in shared virtual memory systems. In *Proceedings of the Fifth Annual ACM Symposium on Principles of Distributed Computing*, pages 229–239. ACM, August 1986.
- [14. **Carriero85a**] Nicholas Carriero and David Gelernter. The S/Net’s Linda kernel. In *Proceedings of the Tenth Symposium on Operating Systems Principles*, pages 110–129. ACM, December 1985.

Class 6 (Feb. 8): Causality [Lamport78a, Jefferson85a] .

- [15. **Lamport78a**] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.
- [16. **Jefferson85a**] David R. Jefferson. Virtual time. *ACM Transactions on Programming Languages and Systems*, 7(3):404–425, July 1985.

Class 7 (Feb. 10): ISIS and schedulers [Birman93a, Waldspurger94a] .

- [17. **Birman93a**] K. P. Birman. The process group approach to reliable distributed computing. *Communications of the ACM*, 36(12):36–53, December 1993.
- [18. **Waldspurger94a**] Carl A. Waldspurger and William E. Weihl. Lottery scheduling: Flexible proportional-share resource management. In *Proceedings of the First USENIX Symposium on Operating Systems Design and Implementation*, pages 1–11, Monterrey, CA, November 1994. USENIX.

3 Naming/Resource Discovery

Class 8 (Feb. 17): classic naming [Danzig92a, Saltzer82a] . Homework 2 due, homework 3 given out.

- [19. **Danzig92a**] Peter B. Danzig, Katia Obraczka, and Anant Kumar. An analysis of wide-area name server traffic: A study of the domain name system. In *Proceedings of the ACM SIGCOMM '92*, pages 281–292, January 1992.
- [20. **Saltzer82a**] Jermome H. Saltzer. On the naming and binding of network destinations. In *International Symposium on Local Computer Networks*, pages 311–317, April 1982.

Class 9 (Feb. 22): Naming extensions [Pike92a, Sechrest92a, Birrell82a] .

- [21. **Pike92a**] Rob Pike, Dave Presotto, Ken Thompson, Howard Trickey, and Phil Winterbottom. The use of name spaces in Plan 9. In *Proceedings of the 5th ACM SIGOPS European Workshop*, pages 72–76, Mont Saint-Michel, 1992. ACM.
- [22. **Sechrest92a**] Stuart Sechrest and Michael McClennen. Blending hierarchical and attribute-based file naming. In *Proceedings of the 12th International Conference on Distributed Computing Systems*, June 1992.
- [23. **Birrell82a**] 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.

Class 10 (Feb. 24): Internet naming [Neuman92b, Neuman89b, Obraczka93b] .

- [24. **Neuman92b**] B. Clifford Neuman. The Prospero File System: A global file system based on the virtual system model. *Computing Systems*, 5(4):407–432, Fall 1992.
- [25. **Neuman89b**] B. Clifford Neuman. The need for closure in large distributed systems. *Operating System Review*, 23(4):28–30, October 1989.
- [26. **Obraczka93b**] Katia Obraczka, Peter B. Danzig, and Shi-Hao Li. Internet resource discovery services. *IEEE Computer*, 26(9):8–22, September 1993.

4 File Systems

Class 11 (Mar. 1): Physical file systems I [McKusick84a, Rosenblum91a, Seltzer95a] . (Optional: Text Chapter 7 (File Service: A Model).)

- [27. **McKusick84a**] Marshall McKusick, William Joy, Samuel Leffler, and R. Fabry. A fast file system for UNIX. *ACM Transactions on Computer Systems*, 2(3):181–197, August 1984.
- [28. **Rosenblum91a**] Mendel Rosenblum and John K. Ousterhout. The design and implementation of a log-structured file system. In *Proceedings of the 13th Symposium on Operating Systems Principles*, pages 1–15. ACM, October 1991.

- [29. **Seltzer95a**] Margo Seltzer, Keith A. Smith, Hari Balakrishnan, Jacqueline Chang, Sara McMains, and Venkata Padmanabhan. File system logging versus clustering: A performance comparison. In *USENIX Conference Proceedings*, pages 249–264, New Orleans, LA, January 1995. USENIX.

Class 12 (Mar. 3): Physical file systems II [Sweeney96a, Patterson88a] .

- [30. **Sweeney96a**] Adam Sweeney, Doug Doucette, Wei Hu, Curtis Anderson, Mike Nishimoto, and Geoff Peck. Scalability in the XFS file system. In *USENIX Conference Proceedings*, pages 1–14. USENIX, January 1996.
- [31. **Patterson88a**] David A. Patterson, Garth Gibson, and Randy H. Katz. A case for redundant arrays of inexpensive disks (RAID). *Proceedings of the 1988 ACM SIGMOD International Conference on Management of Data*, pages 109–116, June 1988.

Class 13 (Mar. 8): Distributed file-systems I [Sandberg85a, Gray89a, Nelson88a] . (Optional: Text Chapter 8 (File Service: Case Studies).) Homework 3 due, homework 4 given out.

- [32. **Sandberg85a**] Russel Sandberg, David Goldberg, Steve Kleiman, Dan Walsh, and Bob Lyon. Design and implementation of the Sun Network File System. In *USENIX Conference Proceedings*, pages 119–130. USENIX, June 1985.
- [33. **Gray89a**] Cary Gray and David Cheriton. Leases: An efficient fault-tolerant mechanism for distributed file cache consistency. In *Proceedings of the Twelfth Symposium on Operating Systems Principles*, pages 202–210. ACM, December 1989.
- [34. **Nelson88a**] 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.

Midterm: Mar. 10. Paper proposal due. Midterm exam.

(Spring break.)

Class 14 (Mar. 22): Distributed file-systems II [Howard88a, Walker83a, Anderson95a] .

- [35. **Howard88a**] John Howard, Michael Kazar, Sherri Menees, David Nichols, Mahadev Satyanarayanan, Robert Sidebotham, and Michael West. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems*, 6(1):51–81, February 1988.
- [36. **Walker83a**] Bruce Walker, Gerald Popek, Robert English, Charles Kline, and Greg Thiel. The LOCUS distributed operating system. In *Proceedings of the Ninth Symposium on Operating Systems Principles*, pages 49–70. ACM, October 1983.
- [37. **Anderson95a**] Thomas E. Anderson, Michael D. Dahlin, Jeanna M. Neefe, David A. Patterson, Drew S. Roselli, and Randolph Y. Wang. Serverless network file systems. In *Proceedings of the 15th Symposium on Operating Systems Principles*, pages 109–126, Copper Mountain Resort, Colorado, December 1995. ACM.

5 Replication

Class 15 (Mar. 24): Voting, Coda, Ficus [Gifford79a, Kistler92a, Guy90b] .

- [38. **Gifford79a**] David K. Gifford. Weighted voting for replicated data. In *Proceedings of the Seventh Symposium on Operating Systems Principles*, pages 150–162. ACM, December 1979.
- [39. **Kistler92a**] James J. Kistler and Mahadev Satyanarayanan. Disconnected operation in the Coda file system. *ACM Transactions on Computer Systems*, 10(1):3–25, 1992.
- [40. **Guy90b**] Richard G. Guy, John S. Heidemann, Wai Mak, Thomas W. Page, Jr., Gerald J. Popek, and Dieter Rothmeier. Implementation of the Ficus replicated file system. In *USENIX Conference Proceedings*, pages 63–71, Anaheim, CA, June 1990. USENIX.

6 Distributed state

Class 16 (Mar. 29) Distributed state [Chandy85a, Lamport82a]

- [41. **Chandy85a**] K. Mani Chandy and Leslie Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, February 1985.
- [42. **Lamport82a**] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.

7 Security

Class 17 (Mar. 31): Security overview [Needham78a, Voydock83a, Schneier96a] . (Optional: Text Chapter 16 (Security).) Homework 4 due, homework 5 given out.

- [43. **Needham78a**] Roger M. Needham and Michael D. Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12):993–999, December 1978.
- [44. **Voydock83a**] V. L. Voydock and S. T. Kent. Security mechanisms in high-level network protocols. *ACM Computing Surveys*, 15(2):135–171, June 1983.
- [45. **Schneier96a**] Bruce Schneier. Why cryptography is harder than it looks. *Risks-Forum Digest (comp.risks)*, 18(61), 15 November 1996.

Class 18 (Apr. 5): Key distribution, confinement, logic [Neuman94b, Lampson73a, Burrows90a] .

- [46. **Neuman94b**] B. Clifford Neuman and Theodore Ts'o. Kerberos: An authentication service for computer networks. *IEEE Communications Magazine*, pages 33–38, September 1994.
- [47. **Lampson73a**] Butler W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16(10):613–615, October 1973.

[48. **Burrows90a**] Michael Burrows, Martín Abadi, and Roger Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, February 1990.

8 Databases

Class 19 (Apr. 7): Databases and OS [Stonebraker81a, Spector85a] . (Optional: Text Chapter 12 (Shared Data and Transactions), 14 (Distributed Transactions))

[49. **Stonebraker81a**] Michael Stonebraker. Operating system support for database management. *Communications of the ACM*, 24(7):412–418, July 1981.

[50. **Spector85a**] A. Spector, D. Daniels, D. Duchamp, J. Eppinger, and R. Pausch. Distributed transactions for reliable systems. In *Proceedings of the Tenth Symposium on Operating Systems Principles*, pages 127–146, December 1985.

9 Kernels

Class 20 (Apr. 12): microkernels [Black92a, Wulf74a, Liedtke93a] . (Optional: Text Chapter 18.1–6 (Case Studies).) Paper proposals must be accepted by this date.

[51. **Black92a**] 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 USENIX Symposium on Microkernels and Other Kernel Architectures*, pages 11–30, April 1992.

[52. **Wulf74a**] 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.

[53. **Liedtke93a**] Jochen Liedtke. Improving IPC by kernel design. In *Proceedings of the 14th Symposium on Operating Systems Principles*, pages 175–188, Asheville, North Carolina, December 1993. ACM.

Class 21 (Apr. 14): layering [Hutchinson91a, Ritchie84a] .

[54. **Hutchinson91a**] 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.

[55. **Ritchie84a**] Dennis M. Ritchie. A stream input-output system. *AT&T Bell Laboratories Technical Journal*, 63(8):1897–1910, October 1984.

Class 22 (Apr. 19): abstraction [Bershad95a, Bugnion97a] . Homework 5 due, homework 6 given out.

[56. **Bershad95a**] Brian N. Bershad, Stefan Savage, Przemyslaw Pardyak, Emin Gün Sirer, Marc Fiuczynski, David Becker, Susan Eggers, and Craig Chambers. Extensibility, safety and performance in the SPIN operating system. In *Proceedings of the 15th Symposium on Operating Systems Principles*, pages 267–284, Copper Mountain Resort, Colorado, December 1995. ACM.

- [57. **Bugnion97a**] Edouard Bugnion, Scott Devine, and Mendel Rosenblum. Disco: Running commodity operating systems on scalable multiprocessors. In *Proceedings of the 16th Symposium on Operating Systems Principles*, pages 143–156, St. Malo, France, October 1997. ACM.

10 Case Studies and Performance Analysis

Class 23 (Apr. 21): Unix, Plan-9, Condor [Ritchie74a, Pike90a, Litzkow88a] .

- [58. **Ritchie74a**] Dennis M. Ritchie and Ken Thompson. The UNIX time-sharing system. *Communications of the ACM*, 17(7):365–375, October 1974.
- [59. **Pike90a**] Rob Pike, Dave Presotto, Ken Thompson, and Howard Trickey. Plan 9 from Bell Labs. In *Proceedings of the Summer '90 UKUUG Conference*, pages 21–33. UKUUG, July 1990.
- [60. **Litzkow88a**] M. Litzkow, M. Livney, and M. Mutka. Condor—a hunter of idle workstations. In *Proceedings of the 8th International Conference on Distributed Computing Systems*, pages 104–111. IEEE, June 1988.

Class 24 (Apr. 26): Sprite, Amoeba [Ousterhout88a, Tanenbaum90a] .

- [61. **Ousterhout88a**] John K. Ousterhout, Andrew R. Cherenon, Frederick Douglass, Michael N. Nelson, and Brent B. Welch. The Sprite network operating system. *IEEE Computer*, pages 23–36, February 1988.
- [62. **Tanenbaum90a**] Andrew S. Tanenbaum, Robbert van Renesse, Hans van Stavern, Gregory J. Sharp, Sape J. Mullender, Jack Jansen, and Guido van Rossum. Experiences with the Amoeba distributed operating system. *Communications of the ACM*, 33(12):46–63, December 1990.

Class 25 (Apr. 28): Performance studies [Chen96a, Ousterhout90a] .

- [63. **Chen96a**] J. Bradley Chen, Yasuhiro Endo, Kee Chan, David Mazières, 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.
- [64. **Ousterhout90a**] John K. Ousterhout. Why aren't operating systems getting faster as fast as hardware? In *USENIX Conference Proceedings*, pages 247–256. USENIX, June 1990.

Class 26 (May. 3): Performance studies, Athena [Chen93b, Champine90a] . Homework 6 due, paper due.

- [65. **Chen93b**] J. Bradley Chen and Brian N. Bershad. The impact of operating system structure on memory system performance. In *Proceedings of the 13th Symposium on Operating Systems Principles*, pages 120–133. ACM, December 1993.
- [66. **Champine90a**] George A. Champine, Jr. Daniel E. Geer, and William N. Ruh. Project Athena as a distributed computer system. *IEEE Computer*, 23(9):40–50, September 1990.

Apr. 29 (stop day). Optional review (if classroom available and students interested).

Final exam (May 10), 8–10am. (The original syllabus incorrectly listed the final as on May 4.)