

# CS555 Syllabus—Spring 1997

Ted Faber and John Heidemann

April 14, 1997

Class begins January 9 and ends April 29.

Text: [1].

## 1 Introduction

Class 0 (Jan. 9), both: Diagnostic exam

Class 1 (Jan. 14), Faber: Introduction Overview/Reading: [2, 3]. Introduction: [4, 5, 6], Text Chapter 2 (Design Goals).

## 2 Concurrency

Class 2 (Jan. 16), Faber: Monitors and messages [7, 8, 9], for next class: Text Chapter 10 (Time and coordination).

Class 3 (Jan. 21), Faber: Time and RPC [10, 11], Text Chapter 5 (RPC).

Class 4 (Jan. 23), Heidemann: Threads (Memory consistency intro) carry-over: [12], [13, 14].

Class 5 (Jan. 28), Heidemann: DSM carry-over: [15], [16, 17], Text Chapter 17 (DSM). Homework 1 given out. Photographs.

## 3 Distributed State

Class 6 (Jan. 30), Heidemann: CATOCS [19], Text Chapter 11 (Replication), carry-over from DSM: [18]. Moved from next class: [22].

Class 7 (Feb. 4), Faber: Byzantine Agreement, Virtual Synchrony [21, 20], for next class: Text Chapter 9 (Name Services).

## 4 Naming/Resource Discovery

Class 8 (Feb. 6), Faber: Classic naming [23, 24, 25].

Class 9 (Feb. 11), Faber: Naming extensions [26, 27, 28]. Homework 1 due, homework 2 given out.

Class 10 (Feb. 13), Faber: Internet naming [29, 30].

## 5 Security

Class 11 (Feb. 18), Heidemann: Security overview [31, 32, 33], Text Chapter 16 (Security). Homework 1 returned.

Class 12 (Feb. 20), Heidemann: Key distribution, confinement, logic [34, 35, 36]

## 6 File Systems

Class 13 (Feb. 25), Faber: RAID and HSM [37, 39], Text Chapter 7 (File Service: A Model). Homework 2 due.

Class 14 (Feb. 27), Faber: log-structured filing [38, 40].

Catch-up/review: Mar. 4. Homework 2 returned.

**Midterm** (Mar. 6). Paper proposals due.

Class 15 (Mar. 18), Heidemann: Distributed file-systems I (NFS, Leases) [41, 43], Text Chapter 8 (File Service: Case Studies). Homework 3 given out.

Class 16 (Mar. 20), Heidemann: Distributed file-systems II (Sprite, AFS) [42, 44] Midterm returned.

Class 17 (Mar. 25), Heidemann: Distributed file-systems III (Locus, xFS) [45, 46]

## 7 Replication

Class 18 (Mar. 27), Heidemann: Voting, Coda, Ficus [47, 48, 49].

Class 19 (Apr. 1), Heidemann: Lazy consistency, Troupes [50, 51]. Homework 3 due, homework 4 given out.

## 8 Databases

Class 20 (Apr. 3), Faber: Databases and OS [52, 53], Text Chapter 12 (Shared Data and Transactions), 14 (Distributed Transactions)

## 9 Kernels

Class 21 (Apr. 8), Faber: microkernels [54, 55], Text Chapter 18.1–6 (Case Studies) Homework 3 returned.

Class 22 (Apr. 10), Heidemann: layering [57, 58, 59]

## 10 Case Studies and Performance Analysis

Class 23 (Apr. 15), Heidemann: Sprite, Amoeba [61, 64]. Class and professor evaluations. Homework 4 due, homework 5 given out.

Class 24 (Apr. 17), Heidemann: Performance studies [67, 68, 69]

Class 25 (Apr. 22), Faber: Unix, Plan-9, Condor [60, 62, 63]

Class 26 (Apr. 24), Faber: Athena, Emerald [65, 66] Homework 5 due.

Class 27 (Apr. 29): optional review class, subject to student interest and available space. (Originally was [56]; that paper is now optional.)

**final exam** (May 1), 11am.

## References

- [1] George Coulouris, Jean Dollimore, and Tim Kindberg. *Distributed Systems: Concepts and Design*. Addison-Wesley, second edition, 1994.
- [2] Michael J. Hanson. Efficient reading of papers in science. Brochure of unknown origin, 1989.
- [3] 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] 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.
- [5] 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.
- [6] 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.
- [7] B.W. Lampson and D.D. Ridell. Experiences with processors and monitors in Mesa. *Communications of the ACM*, 23(2):105–117, February 1980.

- [8] C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, August 1978.
- [9] Hugh C. Lauer and Roger M. Needham. On the duality of operating system structures. In *Proceedings of the Second Symposium on Operating Systems Principles*, pages 3–19. ACM, October 1978.
- [10] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, July 1978.
- [11] David R. Jefferson. Virtual time. *ACM Transactions on Programming Languages and Systems*, 7(3):404–425, July 1985.
- [12] A. Birrell and B. Nelson. Implementing remote procedure calls. *ACM Transactions on Computer Systems*, 2(1):39–59, February 1984.
- [13] 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.
- [14] 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.
- [15] D. Mosberger. Memory consistency models. *ACM Operating Systems Review*, 27(1):18–26, January 1993.
- [16] 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.
- [17] David L. Black, Avadis Tevanian, Jr., David B. Golub, and Michael W. Young. Locking and reference counting in the mach kernel. In *Proceedings of the 1991 International Conference on Parallel Processing*, volume II, pages 167–173. CRC Press, Inc., August 1991.
- [18] Nicholas Carriero and David Gelernter. The S/Net’s Linda kernel. In *Proceedings of the Tenth Symposium on Operating Systems Principles*, page 160. ACM, December 1985.
- [19] E. Chang. Echo algorithms: depth parallel operations on graphs. *IEEE Transactions on Software Engineering*, 8(4):391–400, July 1982.
- [20] 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.

- [21] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.
- [22] Kenneth P. Birman and Thomas A. Joseph. Exploiting virtual synchrony in distributed systems. In *Proceedings of the Eleventh Symposium on Operating Systems Principles*, pages 123–138. ACM, November 1987.
- [23] Douglas B. Terry, Mark Painter, David W. Riggle, and Songnian Zhou. The Berkeley Internet name domain server. In *USENIX Conference Proceedings*, pages 23–31. USENIX, June 1994.
- [24] Jermome H. Saltzer. On the naming and binding of network destinations. In *International Symposium on Local Computer Networks*, pages 311–317, April 1982.
- [25] B. Clifford Neuman. The need for closure in large distributed systems. *Operating System Review*, 23(4):28–30, October 1989.
- [26] Stuart Sechrest and Michael McClellan. Blending hierarchical and attribute-based file naming. In *Proceedings of the 12th International Conference on Distributed Computing Systems*, June 1992.
- [27] 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.
- [28] 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.
- [29] Katia Obraczka, Peter B. Danzig, and Shi-Hao Li. Internet resource discovery services. *IEEE Computer*, 26(9):8–22, September 1993.
- [30] T. Berners-Lee, R. Cailliau, J-F. Groff, and B. Pollermann. World-Wide Web: An information infrastructure for high-energy physics. In D. Perret-Gallix, editor, *Proceedings of the Workshop on Software Engineering, Artificial Intelligence and Expert Systems for High Energy and Nuclear Physics*. World Scientific, Singapore, January 1992.
- [31] 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.
- [32] V. L. Voydock and S. T. Kent. Security mechanisms in high-level network protocols. *ACM Computing Surveys*, 15(2):135–171, June 1983.
- [33] Bruce Schneier. Why cryptography is harder than it looks. *Risks-Forum Digest (comp.risks)*, 18(61), 15 November 1996.

- [34] B. Clifford Neuman and Theodore Ts'o. Kerberos: An authentication service for computer networks. *IEEE Communications Magazine*, pages 33–38, September 1994.
- [35] Butler W. Lampson. A note on the confinement problem. *Communications of the ACM*, 16(10):613–615, October 1973.
- [36] Michael Burrows, Martin Abadi, and Roger Needham. A logic of authentication. *ACM Transactions on Computer Systems*, 8(1):18–36, February 1990.
- [37] David A. Patterson, Garth Gibson, and Randy H. Katz. A case for redundant arrays of inexpensive disks (RAID). *Proceedings of the ACM SIGMOD*, pages 109–116, June 1988.
- [38] 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.
- [39] Samuel S. Coleman and Richard W. Watson. The emerging paradigm shift in storage system architectures. *Proceedings of the IEEE*, 81(4):607–620, April 1993.
- [40] Margo Seltzer, Keith A. Smith, Hari Balakrishnan, Jacqueline Chang, Sara McMains, and Venkata Padmanabhan. File system logging versus clustering: A performance comparison. In *Proceedings of the USENIX Conference Proceedings*, pages 249–264, New Orleans, LA, January 1995. USENIX.
- [41] 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.
- [42] 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.
- [43] 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.
- [44] Mahadev Satyanarayanan. Scalable, secure, and highly available distributed file access. *IEEE Computer*, 23(5):9–21, May 1990.
- [45] 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.

- [46] 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.
- [47] David K. Gifford. Weighted voting for replicated data. In *Proceedings of the Seventh Symposium on Operating Systems Principles*, pages 150–162. ACM, December 1979.
- [48] James J. Kistler and Mahadev Satyanarayanan. Disconnected operation in the Coda file system. *ACM Transactions on Computer Systems*, 10(1):3–25, 1992.
- [49] 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. University of California, Los Angeles, USENIX, June 1990.
- [50] 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.
- [51] Eric C. Cooper. Replicated distributed programs. In *Proceedings of the Tenth Symposium on Operating Systems Principles*, pages 63–78. ACM, December 1985.
- [52] Michael Stonebraker. Operating system support for database management. *Communications of the ACM*, 24(7):412–418, July 1981.
- [53] 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.
- [54] 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.
- [55] W. Wulf, E. Cohen, W. Corwin, A. Jones, R. Levin, C. Pierson, and F. Pollock. HYDRA: the kernel of a multiprocessor operating system. *Communications of the ACM*, 17(6):337–345, June 1974.
- [56] 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.

- [57] Norman C. Hutchinson, Larry L. Peterson, and Herman Rao. The *x*-Kernel: An open operating system design. *IEEE Operating Systems Technical Committee Newsletter*, 3(3):19–23, Fall 1989.
- [58] Brian N. Bershad, Stefan Savage, Przemyslaw Pardyak, Emin Gün Sirer, Marc Fiuczynski, David Becker, and Craig Chambers Susan Eggers. 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.
- [59] Dennis M. Ritchie. A stream input-output system. *AT&T Bell Laboratories Technical Journal*, 63(8):1897–1910, October 1984.
- [60] Dennis M. Ritchie and Ken Thompson. The UNIX time-sharing system. *Communications of the ACM*, 17(7):365–375, October 1974.
- [61] 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] 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.
- [63] 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.
- [64] 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.
- [65] 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.
- [66] Eric Jul, Henry Levy, Norman Hutchinson, and Andrew Black. Fine-grained mobility in the Emerald system. *ACM Transactions on Computer Systems*, 6(1):109–133, February 1988.
- [67] 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*, page 120–133. ACM, December 1993.
- [68] 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.



- [69] John K. Ousterhout. Why aren't operating systems getting faster as fast as hardware? In *USENIX Conference Proceedings*, pages 247–256. USENIX, June 1990.