

Syllabus for CSCI555 Spring 2011

Class Information

People

Name:	Ted Faber ¹	Sang Su Lee	
Office:	SAL 234 or ISI 1141	SAL 235	
Office Hours:	TuTh 10:50 - 11:50 AM	MW 10 - 11 AM	
email:	<faber@isi.edu>	<sangsl@usc.edu>	
Phone:	310-448-9190	TBA	

Classroom

Tuesdays and Thursdays, 9:30-10:50 AM, OHE136

Description

This class is a discussion of advanced topics in operating systems and distributed systems. Students will read and discuss relevant papers in the field and write a term paper² on an area of operating systems research.

Prerequisites

CS402x or the equivalent as well as passing the department OS screening exam.

Academic Integrity

All students are required to abide by the USC code of Student Conduct³, in particular the elements of that code dealing with academic integrity. Violation of that Code will be dealt with as described in SCAM-PUS.⁴ Understand that I will assign an *F* in CSCI555 for any offense described therein. If you have *any* questions about the responsibilities of either students, faculty, or graders under this policy, contact the instructor or the Office of Student Judicial Affairs and Community Standards.⁵ You can also investigate the excellent publications⁶ that SJACS provides. Even if you are confident that you understand academic integrity issues, I encourage you to read *Trojan Integrity: A Guide for Avoiding Plagiarism*⁷, which is a short review/overview of the issues.

Grading

Homework:	10%
-----------	-----

¹ <http://www.isi.edu/~faber>

² <http://www.isi.edu/~faber/cs555/papers.html>

³ <http://web-app.usc.edu/scampus/university-student-conduct-code/>

⁴ http://web-app.usc.edu/scampus/wp-content/uploads/2007/08/appendix_a.pdf

⁵ <http://www.usc.edu/student-affairs/SJACS/>

⁶ <http://www.usc.edu/student-affairs/SJACS/pages/students/publications.html>

⁷ <http://www.usc.edu/student-affairs/SJACS/forms/tig.pdf>

Term Paper:	35%
Final:	30%
Midterm:	25%

The term paper accounts for a significant fraction of the grade, and has many intermediate deadlines. Allocate your time, and arrange your class schedule accordingly

Should you find an error in the grading of an assignment, bring it first to the attention of the primary grader for that assignment. For homeworks, this will be the TA and for examinations and the term paper that will be the instructor⁸. While the instructor is happy to review and discuss your answers to exam and homework questions in office hours or e-mail, few grading errors (other than arithmetic errors) are found.

Examinations

Exams are closed book, and will cover the lectures, readings and homework assignments. All students *must* take the exams at USC. This includes all remote students except for those at Palmdale, NASA-Ames, San Diego and Phoenix.

- The Midterm exam is scheduled for Thursday 10 Mar 2011 during class time (9:30 - 10:50 AM PST), Room TBA.
- The Final exam is scheduled for Tuesday, 10 May 2011 from 8:00 AM to 10:00 AM PDT. Room is TBA.

Term Paper

The term paper will be a 12-page paper on a current topic in OS research. Details will be given in class. A companion document⁹ describes the paper requirements in detail.

Web Page

The class web page is at <http://www.isi.edu/~faber/cs555>¹⁰ This is the primary venue for announcements made outside of class, so check it often.

Classes

11 Jan Class Goals and Requirements. A general preview of the class structure and goals. Definition of student and instructor responsibilities. Papers on reading papers.[Keshav07][Levin83]

Distributed Systems

13 Jan System Organization. The end-to-end argument[Saltzer81], local autonomy (loosely coupled systems)[Clark80], and Athena[Champine90].

18 Jan Communication. Implementation of monitors[Lampson80] and communicating processes[Hoare78].

20 Jan Communication. Remote Procedure Calls[Birrell84], Linda[Carriero85], and JINI[Waldo99].

⁸ <mailto:faber@isi.edu>

⁹ <http://www.isi.edu/~faber/cs555/papers.html>

¹⁰ <http://www.isi.edu/~faber/cs555>

25 Jan Distributed Shared Memory. Memory coherence models[Li89], Content Addressable Networks[Ratnasamy01], and FreeNet[Clarke00].

27 Jan Names. The naming problem[Saltzer82][Sollins03] and closures[Neuman89]

1 Feb Names. Hierarchical naming[Ritchie74], user-centric naming[Pike92], and Intentional naming[Adjie-Winoto99]

3 Feb Searching. Google[Brin98] and Connections[Soules05]

8 Feb Consistency. The Byzantine Generals Problem[Lamport82]and LOCUS[Walker83].

10 Feb Consistency. Grapevine[Birrell82] and DNS[Mockapetris87].

15 Feb Consistency. Weighted voting[Gifford79] and ACID vs. BASE[Fox97].

17 Feb Consistency. FICUS[Guy90] and Coda[Kistler92].

22 Feb Consistency & Distributed Properties. Logical clock synchronization[Lamport78] and Distributed Sanpshots[Chandy85].

24 Feb Distributed Properties. Virtual Time[Jefferson85] and ISIS[Birman93].

Security

1 Mar Cryptography. Crypto overview[Voydock83] and Needham-Schroeder authentication[Needham78].

3 Mar Authentication. Kerberos[Neunman94] and authentication logic[Burrows90]

8 Mar Access Controls. Flask[Spencer99], confinement[Lampson73], and Plan 9[Cox02].

22 Mar Capabilities. Hydra[Wulf74], Amoeba[Tanenbaum90], and SPIN[Bershad95]

Systems

24 Mar File Systems. Practical design[Lampson83], and the Network File System (NFS)[Sandberg85].

29 Mar File Systems. The Log Structured File System (LFS)[Rosenblum91], LFS performance[Seltzer95].

31 Mar File Systems and Disk Layout. Soft-Updates[McKusick99], and RAID disk systems[Patterson88].

5 Apr More File Systems. XFS[Sweeney96], and the Google File System[Ghemawat03]

7 Apr Microkernels. Mach[Black92], and microkernel performance[Liedtke93].

12 Apr Microkernels/Extensibility. The x-kernel[Hutchinson91], streams[Ritchie84].

14 Apr Schedulers. Lottery Scheduling[Waldspurger94] and Quincy[Isard09]

19 Apr Resource-constrained Systems. Emeralds[Zuberi99] and TinyOS[Hill00].

21 Apr Process Migration and Transparency. Condor[Litzkow88] and Sprite[Ousterhout88].

26 Apr Virtualization. Xen[Barham03] and NetLab[White02].

Bibliography

- [Keshav07] S Keshav, "How to Read a Paper," *ACM SIGCOMM Computer Communications Review*, vol. 37, no. 3, July 2009, 83-84,
- [Levin83] 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*, vol. 17, no. 3, July 1983, 35-40,
- [Saltzer81] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-end Arguments In System Design," *ACM Transactions on Computer Systems*, ACM Press, vol. 2, no. 4, 1984, 277-288,
- [Clark80] David D. Clark and Liba Svobodova, "Design of Distributed Systems Supporting Local Autonomy," 20th IEEE COMPCON, IEEE, February 1980, 438-444,
- [Champine90] George A. Champine, Daniel E. Geer, Jr., and William N. Ruh, "Project Athena as a Distributed Computer System," *IEEE Computer*, IEEE, vol. 23, September 1990, 40-50,
- [Lampson80] B.W. Lampson and D.D. Ridell, "Experiences with Processors and Monitors in Mesa," *Communications of the ACM*, vol. 23, no. 2, February 1980, 105-117,
- [Hoare78] C. A. Hoare, "Communicating Sequential Processes," *Communications of the ACM*, vol. 21, no. 8, August 1978, 666-677,
- [Birrell84] A. Birrell and B. Nelson, "Implementing Remote Procedure Calls," *ACM Transactions on Computer Systems*, ACM, vol. 2, no. 1, February 1984, 39-59,
- [Carriero85] Nicholas Carriero and David Gelernter, "The S/Net's Linda Kernel," *ACM Transactions on Computer Systems*, ACM Press, vol. 4, no. 2, 1986, 110-129,
- [Waldo99] Jim Waldo, "The Jini Architecture for Network-centric Computing," *Communications of the ACM*, ACM, vol. 42, no. 7, July 1999, 76-82,
- [Li89] Kai Li and Paul Hudak, "Memory Coherence in Shared Virtual Memory Systems," *ACM Transactions on Computer Systems*, ACM, vol. 7, no. 4, November 1989, 321-359,
- [Ratnasamy01] Sylvia Ratnasamy, Paul Francis, Mark Handley, Richard Karp, and Scott Schenker, "A scalable content-addressable network," SIGCOMM 2001, ACM Press, 2001, 161-172,
- [Clarke00] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong, "Freenet: A Distributed Anonymous Information Storage Retrieval System," ICSI Workshop on Design Issues in Anonymity and Unobservability, July 2000,
- [Saltzer82] Jermome H. Saltzer, "On the Naming and Binding of Network Destinations," *RFC-1498*, August 1993,
- [Sollins03] Karen R. Sollins, "Effective Design of Naming Systems for Networks," *unpublished draft*, February 2003,
- [Neuman89] B. Clifford Neuman, "The Need for Closure in Large Distributed Systems," *Operating System Review*, ACN, vol. 23, no. 4, October 1989, 28-30,
- [Ritchie74] Dennis M. Ritchie and Ken Thompson, "The UNIX Time-sharing System," *Communications of the ACM*, ACM, vol. 17, no. 7, July 1974, 365-375,
- [Pike92] Rob Pike, Dave Presotto, Ken Thompson, Howard Trickey, and Phil Winterbottom, "The Use of Name Spaces in Plan 9," 5th ACM SIGOPS European Workshop, ACM, 1992, 72-76,

- [Adjie-Winoto99] William Adjie-Winoto, Elliot Schwartz, Hari Balakrishnan, and Jeremy Lilley, "The Design and Implementation of an Intentional Naming System," Proc 17th Symposium on Operating Systems Principles, ACM, December 1999, 186-201,
- [Brin98] Sergey Brin and Lawrence Page, "The Anatomy of a Large-scale Hypertextual Web Search Engine," 7th World Wide Web Conference, April 1998, 107-117,
- [Soules05] Craig A. Soules and Gregory R. Ganger, "Connections: using context to enhance file search," 20th ACM symposium on Operating systems principles, ACM Press, 2005, 119-132,
- [Lamport82] Leslie Lamport, Robert Shostak, and Marshall Pease, "The Byzantine Generals Problem," *ACM Transactions on Programming Languages and Systems*, ACM, vol. 4, no. 3, July 1982, 382-401,
- [Walker83] Bruce Walker, Gerald Popek, Robert English, Charles Kline, and Greg Thiel, "The LOCUS Distributed Operating System," Ninth Symposium on Operating Systems Principles, ACM, October 1983, 49-70,
- [Birrell82] Andrew D. Birrell, Roy Levin, Roger M. Needham, and Michael D. Schroeder, "Grapevine: An Exercise in Distributed Computing," *Communications of the ACM*, ACM, vol. 25, no. 4, April 1982, 260-274,
- [Mockapetris87] Paul Mockapetris, "Domain Names - Concepts And Facilities," *RFC-1034*, November 1987,
- [Gifford79] David K. Gifford, "Weighted Voting for Replicated Data," Seventh Symposium on Operating Systems Principles, ACM, December 1979, 150-162,
- [Fox97] Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, and Paul Gauthier, "Cluster-Based Scalable Network Services," *Proceedings of the 16th Symposium on Operating Systems Principles*, ACM, October 1997, 78-91,
- [Birman93] Kenneth P. Birman, "The Process Group Approach To Reliable Distributed Computing," *Communications of the ACM*, ACM, vol. 36, no. 12, December 1993, 36-53,
- [Guy90] 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," *USENIX Conference Proceeding*, USENIX, June 1990, 63-71,
- [Kistler92] James J. Kistler and Mahadev Satyanarayanan, "Disconnected Operation in the Coda File System," *ACM Transactions on Computer Systems*, ACM, vol. 10, no. 1, 1992, 3-25,
- [Chandy85] K. Mani Chandy and Leslie Lamport, "Distributed Snapshots: Determining Global States of Distributed Systems," *ACM Transactions on Computer Systems*, vol. 3, no. 1, February 1985, 63-75,
- [Lamport78] Leslie Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Communications of the ACM*, vol. 21, no. 7, July 1978, 558-565,
- [Jefferson85] David R. Jefferson, "Virtual Time," *ACM Transactions on Programming Languages and Systems*, ACM, vol. 7, no. 3, July 1985, 404-425,
- [Voydock83] V. L. Voydock and S. T. Kent, "Security Mechanisms in High-Level Network Protocols," *ACM Computing Surveys*, ACM, vol. 15, no. 2, June 1983, 135-171,
- [Needham78] Roger M. Needham and Michael D. Schroeder, "Using Encryption for Authentication in Large Networks of Computers," *Communications of the ACM*, ACM, vol. 21, no. 12, December 1978, 993-999,
- [Neuman94] B. Clifford Neuman and Theodore Ts'o, "Kerberos: An Authentication Service for Computer Networks," *IEEE Communications Magazine*, IEEE, September 1994, 33-38,
- [Burrows90] Michael Burrows, Marti Abadi, and Roger Needham, "A Logic of Authentication," *ACM Transactions on Computer Systems*, ACM, vol. 8, no. 1, February 1990, 18-36,

- [Spencer99] Ray Spencer, Stephen Smalley, Peter Loscocco, Mike Hibler, David Andersen, and Jay Lepreau, "The Flask Security Architecture: System Support for Diverse Security Policies," The Eighth USENIX Security Symposium, USENIX, August 1999, 123-139,
- [Lampson73] Butler W. Lampson, "A Note on the Confinement Problem," *Communications of the ACM*, ACM, vol. 16, no. 10, October 1973, 613-615,
- [Cox02] Russ Cox, Eric Grosse, Rob Pike, Dave Presotto, and Sean Quinlan, "Security in Plan 9," 11th USENIX Security Symposium, USENIX, August 2002,
- [Wulf74] 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*, vol. 17, no. 6, June 1974, 337-345,
- [Tanenbaum90] 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*, vol. 33, no. 12, December 1990, 46-63,
- [Bershad95] Brian N. Bershad, Stefan Savage, Przemyslaw Pardyak, Emin G\ Sirer, Marc Fiuczynski, David Becker, Susan Eggers, and Craig Chambers, "Extensibility, Safety and Performance in the SPIN Operating System," 15th Symposium on Operating Systems Principles, ACM, December 1995, 267-284,
- [Lampson83] B. W. Lampson, "Hints for Computer System Design," 9th Symposium on Operating Systems Principles, ACM, October 1983, 33-48,
- [Sandberg85] Russel Sandberg, David Goldberg, Steve Kleiman, Dan Walsh, and Bob Lyon, "Design and Implementation of the Sun Network File System," USENIX Conference, USENIX, June 1985, 119-130,
- [Rosenblum91] Mendel Rosenblum and John K. Ousterhout, "The design and implementation of a log-structured file system," 13th Symposium on Operating Systems Principles, ACM, October 1991, 1-15,
- [Seltzer95] Margo Seltzer, Keith A. Smith, Hari Balakrishnan, Jacqueline Chang, Sara McMains, and Venkata Padmanabhan, "File System Logging Versus Clustering: A Performance Comparison," USENIX Conference, January 1995, ,
- [McKusick99] Marshall Kirk McKusick and Gregory R. Ganger, "Soft Updates: A Technique for Eliminating Most Synchronous Writes in the Fast Filesystem," FreeNIX Track: 1999 USENIX Annual Technical Conference, USENIX, June 1999, ,
- [Patterson88] 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, June 1988, 109-116,
- [Sweeney96] Adam Sweeney, Doug Doucette, Wei Hu, Curtis Anderson, Mike Nishimoto, and Geoff Peck, "Scalability in the XFS File System," *USENIX Conference Proceedings*, USENIX, January 1996, 1-14,
- [Ghemawat03] Sanjay Ghemawat, Howard Gobioff, and Shun Leung, "The Google file system," Proceedings of the 19th ACM symposium on Operating systems principles, ACM, 2003, 29-43,
- [Black92] David L. Black, David B. Golub, Daniel P. Julin, Richard F. Rashid, Richard P. Draves, Randall W. Dean, Alessandro Florin, Joesph Barrera, Hideyuki Tokuda, Gerald Malan, and David Bohman, "Microkernel Operating System Architecture and Mach," USENIX Workshop on Microkernels and Other Kernel Architectures, April 1992, 11-30,
- [Liedtke93] Jochen Liedtke, "Improving IPC by Kernel Design," 14th Symposium on Operating Systems Principles, ACM, December 1993, 175-188,
- [Hutchinson91] Norman C. Hutchinson and Larry L. Peterson, "The x-Kernel: An architecture for implementing network protocols," *IEEE Transactions on Software Engineering*, IEEE, vol. 17, no. 1, January 1991, 64-76,

- [Ritchie84] Dennis M. Ritchie, "A Stream Input-Output System," *AT&T Bell Laboratories Technical Journal*, vol. 63, October 1984, 1897-1910,
- [Zuberi99] Khawar M. Zuberi, Padmanabhan Pillai, and Kang G. Shin, "EMERALDS: A Small-Memory Real-Time Microkernel," *Proc 17th Symposium on Operating Systems Principles*, ACM, December 1999, 277-291,
- [Hill00] Jason Hill, Robert Szewczyk, Alex Woo, Seth Hollar, David Culler, and Kristofer Pister, "System Architecture Directions for Network Sensors," *9th International Conference on Architectural Support for Programming Languages and Operating Systems*, ACM, November 2000, 93-104,
- [Waldspurger94] Carl Waldspurger and William Weihl, "Lottery Scheduling: Flexible Proportional-Share Resource Management," *USENIX 1994 Operating Systems Design and Implementation Proceedings*, USENIX/ACM, November 1994, 1-11,
- [Isard09] Michael Isard, Vijayan Prabhakaran, Jon Currey, Udi Wieder, Kunal Talwar, and Andrew Goldberg, "Quincy: fair scheduling for distributed computing clusters," *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*, ACM, October 2009, 261-276,
- [Litzkow88] M. Litzkow, M. Livny, and M. Mutka, "Condor - A Hunter of Idle Workstations," *8th International Conference on Distributed Computing System*, IEEE, June 1988, 104-111,
- [Ousterhout88] John K. Ousterhout, Andrew R. Cherenon, Frederick Douglass, Michael N. Nelson, and Brent B. Welch, "The Sprite Network Operating System," *IEEE Computer*, IEEE, February 1988, 23-36,
- [Barham03] 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*, *Proc 19th Symposium on Operating Systems Principles*, ACM, December 2003, pp. 164-177,
- [White02] Brian White, Jay Lepreau, Leigh Stoller, Robert Ricci, Shashi Guruprasad, Mac Newbold, Mike Hibler, Chad Barb, and Abhijeet Joglekar, "An integrated experimental environment for distributed systems and networks," *SIGOPS Oper. Syst. Rev.*, ACM Press, vol. 36, no. SI, 2002, 255-270,
- [Chou01] Andy Chou, Junfeng Yang, Benjamin Chelf, Seth Hallem, and Dawson Engler, "An empirical study of operating systems errors," *Proceedings of the 18th ACM symposium on Operating systems principles*, ACM Press, 2001, 73-88,