

1. Course Information

Staff

Faculty

Butler Lampson NE43-535 547-9580 blampson@microsoft.com

Martin Rinard NE43-620A x8-6922 rinard@lcs.mit.edu

Teaching Assistant

Kyle Jamieson NE43-525 x3-3058 jamieson@mit.edu

Course Secretary

Rachel Allen NE43-620 x3-9620 rachela@cag.lcs.mit.edu

Office Hours

Messrs. Lampson and Rinard will arrange individual appointments. Kyle Jamieson will hold scheduled office hours in the 5th floor lounge of NE43 at a time to be announced. In addition to holding regularly scheduled office hours, the TA will also be available by appointment.

Lectures and handouts

Lectures are held on Tuesdays and Thursdays from 1:00 to 2:30PM in room 4-149. Messrs. Lampson and Rinard will split the lectures. The tentative schedule is at the end of this handout.

The source material for this course is an extensive set of handouts. There are about 400 pages of topic handouts that take the place of a textbook; you will need to study them to do well in the course. Since we don't want to simply repeat the written handouts in class, we will hand out the material for each lecture one week in advance. *We expect you to read the day's handouts before the class and come prepared to ask questions, discuss the material, or follow extensions of it or different ways of approaching the topic.*

Seven research papers supplement the topic handouts. In addition there are 5 problem sets, and the project described below. Solutions for each problem set will be available shortly after the due date.

There is a course Web page, at web.mit.edu/6.826/www. Last year's handouts can be found from this page. Current handouts will be placed on the Web as they are produced, as Postscript (.ps), Acrobat (.pdf), browser (.html), and Word (.rtf) files.

Current handouts will generally be available in lecture. If you miss any in lecture, you can obtain them afterwards from the course secretary. She keeps them in a file cabinet outside her office.

Problem sets

There is a problem set approximately once a week for the first half of the course. Problem sets are handed out on Wednesdays and are due in class the following Wednesday. They normally cover the material discussed in class during the week they are handed out. Delayed submission of the solutions will be penalized, and no solutions will be accepted after Thursday 5:00PM.

Students in the class will be asked to help grade the problem sets. Each week a team of students will work with the TA to grade the week's problems. This takes about 3-4 hours. Each student will probably only have to do it once during the term.

We will try to return the graded problem sets, with solutions, within a week after their due date.

Policy on collaboration

We encourage discussion of the issues in the lectures, readings, and problem sets. However, if you collaborate on problem sets, you must tell us who your collaborators are. And in any case, you must write up all solutions on your own.

Project

During the last half of the course there is a project in which students will work in groups of three or so to apply the methods of the course to their own research projects. Each group will pick a real system, preferably one that some member of the group is actually working on but possibly one from a published paper or from someone else's research, and write:

A specification for it.

A high-level implementation that captures the novel or tricky aspects of the actual implementation.

The abstraction function and key invariants for the correctness of the implementation. This is not optional; if you can't write these things down, you don't understand what you are doing.

Depending on the difficulty of the specification and implementation, the group may also write a correctness proof for the implementation.

Projects may range in style from fairly formal, like handout 18 on consensus, in which the 'real system' is a simple one, to fairly informal (at least by the standards of this course), like the section on copying file systems in handout 7. These two handouts, along with the ones on naming, sequential transactions, concurrent transactions, and caching, are examples of the appropriate size and possible styles of a project.

The result of the project should be a write-up, in the style of one of these handouts. During the last two weeks of the course, each group will give a 25-minute presentation of its results. We have allocated four class periods for these presentations, which means that there will be twelve or fewer groups.

The projects will have five milestones. The purpose of these milestones is not to assign grades, but to make it possible for the instructors to keep track of how the projects are going and give everyone the best possible chance of a successful project

1. We will form the groups around March 1, to give most of the people that will drop the course a chance to do so.
2. Each group will write up a 2-3 page project proposal, present it to one of the instructors around spring break, and get feedback about how appropriate the project is and suggestions on how to carry it out. Any project that seems to be seriously off the rails will have a second proposal meeting a week later.
3. Each group will submit a 5-10 page interim report in the middle of the project period.
4. Each group will give a presentation to the class.
5. Each group will submit a final report, which is due on Friday, May 12, the last day allowed by MIT regulations. Of course you are free to submit it early.

Half the groups will be ‘early’ ones that give their presentations on May 2 or May 4; the other half will be ‘late’ ones that give their presentations on May 9 or May 11. The due dates of proposals and interim reports will be spread out over two weeks in the same way.

Grades

There are no exams. Grades are based 30% on the problem sets, 50% on the project, and 20% on class participation and quality and promptness of grading.

Course mailing list

A mailing list for course announcements—6.826-students@mit.edu—has been set up to include all students and the TA. If you do not receive any email from this mailing list within the first week, check with the TA. Another mailing list, 6.826-staff@mit.edu, sends email to the entire 6.826 staff.

Course Schedule

Date	No	By	HO	Topic	PS out	PS due
Tues., Feb. 1	1	L		Overview. The Spec language. State machine semantics. Examples of specifications and implementations.	1	
				1 Course information		
				2 Background		
				3 Introduction to Spec		
				4 Spec reference manual		
				5 Examples of specs and implementations		
Thurs., Feb. 3	2	L		Specification and implementation for sequential programs. Correctness notions and proofs. Proof methods: abstraction functions and invariants.		
				6 Abstraction functions		
Tues., Feb. 8	3	L		File systems 1: Disks, simple sequential file system, caching, logs for crash recovery.	2	1
				7 Disks and file systems		
Thurs., Feb. 10	4	L		File systems 2: Copying file system.		
Tues., Feb. 15	5	R		Proof methods: History and prophecy variables; abstraction relations.	3	2
				8 History variables		
Thurs., Feb. 17	6	R		Semantics and proofs: Formal sequential semantics of Spec.		
				9 Atomic semantics of Spec		
Tues., Feb. 22				No class, Monday schedule from Presidents’ Day		
Thurs., Feb. 24	7	L		Performance: How to get it, how to analyze it.	4	3
				10 Performance		
				11 Paper: Michael Schroeder and Michael Burrows, Performance of Firefly RPC, <i>ACM Transactions on Computer Systems</i> 8 , 1, February 1990, pp 1-17.		
Tues., Feb. 29	8	L		Naming: Specs, variations, and examples of hierarchical naming.		Form groups
				12 Naming		
				13 Paper: David Gifford et al, Semantic file systems, <i>Proc. 13th ACM Symposium on Operating System Principles</i> , October 1991, pp 16-25.		

Date	No	By	HO	Topic	PS out	PS due
Thurs., Mar. 2	9	R		Concurrency 1: Practical concurrency, easy and hard. Easy concurrency using locks and condition variables. Problems with it: scheduling, deadlock.	5	4
			14	Practical concurrency		
			15	Concurrent disks		
			16	Paper: Andrew Birrell, An Introduction to Programming with Threads, Digital Systems Research Center Report 35, January 1989.		
Tues., Mar. 7	10	R		Concurrency 2: Concurrency in Spec: threads and non-atomic semantics. Big atomic actions. Safety and liveness. Examples of concurrency.		
			17	Formal concurrency		
Thurs., Mar. 9	11	R		Concurrency 3: More examples. Proving correctness of concurrent programs.		5
Tues., Mar. 14	12	R		Concurrency 4: Examples of correctness proofs		
Thurs., Mar. 16	13	L		Distributed consensus in the presence of faults. Paxos algorithm for asynchronous consensus.	Early proposals	
			18	Consensus		
Mar. 18-26				Spring Vacation		
Tues., Mar. 28	14	R		Sequential transactions with caching.		
			19	Sequential transactions		
Thurs., Mar. 30	15	R		Concurrent transactions: Specs for serializability. Ways to implement the specs.	Late proposals	
			20	Concurrent transactions		
Tues., Apr. 4	16	R		Distributed transactions: Commit as a consensus problem. Two-phase commit. Optimizations.	Late interim reports	
			27	Distributed transactions		
Thurs., Apr. 6	17	R		Introduction to distributed systems: Characteristics of distributed systems. Physical, data link, and network layers. Design principles. Networks 1: Links. Point-to-point and broadcast networks.		
			21	Distributed systems		
			22	Paper: Michael Schroeder et al, Autonet: A high-speed, self-configuring local area network, <i>IEEE Journal on Selected Areas in Communications</i> 9 , 8, October 1991, pp 1318-1335.		

Date	No	By	HO	Topic	PS out	PS due
			23	Networks: Links and switches		
Tues., Apr. 11	18	L		Networks 2: Links cont'd: Ethernet. Token Rings. Switches. Implementing switches. Routing. Learning topologies and establishing routes.		
Thurs., Apr. 13	19	L		Networks 3: Network objects and remote procedure call (RPC).		Early interim reports
			24	Network objects		
			25	Paper: Andrew Birrell et al., Network objects, <i>Proc. 14th ACM Symposium on Operating Systems Principles</i> , Asheville, NC, December 1993.		
Tues., Apr. 18				Patriot's Day, no class.		
Thurs., Apr. 20	20	J		Networks 4: Reliable messages. 3-way handshake and clock implementations. TCP as a form of reliable messages.		Late interim reports
			26	Paper: Butler Lampson, Reliable messages and connection establishment. In <i>Distributed Systems</i> , ed. S. Mullender, Addison-Wesley, 1993, pp 251-281.		
Tues., Apr. 25	21	L		Replication and availability: Implementing replicated state machines using consensus. Applications to replicated storage.		
			28	Replication		
			29	Paper: Jim Gray and Andreas Reuter, Fault tolerance, in <i>Transaction Processing: Concepts and Techniques</i> , Morgan Kaufmann, 1993, pp 93-156.		
Thurs., Apr. 27	22	R		Caching: Maintaining coherent memory. Broadcast (snoopy) and directory protocols. Examples: multiprocessors, distributed shared memory, distributed file systems.		
			30	Concurrent caching		
Tues., May 2	23			Early project presentations		
Thurs., May 4	24			Early project presentations		
Tues., May 9	25			Late project presentations		
Thurs., May 11	26			Late project presentations		
Fri., May 12				Final reports due		
May 15-19				Finals week		