CSCI 499: Introduction to Distributed Systems

Units: 4.0
Fall 2024 MonWed 4:00-5:50pm

Location: TBD

Instructor: Harsha V. Madhyastha
Office: SAL 326
Office Hours: 4 hours per week
Contact Info: Post privately on Piazza to contact the instructor.

Course Description

Application deployments spread across many computers offer higher performance, greater fault-tolerance, and better scalability than single-computer systems, but are challenging to develop. In this course, which is intended for Computer Science majors who are in their junior or senior year, students will learn the fundamental techniques that application providers such as Google, Amazon, Facebook, Microsoft, etc. employ in developing the distributed systems on which their planet-scale services are hosted. Topics covered will include various forms of replication, different data consistency models, and impossibility results on the tradeoffs between performance, fault-tolerance, and consistency. Students will get to see how the principles they learn have been used in real-world systems. A substantial amount of learning in this class will be from the five projects that the students will work on.

Learning Objectives

By the end of this course, students will have gained experience with

- a variety of replication strategies, e.g., using logical clocks, primary-backup replication, consensus protocols such as Paxos, etc.
- a number of data consistency semantics, such as linearizability, serializability, eventual consistency
- fundamental tradeoffs between performance, fault-tolerance, and consistency
- approaches for scaling up distributed systems, such as consistent hashing, distributed hash tables
- implementing distributed systems in Golang

Recommended Preparation

CSCI 201 and CSCI 350

Course Notes

Students enrolled in this course will receive a letter grade.
All lectures will be recorded to enable students to review the lectures, but not to use as a substitute for attending lecture.
Lecture slides and recordings, as well as project specifications and homeworks will be posted on a website dedicated for this course.
Piazza will be used for announcements and discussions regarding lectures, projects, and exams.

Technological Proficiency and Hardware/Software Required

To work on the projects, students may use any computer which has an installation of Go (version 1.13 or greater).

Optional Readings and Supplementary Materials

Distributed Systems (3rd Edition), by Maarten Van Steen and Andrew S. Tanenbaum

For every case study discussed in lecture, a corresponding publication will be posted on the course website, which students are encouraged to read.

Description and Assessment of Assignments

Students will work on five projects in this class, all of which will be in the Go programming language.
Students will work in groups of two for all projects, except the first one.

- Project 1: Simple implementation of MapReduce to get familiar with Go
- Project 2: Key,value store using primary/backup replication
- Project 3: Key,value store using Paxos
- Projects 4 and 5: Key,value stores using sharding for scalability

There will also be five homeworks. In each of these, students will answer 2-3 questions which will help them confirm their understanding of the lecture material.

Participation

Class participation will count towards 5% of a student's grade.

Grading Breakdown

Table 1 Grading Breakdown

Assessment Tool (assignments)	Points	% of Grade
Projects		60%
Homeworks		10%
Midterm exam		15%
Final exam		10%
Class participation		5%
TOTAL		100%

Table 2 Course Grading Scale

Letter grade	Corresponding numerical point range
A	93-100
A-	90-92
B+	87-89
B	83-86
B-	80-82
C+	77-79
C	73-76
C-	70-72
D+	67-69
D	63-66
D-	60-62
F	59 and below

Assignment Submission Policy

Students will submit their solutions to the projects through a web form. Projects will be due at 11:59 pm on the due date. Every student will have a total of 2 late days across all projects.

Homeworks and exams will be submitted in-person to the instructor in class.

Grading Timeline

All projects will be graded by an autograder. The autograder typically finishes grading a group's submission in less than an hour.
Scores for homeworks and exams will be released a week after submission.

Attendance

Students are highly encouraged to attend lecture in-person.

Classroom Norms

Lectures will make heavy use of in-class discussions. Students are expected to actively participate in these discussions. **We ask that all students treat each other with respect.**

Course Evaluations

Students will be asked to submit their evaluation of the course both in the middle and at the end of the semester.

Course Schedule

A homework will go out roughly once every three weeks, and each homework will be due a week after it goes out.

Table 3 Course schedule

	Topics/Daily Activities	Readings/ Preparation	Deliverables
Week 1	Course overview and logistics, MapReduce, Clock synchronization		
Week 2	Replicated state machines, Lamport clock, Vector clock		Project 1 out
Week 3	Primary backup replication, external consistency, and developer-transparent replication		
Week 4	Linearizability and other data consistency models		Project 1 due; Project 2 out
Week 5	Consensus protocols such as Paxos and Raft		
Week 6	Consensus-based replicated state machines		Project 2 due; Project 3 out
Week 7	Fundamental tradeoffs between availability, consistency, and performance		
Week 8	Eventual consistency		Project 3 due
Week 9	Midterm exam		
Week 10	Scalability, consistent hashing, and distributed hash tables		Project 4 goes out
Week 11	Amazon Dynamo and Microsoft Azure Storage		
Week 12	Distributed transactions, two-phase commit, and Google Spanner		Project 4 due; Project 5 out
Week 13	Performance at Scale, RAMCloud, Scalable causal consistency		
Week 14	Bitcoin		
Week 15	Current research in distributed systems		Project 5 due
FINAL			Refer to the final exam schedule in the USC *Schedule of Classes* at classes.usc.edu.

Statement on Academic Conduct and Support Systems

Academic Integrity:
The University of Southern California is a learning community committed to developing successful scholars and researchers dedicated to the pursuit of knowledge and the dissemination of ideas. Academic misconduct, which includes any act of dishonesty in the production or submission of academic work, comprises the integrity of the person who commits the act and can impugn the perceived integrity of the entire university community. It stands in opposition to the university's mission to research, educate, and contribute productively to our community and the world.

All students are expected to submit assignments that represent their own original work, and that have been prepared specifically for the course or section for which they have been submitted. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s).

Other violations of academic integrity include, but are not limited to, cheating, plagiarism, fabrication (e.g., falsifying data), collusion, knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

The impact of academic dishonesty is far-reaching and is considered a serious offense against the university. All incidences of academic misconduct will be reported to the Office of Academic Integrity and could result in outcomes such as failure on the assignment, failure in the course, suspension, or even expulsion from the university.

For more information about academic integrity see the student handbook or the Office of Academic Integrity's website, and university policies on Research and Scholarship Misconduct.

Please ask your instructor if you are unsure what constitutes unauthorized assistance on an exam or assignment, or what information requires citation and/or attribution.

Course Content Distribution and Synchronous Session Recordings Policies
USC has policies that prohibit recording and distribution of any synchronous and asynchronous course content outside of the learning environment.

Recording a university class without the express permission of the instructor and announcement to the class, or unless conducted pursuant to an Office of Student Accessibility Services (OSAS) accommodation. Recording can inhibit free discussion in the future, and thus infringe on the academic freedom of other students as well as the instructor. (Living our Unifying Values: The USC Student Handbook, page 13).

Distribution or use of notes, recordings, exams, or other intellectual property, based on university classes or lectures without the express permission of the instructor for purposes other than individual or group study. This includes but is not limited to providing materials for distribution by services publishing course materials. This restriction on unauthorized use also applies to all information, which had been distributed to students or in any way had been displayed for use in relationship to the class, whether obtained in class, via email, on the internet, or via any other media. (Living our Unifying Values: The USC Student Handbook, page 13).

Students and Disability Accommodations:
USC welcomes students with disabilities into all of the University's educational programs. The Office of Student Accessibility Services (OSAS) is responsible for the determination of appropriate accommodations for students who encounter disability-related barriers. Once a student has completed the OSAS process (registration, initial appointment, and submitted documentation) and accommodations are determined to be reasonable and appropriate, a Letter of Accommodation (LOA) will be available to generate for each course. The LOA must be given to each course instructor by the student and followed up with a discussion.

This should be done as early in the semester as possible as accommodations are not retroactive. More information can be found at osas.usc.edu. You may contact OSAS at (213) 740-0776 or via email at osasfrontdesk@usc.edu.

Support Systems:

Counseling and Mental Health - (213) 740-9355 – 24/7 on call
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

988 Suicide and Crisis Lifeline - 988 for both calls and text messages – 24/7 on call
The 988 Suicide and Crisis Lifeline (formerly known as the National Suicide Prevention Lifeline) provides free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week, across the United States. The Lifeline is comprised of a national network of over 200 local crisis centers, combining custom local care and resources with national standards and best practices. The new, shorter phone number makes it easier for people to remember and access mental health crisis services (though the previous 1 (800) 273-8255 number will continue to function indefinitely) and represents a continued commitment to those in crisis.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-9355(WELL) – 24/7 on call
Free and confidential therapy services, workshops, and training for situations related to gender- and power-based harm (including sexual assault, intimate partner violence, and stalking).

Office for Equity, Equal Opportunity, and Title IX (EEO-TIX) - (213) 740-5086
Information about how to get help or help someone affected by harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants.

Reporting Incidents of Bias or Harassment - (213) 740-5086 or (213) 821-8298
Avenue to report incidents of bias, hate crimes, and microaggressions to the Office for Equity, Equal Opportunity, and Title for appropriate investigation, supportive measures, and response.

The Office of Student Accessibility Services (OSAS) - (213) 740-0776
OSAS ensures equal access for students with disabilities through providing academic accommodations and auxiliary aids in accordance with federal laws and university policy.

USC Campus Support and Intervention - (213) 740-0411
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.

Diversity, Equity and Inclusion - (213) 740-2101
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-1200 – 24/7 on call
Non-emergency assistance or information.

Office of the Ombuds - (213) 821-9556 (UPC) / (323-442-0382 (HSC)
A safe and confidential place to share your USC-related issues with a University Ombuds who will work with you to explore options or paths to manage your concern.

Occupational Therapy Faculty Practice - *(323) 442-2850 or* otfp@med.usc.edu
Confidential Lifestyle Redesign services for USC students to support health promoting habits and routines that enhance quality of life and academic performance.