

Course Syllabus



San José State University

College of Science/Computer Science Department CS 185C, Distributed Systems

Course and Contact Information

Instructor(s): Ben Reed

Office Location: MH 213

Telephone: (408) 924-517

Email: ben.reed@sjsu.edu

Office Hours:

Mondays 6-7 PM in MH 213

Tuesdays 10-11:30 AM over zoom <https://sjsu.zoom.us/j/4077267356>
(<https://sjsu.zoom.us/j/4077267356>)

Wednesdays 1-2:30 in MH 213

Thursdays 10-11:30 AM over zoom <https://sjsu.zoom.us/j/4077267356>
(<https://sjsu.zoom.us/j/4077267356>)

Class Days/Time: MW 4:30-5:45

Classroom: MH 222

Prerequisites: CS 149

Course Description

Current issues in operating systems, including multiprocessor systems and distributed computing, networks,

security and performance. Case studies of current operating systems.

Course Format

This is an in-person course.

Course Learning Outcomes (CLO)

Upon successful completion of this course, students will be able to:

1. define the terminology and common ideas of distributed computing.
2. explain fundamental ideas of distributed computing.
3. apply distributed computing principles to modern applications.
4. explain the challenges present in a distributed environment.
5. discuss the differences and trade-offs between various solutions for distributed systems problems.

Required Texts/Readings

Textbook

Designing Data-Intensive Applications by Martin Kleppmann publish by O'Reilly. Available online at the SJSU library.

Course Requirements and Assignments

Homework will be given, but will not be graded. It is intended for self-evaluation and will be the basis for future exams. I encourage students to work on homework in groups and discuss possible solutions together. We will take time at the beginning of each class to discuss any difficulties students have completing the homework. There will be discussion assignments associated with each set of homework questions to encourage participation and preparation before the in-class discussions.

Ethical assignment (10%)

Along with technical questions in the homework, we will also discuss ethical issues related to distributed systems. We want you to understand that along with technical choices come moral implications, and we want to be able to identify and reason about them. There will be 2 written (1 page) assignments to discuss contemporary ethical issues.

Programming assignments (BAKEOFF) (30%)

We will be doing individual programming assignments. At the end of many of the assignments we will be doing a bakeoff (see RFC 1025 <https://datatracker.ietf.org/doc/html/rfc1025> (<https://datatracker.ietf.org/doc/html/rfc1025>)) in class to see if our implementations can interoperate. THERE ARE NO LATE ASSIGNMENTS. THE BAKE-OFF IS DONE IN CLASS. Individual programming assignments are not group projects. If students get help on assignments, even to resolve a stupid

problem, it must be documented in the code with the name of the person rendering the help and a brief description of the help provided. Extensive help on a project will result in a reduced grade. Failure to document help, or any other forms of cheating will result in a failing grade on the assignment at a minimum and may result in failure of the course. All incidents will be reported to the Office of Student Conduct & Ethical Development. Even in open source, you cannot copy code from one open source project to another without attribution. Sharing solutions with other students, even if it is indirectly through public source repositories, falls under "aiding and abetting".

The University Policy S16-9, Course Syllabi (<http://www.sjsu.edu/senate/docs/S16-9.pdf>) requires the following language to be included in the syllabus: "Success in this course is based on the expectation that students will spend, for each unit of credit, a minimum of 45 hours over the length of the course (normally three hours per unit per week) for instruction, preparation/studying, or course related activities, including but not limited to internships, labs, and clinical practica. Other course structures will have equivalent workload expectations as described in the syllabus

Final Examination and exams (60% total)

All exams, including the final, will be taken in class using lockdown browser. This course will have a cumulative final exam given during exam week. There will be three in-class exams given in the semester (the last being the final exam :). The second exam will have two questions derived from the previous exam, and the final exam will have two questions derived from the first exam and two questions derived from the second exam.

I do not grade on a curve. The exams and projects measure what you are expected to have learned. There aren't many opportunities for extra credit, but there are bonus questions on exams.

Grading Information

Grades will be calculated based on the individual project grades, the two mid-semester exams, the final, discussion participation, the video explanation of a paper, and the ethic assignment. Thus, the grade distribution is

programming assignment	30%
Midterm 1	20%
Midterm 2	20%
Final	20%
Ethics assignment	10%

I do not curve grades. I do round the reported grade up. I don't feel that it is fair that only students who beg for their grade to be rounded up to get it, so I automatically round-up for everyone. So, if you get an 89.1, that will be rounded to an A-. However, do not expect a 68.9 to be rounded to a C-, even if you can go back and argue for some points to get it to a 69.1. The whole point of rounding is to find people that got close enough the grade they deserve. Getting a 68.9 is not close to a C-.

Grade	Percentage
A plus	96 to 100%
A	93 to 95%
A minus	90 to 92%
B plus	86 to 89 %
B	83 to 85%
B minus	80 to 82%
C plus	76 to 79%
C	73 to 75%
C minus	70 to 72%
D plus	66 to 69%
D	63 to 65%
D minus	60 to 62%

Classroom Protocol

This is your class. Please ask questions. Please come prepared. Do not engage in activity that may distract other students. I do not take attendance except for the first two classes. Students not attending either of the first two classes will be dropped to make room for students on the waiting list. Attempting to get marked as present (by having someone else attend in your place or using technological deceptions) will be considered academic dishonesty and at a minimum will result in you getting dropped from the course.

University Policies

Per [University Policy S16-9](http://www.sjsu.edu/senate/docs/S16-9.pdf) (<http://www.sjsu.edu/senate/docs/S16-9.pdf>), relevant university policy concerning all courses, such as student responsibilities, academic integrity, accommodations, dropping and adding, consent for recording of class, etc. and available student services (e.g. learning assistance, counseling, and other resources) are listed on [Syllabus Information](https://www.sjsu.edu/curriculum/courses/syllabus-info.php) (<https://www.sjsu.edu/curriculum/courses/syllabus-info.php>) [web page](https://www.sjsu.edu/curriculum/courses/syllabus-info.php) (<https://www.sjsu.edu/curriculum/courses/syllabus-info.php>) (<https://www.sjsu.edu/curriculum/courses/syllabus-info.php>). Make sure to visit this page to review and be aware of these university policies and resources.

COVID-19 and Monkeypox

Students registered for a College of Science (CoS) class with an in-person component should view the [CoS COVID-19 and Monkeypox Training](https://drive.google.com/drive/folders/1Vmp39U9-CNpbwRobtZsGIZPTgRwV_Nh6) (https://drive.google.com/drive/folders/1Vmp39U9-CNpbwRobtZsGIZPTgRwV_Nh6) slides for updated CoS, SJSU, county, state and federal information and guidelines, and more information can be found on the [SJSU Health Advisories](https://www.sjsu.edu/healthadvisories/) (<https://www.sjsu.edu/healthadvisories/>) website. By working together to follow these safety practices, we can keep our college safer. Failure to follow safety practice(s) outlined in the training, the SJSU Health Advisories website, or instructions from instructors, TAs or CoS Safety Staff may result in dismissal from CoS buildings, facilities or field sites. Updates will be implemented as changes occur (and posted to the same links).

Course Schedule

Week	Date	Topics, Readings, Assignments, Deadlines
1	8/22/2022	Aspects of Reliable, Scalable, and Maintainable (https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch01.html) Applications (https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch01.html)

Aspects of Reliable, Scalable, and Maintainable Applications

- | | | |
|---|-----------|---|
| 1 | 8/24/2022 | https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch01.html
Attendance assignment due |
| 2 | 8/29/2022 | <u>Replication</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch05.html
Gossip Assignment |
| 2 | 8/31/2022 | <u>Replication</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch05.html |
| 3 | 9/7/2022 | <u>Replication</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch05.html
ZooKeeper Leader Assignment |
| 4 | 9/12/2022 | <u>Partitioning</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch06.html
Gossip bakeoff |
| 4 | 9/14/2022 | <u>Partitioning</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch06.html |
| 5 | 9/19/2022 | <u>Partitioning</u>
https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch06.html (online class) |
| 5 | 9/21/2022 | Exam 1 |
| 6 | 9/26/2022 | ZooKeeper Leader Bakeoff |
| 6 | 9/28/2022 | <u>Transactions</u> |

<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch07.html>

Distributed K/V assignment

[Transaction](#)

7 10/3/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch07.html>

[Transactions](#)

7 10/5/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch07.html>

[Transactions](#)

8 10/10/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch07.html>

[Challenges of Distributed Systems](#)

8 10/12/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch08.html>

[Challenges of Distributed Systems](#)

9 10/17/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch08.html>

Distributed K/V bakeoff

[Challenges of Distributed Systems](#)

9 10/19/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch08.html>

[Challenges of Distributed Systems](#)

10 10/24/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch08.html>

[Challenges of Distributed Systems](#)

10 10/26/2022 <https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch08.html>

Transaction log assignment

11 10/31/2022 Exam 2

- 11 11/2/2022 **Consistency and Consensus**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch09.html>
- 12 11/7/2022 **Consistency and Consensus**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch09.html>
- 12 11/9/2022 **Consistency and Consensus**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch09.html>
- 13 11/14/2022 **Consistency and Consensus**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch09.html>
- 13 11/16/2022 **Distributed Batch Processing**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch10.html>
- 14 11/21/2022 **Distributed Batch Processing**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch10.html>
- 15 11/28/2022 **Distributed Batch Processing**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch10.html>
- 15 11/30/2022 **Distributed Batch Processing**
<https://learning.oreilly.com/library/view/designing-data-intensive-applications/9781491903063/ch10.html>

16

12/5/2022 Transaction log bakeoff

Final Exam

Wednesday, December 14, 2022 @ 2:45