



Course Syllabus: Advanced Distributed & Networked Systems - CS 345

| | |
|--|--|
| Division | Computer, Electrical and Mathematical Sciences & Engineering |
| Course Number | CS 345 |
| Course Title | Advanced Distributed & Networked Systems |
| Academic Semester | Spring |
| Academic Year | 2018/2019 |
| Semester Start Date | 01/27/2019 |
| Semester End Date | 05/23/2019 |
| Class Schedule (Days & Time) | 04:00 PM - 05:30 PM Sun Wed |

Instructor(s)

| Name | Email | Phone | Office Location | Office Hours |
|--------------|--------------------|---------------|----------------------------------|--------------|
| Marco Canini | marco@kaust.edu.sa | +966128080489 | 0144, 1, Al-Khawarizmi (bldg. 1) | |

Teaching Assistant(s)

| Name | Email |
|------|-------|
| | |

Course Information

| | |
|--|---|
| Comprehensive Course Description | <p>This class is an advanced course in distributed and networked systems.</p> <p>This course takes the form of a seminar that is based on a selection of papers that either have had a strong impact on distributed and networked systems today, or explore novel ideas that may be important in the future. The syllabus for this course will vary from year to year so as to cover a mixture of older and more contemporary papers. Contemporary papers will be generally selected from the past 5 years, primarily drawn from high quality conferences such as SOSP, SIGCOMM, OSDI, NSDI and EuroSys. Students are expected to read papers before the class, write a summary and "review" of the paper prior to class, and participate in the discussion during the class.</p> <p>The course has one semester-long project that is aimed at producing a conference/workshop-quality research paper and must involve writing some software. The project must address an open research problem. Suggested project ideas will be provided by the instructor. To get ideas for projects, please DO be courageous and read ahead in the course schedule! Your project topic is expected to be distinct from your "regular" research project. It can be related, but try to strike out a new direction that is truly systems research! It is highly recommended that you work in groups. Teams are suggested to be 2 students. All projects will have to submit three versions of reports (proposal, midterm, and final reports); these submissions should read like research papers (or parts thereof). Up to five "best projects" at the end of the semester will be earmarked for expedited submission to a renowned conference, with the help of the instructor's involvement even after the semester is over.</p> |
| Course Description from Program Guide | |

| | |
|----------------------------------|--|
| Goals and Objectives | <p>This class is a graduate seminar that covers design and implementation concepts in distributed systems and networked systems by reviewing a selection of classical and contemporary papers. We will study efficient system design and evaluation, as well as learn high-level system issues, with a focus on exciting topics in distributed and networked systems. Research in these areas also tends to be scattered across disjoint sets of researchers and conferences, and the course attempts to study commonalities. On completion of this course students should have a broad understanding of some key papers and concepts in computer systems research, as well as an appreciation of how to argue for or against any particular idea.</p> <p>By the end of the course, students should be able to:</p> <ul style="list-style-type: none"> -Comprehend and critique relevant research papers in the area of distributed and networked systems. -Present their research both orally in a concise way and within the allotted time as well as in writing. -Defend the research approach, design decisions, and the evaluation methods in a discussion. -Moderate a discussion after a research presentation. |
| Required Knowledge | <p>Students are expected to have good programming skills and must have taken at least one systems-related course (e.g., from operating systems/CS140, distributed systems/CS240, networking/CS244, or databases/CS245), or instructor consent, is required.</p> |
| Reference Texts | <p>There is no textbook for the course. You will be expected to read 1-2 papers per class. The papers are available electronically through the course website.</p> |
| Method of evaluation | <p>20.00% - Scientific review article presentation 50.00% - Research Project 15.00% - Presentation 15.00% - Attendance and Participation</p> |
| Nature of the assignments | <p>The class consists of two main activities:</p> <ol style="list-style-type: none"> 1. Papers and discussions: We will read 1-2 papers for each class (i.e. 2-4 papers per week), and discuss them in class. You will need to spend about 2 hours reading each paper, and making notes, to prepare you to discuss them in class. This only works if you come to class ready to discuss the papers in detail, which is why 15% of your grade is for in-class participation. There will be as much as 10-15 hours of reading per week. Do not take this course unless you are willing and able to do a lot of reading. 2. Project: You will do a semester-long open-ended project (see sidebar for more information). |

Course Policies

Groups

It is highly recommended that you work in groups. Teams are suggested to be 2 students.

Declare your group's membership as well as preferences for papers by February 6, 2019. After this date, we will form groups from the remaining students.

Note that paper preferences are only hints.

Paper presentations

The course will be conducted as a seminar. Only one student will present in each class. Each student will be assigned to present two papers (mandatory and companion) on the same day at least once throughout the semester. Presentations should last at most 45 minutes without interruption. However, presenters should expect questions and interruptions throughout. In the presentation, you should:

- Motivate the paper and provide background.
- Present the high level idea, approach, and/or insight (using examples, whenever appropriate).
- Discuss technical details so that one can understand the key details without carefully reading it.
- Explain the difference between this paper and related work.
- Raise questions throughout the presentation to generate discussion.

The slides for a presentation must be emailed to the instructor at least 24 hours prior to the corresponding class. You should use this powerpoint template for making your slides.

Paper reviews

Each critique should be about a half a page and should roughly cover the following questions:

Read all the papers of each day carefully and write your own paper review. Being able to critically judge others' work is crucial for your understanding. For papers that require summaries, you must address a number of questions including:

- What is the problem addressed by the paper? Is the problem real? Why is this problem important?
- What is the hypothesis of the work?
- What is the proposed solution's main idea, and what key insight guides their solution?
- Why is the solution different from previous work?
 - › Are system assumptions different?
 - › Is workload different?
 - › Is problem new?
- Does the paper (or do **you**) identify any fundamental/hard trade-offs?
- What is one (or more) drawback or limitation of the proposal, and how will you improve it?
- Do you think the work will be influential in 10 years?
 - › Why or why not?

*The paper reviews of each day must be submitted electronically at least 24 hours prior to the corresponding class. **Late reviews will not be counted.*** You can miss up to 4 paper reviews during the term without penalty. Each missing review beyond the fourth one will result in 25% decrease in grade for paper reviews. Meaning, failing to turn in 8 or more reviews on time will result in a zero for all paper reviews.

The grade for this part will be assigned as follows: at the end of the term, 2 of your reviews will be randomly selected and graded. The higher grade of the two will be used for grading.

Participation

Is class participation based solely on attendance? No. Attendance is a necessary but not sufficient condition for good class participation. You are expected to attend all lectures (you may skip up to 2 lectures due to legitimate reasons), and more importantly, participate in class discussions. If you have any concerns about not being able to regularly attend class (e.g., you will have to miss several classes during the quarter) please discuss this as soon as possible with the instructor. Beyond attendance, we evaluate class participation by observing how prepared students are to discuss the covered paper when they come to class. This is not a trivial requirement because we expect papers to have been read thoroughly prior to lecture.

Project

Research projects are the critical aspect of the course. Your goal is to do publishable quality systems research on an instructor-approved problem; that is, to add to our understanding of how to build systems. Research projects must be written up in a term paper.

Exam

This class has no exam.

Late work

There is no policy on late work. If you cannot submit your work by the deadline, it will not be accepted.

Exceptional circumstances

If you have a medical emergency, then email the instructor to request an extension. At times, there may be system problems with the computer infrastructure which may prevent you from submitting your reviews on time. In these cases an extension will be granted to the entire class.

Policies

We will have zero tolerance for academic misconduct. Cheating, plagiarism, and any form of dishonesty will be handled with maximum severity, according to university regulations. If you are ever in doubt about whether an action on your part may constitute unacceptable behavior, please ask the instructor before proceeding—doing so afterward is too late.

Apart from the term project, any work you turn in must be your own and is to be done individually, and the usual code of conduct applies. You must acknowledge any sources of your words, ideas, and software when

| | |
|-------------------------------|---|
| | they are not your own, and you must disclose in advance, without any specific request, any sources you used. Do not use code from a student who took the course in a previous semester. |
| Additional Information | For a more up to date version of this syllabus and the actual lecture schedule, you are required to see the version of the syllabus on the class website. The class website will be announced on the first lecture. An online discussion forum will be offered on the Piazza system. This instructor does not use Blackboard. |

Tentative Course Schedule

(Time, topic/emphasis & resources)

| Week | Lectures | Topic |
|------|----------------------------------|------------------------------------|
| 1 | Sun 01/27/2019 Wed 01/30/2019 | Introduction |
| 2 | Sun 02/03/2019 Wed 02/06/2019 | Background |
| 3 | Sun 02/10/2019 Wed 02/13/2019 | Storage and computation |
| 4 | Sun 02/17/2019 Wed 02/20/2019 | Storage and computation |
| 5 | Sun 02/24/2019 Wed 02/27/2019 | Resource management and scheduling |
| 6 | Sun 03/03/2019 Wed 03/06/2019 | Resource management and scheduling |
| 7 | Sun 03/10/2019 Wed 03/13/2019 | Systems for AI |
| 8 | Sun 03/17/2019 Wed 03/20/2019 | Systems for AI |
| 9 | Sun 03/24/2019 Wed 03/27/2019 | Spring Break |
| 10 | Sun 03/31/2019 Wed 04/03/2019 | Systems for AI |
| 11 | Sun 04/07/2019 Wed 04/10/2019 | Systems for AI |
| 12 | Sun 04/14/2019 Wed 04/17/2019 | Systems for AI |
| 13 | Sun 04/21/2019 Wed 04/24/2019 | Big data |
| 14 | Sun 04/28/2019 Wed 05/01/2019 | Big data |
| 15 | Sun 05/05/2019 Wed 05/08/2019 | New trends |
| 16 | Sun 05/12/2019 Wed 05/15/2019 | New trends |
| 17 | Sun 05/19/2019 Wed 05/22/2019 | Final Exam Week |

Note

The instructor reserves the right to make changes to this syllabus as necessary.