

CS 525: Advanced Distributed Systems

Course Overview

Over the past few decades, distributed systems have come to surround us. Everything is distributed nowadays! Clouds and datacenters, distributed machine learning, IoT/sensor networks, edge systems, the Grid, Internet and the Web are all examples. Some functionalities are being consolidated in datacenters and computing clouds (e.g., AWS, Azure, Google Compute Engine, and many private clouds, etc.), others run on edge devices, and some even run on mobile devices. This course focuses on multiple case study areas, spanning both classical topics to bleeding edge topics of distributed compute systems, including (but not limited to): **cloud computing, distributed machine learning, edge computing, IoT/sensor networks**. We will study efficient protocol design and evaluation, as well as learn high-level system issues, with a focus on exciting topics in distributed systems. Readings are research papers that are a balanced mix of classical papers and very newly published work by authors all over the world, and from different research communities. Research in these areas also tends to be scattered across disjoint sets of researchers and conferences, and the course attempts to study commonalities. Projects can be one of two types - **research projects** or **entrepreneurial projects**.

Prerequisites

Basic Computer Science and basic computer programming skills are essential. Knowledge of Operating Systems (CS 241/340/341 or CS 423), or Networks (CS 438), or Distributed Systems (CS 425), or an equivalent course, or instructor consent, is required.

Course Website <http://courses.engr.illinois.edu/cs525/>

Timings

Class: Wednesdays and Fridays, 11.00 AM - 12:15 PM, 1310 DCL.

Professor Office Hours (class days only): Wednesdays and Fridays 12:15 PM - 1.00 PM (chat after class or in Indy's office 3112 SC)

Course Staff

Dr. Indranil Gupta (Professor)
3112 SC, indy@illinois.edu, 265-5517

TAs: Anna Karanika (annak8); Office hours by appointment only

Course Content

The first few weeks of lectures introduce ground basics in cloud computing, peer to peer systems, distributed algorithms and IoT/sensor networks. Subsequently, over 70 research papers in various areas of distributed systems are presented, discussed, and debated by the students. The selection includes papers from industry and successful companies, as well as contemporary papers from the topmost distributed systems conferences such as SOSP, OSDI, NSDI, VLDB, MLSys, Mobisys, MobiHoc, NeurIPS, ICML, ASPLOS, and many others. Rather than running through all

the papers in a few of such proceedings and journals, we will pick and 1) topics that are relevant and active today, both in industry and academia, and 2) choose publications appropriate to the stated goals of this course.

An essential component of this course is a project that could be either:

1. a **Research Project**, or
2. an **Entrepreneurial Project**

The choice is yours, but you must make this choice at the semester start. We encourage you to do a research project, and treat the entrepreneurial project as a fallback option.

A research project will solve a cutting-edge research problem with innovative ideas, build systems solutions (real implementation required!), and evaluate it experimentally in a deployment scenario. The goal of this project is a conference-quality paper submission and acceptance; many past course projects have been published in conferences and journals, winning best paper awards.

An entrepreneurial project will explore a new idea for a startup or non-profit company. You will have to build the system or working prototype (real implementation required!) and evaluate it experimentally. You will also have to write a Business Plan. Think about who your potential customers are. In order to display the technical viability of your plan you will have to write a technical paper (with similar formatting and goals as the research project).

In all projects (research or entrepreneurial), the core ideas have to come from the students. The instructor will work with you in refining your ideas. for projects.

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 distributed systems!

It is highly recommended that you work in groups. Teams must only be 2 to 3 students. All projects will have to submit three versions of reports (survey, midterm, and final reports), and a final project presentation (which is usually before the final report is due) - for research projects, these submissions should read like research papers (or parts thereof), while for entrepreneurial projects, these submissions should be readable documentations of ideas, systems, implementations and (a required) business plan. For the final project presentation, you will have a short time (typically 5 minutes), so prepare a punchy “lightning” talk!

At the end of the course, a few “best” projects will be selected. All promising projects can be continued after the semester.

This course’s projects in past semesters have been accepted in conferences (e.g., ICDCS, Infocom, Middleware, MMCN, SASO, GRID, etc.), as well as appear in several top journals (e.g., ACM TAAS, ACM TOSN, IEEE TNSM, JSS, Distributed Computing, etc.). Course project papers have won a Best Paper Award at conferences (IC2E 2016, ICAC 2015) and workshops (BigMine 2012).

About the Class

The initial few weeks of class will consist of lecturing, with the intent of building up common knowledge and grounding for the latter half of the course. We will then transition to student-led presentations of papers. Once student-led presentations start, students who are not presenting in that particular session are expected to write short reviews (1-2 pages total) for the “Main Papers” in that session. Active class participation is required, even in the initial part of the course!

Class Evaluation: Project (survey + midterm + final project presentation + final report), reviews of papers, presentation, and class participation (may include a quiz). Tentative splits are 50% (5% + 15% + 7.5% + 22.5%), 20%, 25%, 5% respectively. The class grading is curved separately for undergrad students and for grad students. At this 500-level of courses, you shouldn't worry about grades, and instead enjoy the material and be your creative and hard-working best effort!