

ENGR E-516: Engineering Cloud Computing

Course Overview and Logistics

Prateek Sharma

Course Information

- ENGR-E 516 / CMPSCI-B 649 : Engineering Cloud Computing
- All information on course website

<http://homes.sice.indiana.edu/prateeks/cloud-course.html>

The screenshot shows a web browser window displaying the course website. The browser's address bar shows the URL <https://img.sxkcd.com/>. The website content includes:

- Engineering Cloud Computing**
ENGR-E 516/CMPSCI-B 649 (Fall 2020)
- Announcements** (a section header with a blue underline)
- Course Description** (a section header with a blue underline)
- Course Description Text:**

This course will teach the fundamental concepts, engineering principles, and practical skills pertaining to the effective use of cloud computing. This course will focus on both cloud applications and the design of cloud platforms. We will cover the relevant concepts from operating systems, computer networks, and distributed systems.

This course should be useful to anyone who wants a deeper understanding of how the cloud works, as well those who want to learn how to easily and effectively use the cloud for running their applications at low cost. We will look at a wide spectrum of cloud-based applications such as a parallel data processing (e.g., MapReduce), data storage and caching (e.g., key-value stores), scientific computing, interactive notebooks (e.g., Jupyter), etc.

We will also look at the challenges involved in the efficient operation of large-scale cloud platforms with hundreds of thousands of servers. The course will cover a wide gamut of data center optimization techniques such as hardware virtualization, distributed resource management, and software-defined datacenters.

This course will expose students to popular cloud platforms such as Amazon EC2, Google Cloud Platform, Microsoft Azure, etc., and introduce students to new developments such as serverless computing and edge-clouds.

More details about the course syllabus can be found on the course website.

Prerequisites: The course has no official prerequisites, but requires a high comfort-level with systems programming and debugging. The assignments in this course will include nontrivial programming in the language of your choice.

Logistics: This course is also cross-listed as CSCI-B 490. This year, the class is **online**: lecture videos will be posted on Youtube (see Canvas for links).

At the bottom of the page, there are navigation icons (back, forward, search, etc.) and a page number **2 / 16**.

Online Instruction

- Lecture videos will be posted on Youtube
- Short quizzes on Canvas each week

Welcome!

Me:

- Assistant Professor in Intelligent Systems Engineering
- Research areas: Cloud computing, Operating systems, Distributed systems
- Ph.D in Transient Cloud Computing (UMass-Amherst 2018)

I Like:

- Tinkering around with Free Software: Linux, Emacs, Firefox

How I spent my summer..

Rust

What is this course about?

This course is about how the largest computing systems are built and run, and how you can use them effectively.

Who is this course for?

We will cover **advanced topics** in computer systems:

- 1 Distributed data processing and computation (MapReduce etc.)
- 2 Hardware virtualization and hypervisor design
- 3 Performance modeling (queueing theory, etc.)

To learn any of this, you will need a solid grasp of:

- Operating Systems
- Computer Networks
- Systems programming

- Early parts of the course will cover OS and networking essentials
- Must put in the hard work and do the early “warm up” assignments

Course Overview

- 1 Fundamentals of operating systems, computer networks, and distributed systems
- 2 Cloud applications [network services, data processing]
- 3 Virtualization
- 4 Distributed management of computing resources
- 5 Current and future cloud ecosystems and services
- 6 Distributed data storage: costs, caching, and consistency

Grading Criteria

Programming assignments (~5)	50%
Homework, readings, and paper reviews	15%
Final Exam	25%
Quizzes and other class participation	10%

- **The relative weights of homework, assignments, and exams will change**

Assignments and Projects

- 5 programming assignments
- Software design: most assignments will be loosely specified.
 - Choose your own language and design (**Python/Java** supported by TAs)
 - You are encouraged and required to come up with realistic solutions to loosely defined problems
 - State your design assumptions and *implement well*
- Projects in the latter half will have to be deployed on a public cloud (such as Amazon AWS, Google, or Azure).
- No group-work. You should submit your own work
- No late submissions.

Start Early!

Prerequisites

- Basic knowledge of operating systems and how programs run
- Basics of Linux command line
- Comfortable with programming and debugging
- **Userspace** programming. *Not* an operating systems course
- Be comfortable in at least one of C, Python, Java, etc.

Learning Resources

Reference books (both books available online):

- 1 Distributed Systems: Principles and Paradigms, 3rd Edition. Maarten Van Steen and Andrew Tanenbaum.
 - 2 Operating Systems: Three Easy Pieces. Remzi H. Arpaci-Dusseau and Andrea C. Arpaci-Dusseau
- All slides and lecture notes will be on course website
 - Canvas used only for homework submission, discussion, etc.

How to succeed in this course

- 1 Be interested in learning about computer software systems
- 2 Attend lectures, ask questions
 - Text-book does not have all material we will cover
 - Cloud computing is a relatively new field. Text-book has many gaps.
- 3 Read reference material posted for each lecture

Getting Help

- Office hours (schedule will be posted soon)
- Post questions on Canvas



What this course is *really* about?

Officially:

How the largest computing systems are built and run, and how you can use them effectively.


Reality:

- Implementing fairly complex software systems...
- With a design based on underspecified constraints....
- Tackling the hardest challenges in performance and communication...
- *Mostly on your own.*

But Why?

- Real learning of computer systems only happens if you put in time and effort in understanding how these systems behave and operate
- This course provides a safe, low-risk environment to try and learn difficult things

There's an XKCD For That



Next Time

- Overview of cloud computing, its history, and major challenges