ENGR E-516: Engineering Cloud Computing

Course Overview and Logistics

Prateek Sharma

Fall 2023

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

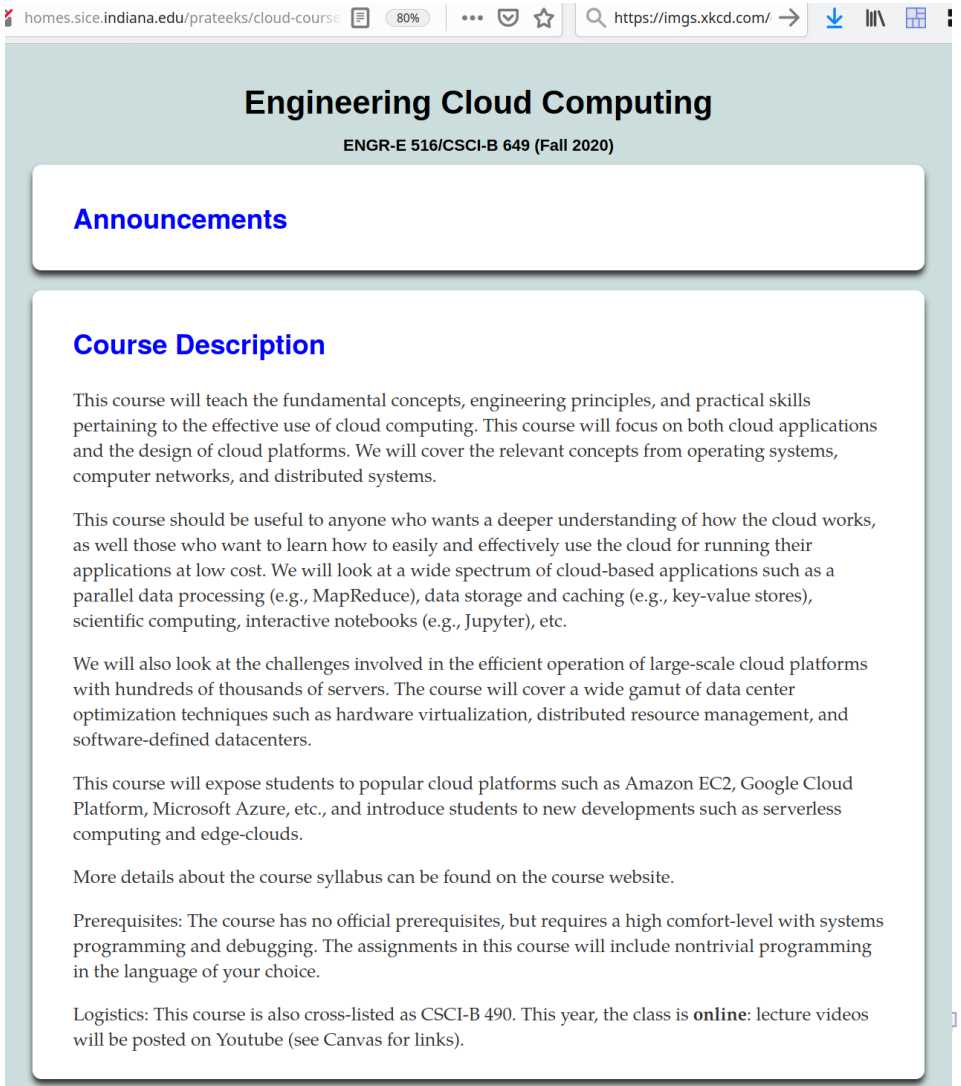

What this course is about

- How cloud platforms and services work
- How they can be used effectively

User centric view:

- How to build and deploy applications on the cloud
- How to improve performance, cost

Who is this course for?

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

1. Hardware virtualization and hypervisors
2. Operating systems and containerization
3. Performance modeling (queueing theory, etc.)
4. Distributed data processing and computation (MapReduce 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 some 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 (~4)	40%
Homework, readings, and paper reviews	10%
Mid term and Final Exams	40%
Lecture notes and class participation	10%

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

Assignments and Projects

- 4 programming assignments
- Software design: most assignments will be loosely specified.
  - Choose your own language and design
  - Must be comfortable with some systems programming language (C, Go, Rust)
  - 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 (Google).
- No group-work. You should submit your own work.
- Late submissions: total 4 late-days. use as you wish.

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 atleast one of C, Python, Java, etc.

Dont take this course for the grades

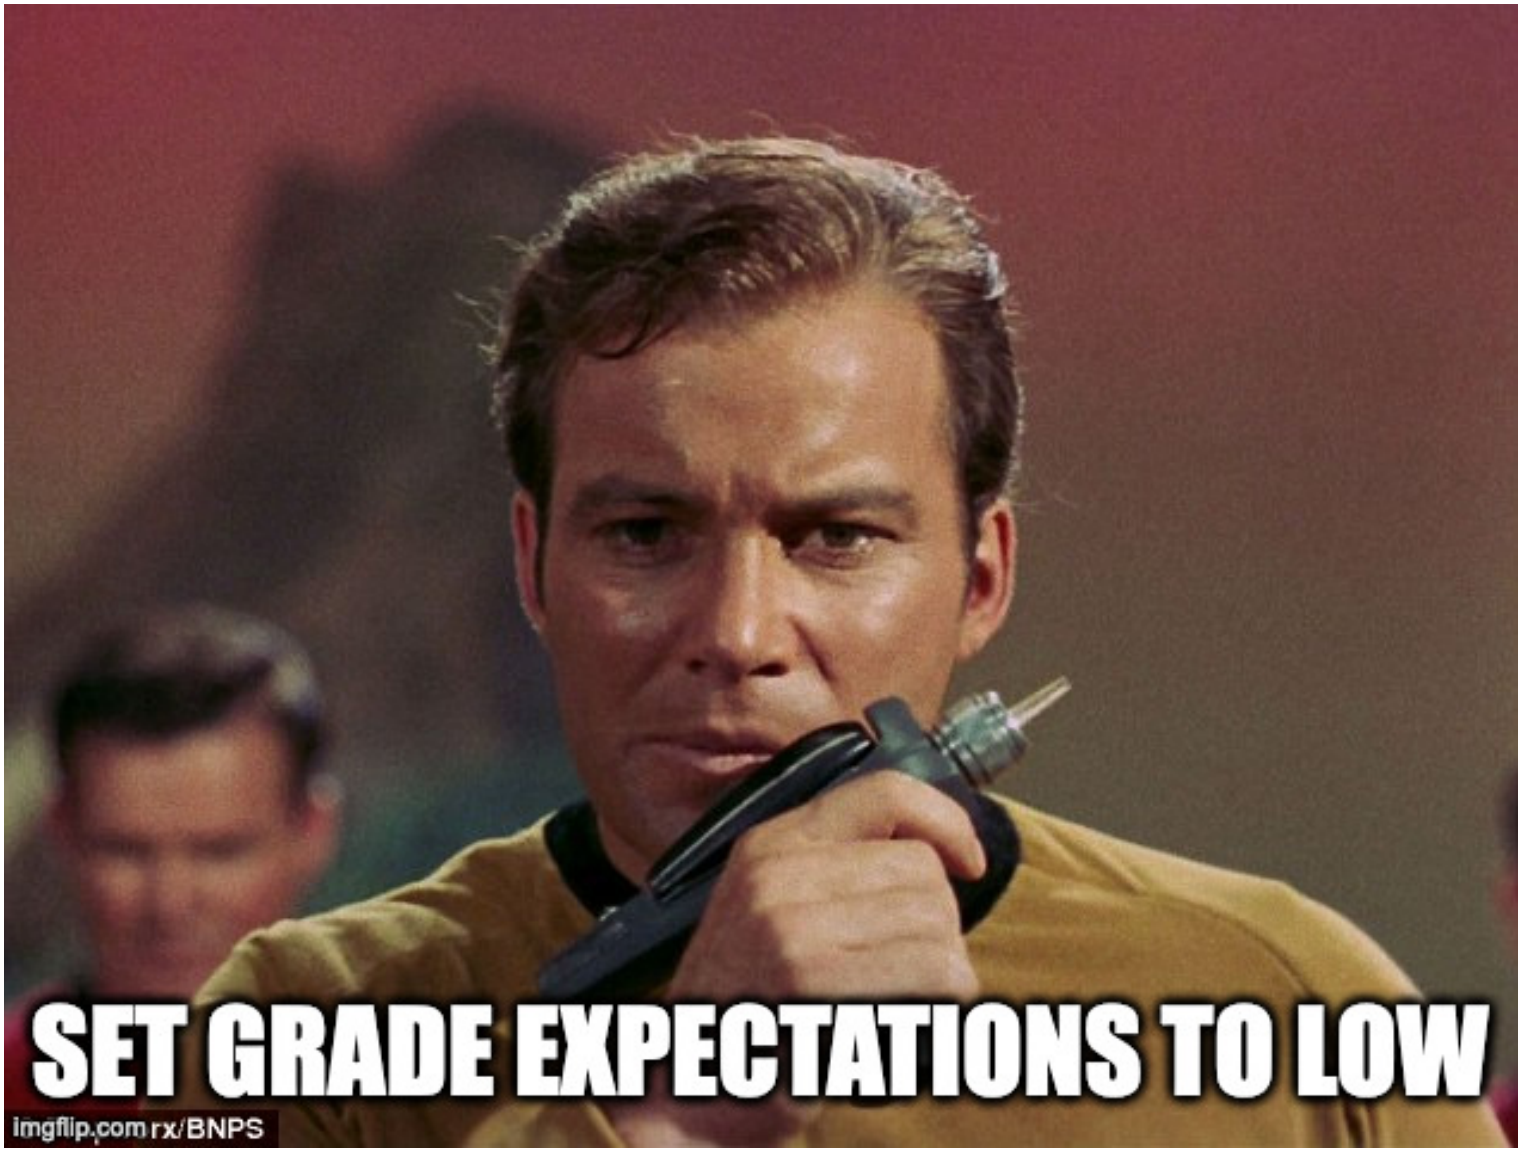

Learning Resources

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

- **Ask questions in class**
- 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

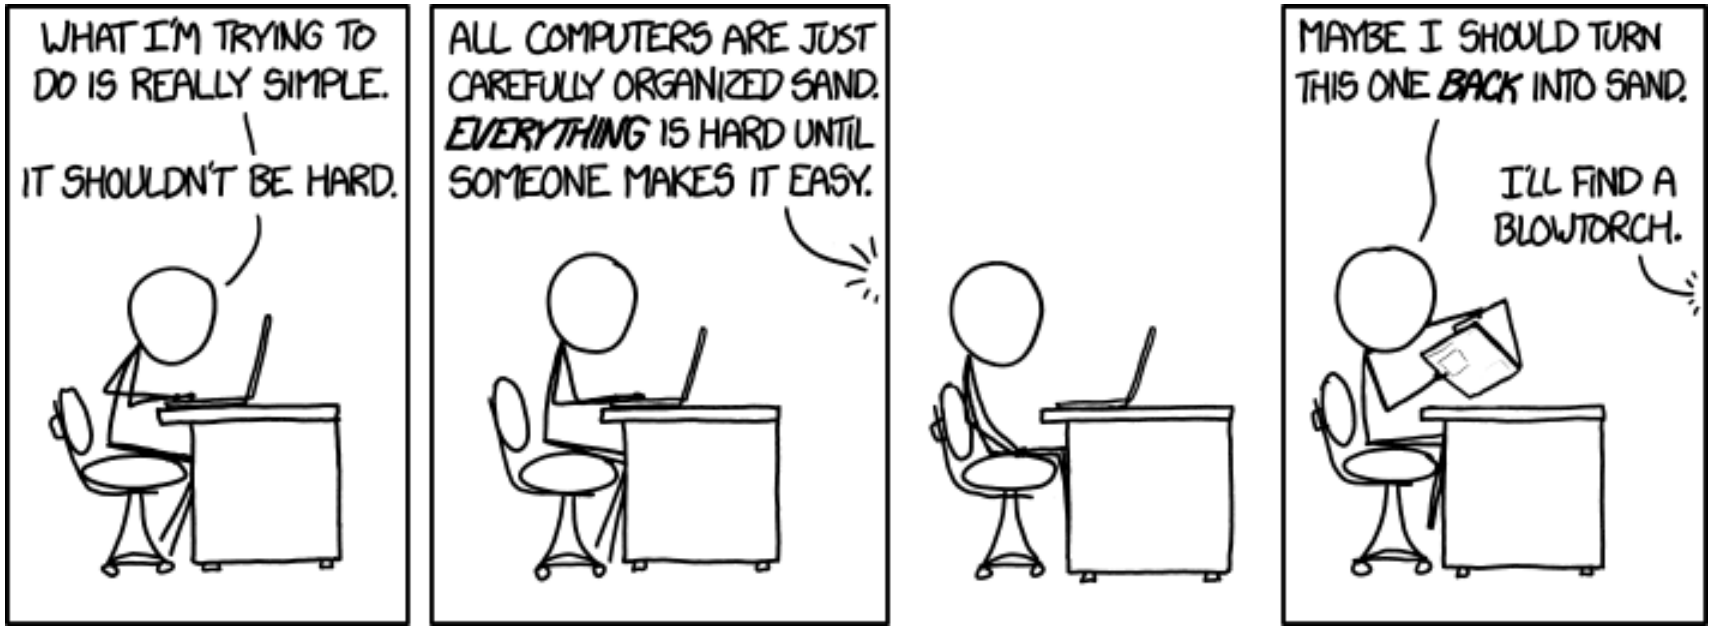

Up Next

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