




Introduction

CS 6410: Advanced Systems

Fall 2024

Hakim Weatherspoon



Cornell Bowers CIS
Computer Science

Inclusion

- We strive to make CS6410 a welcoming, safe, equitable, and respectful environment, consistent with [Cornell's commitments](#)
- We recognize that the society we live in is none of those things, that we have implicit biases, and that we have to work hard every day to counter those biases to create an inclusive environment
- If you witness a bias incident or have been the victim of one, please file a [confidential report](#) with Cornell
- If you have any suggestions such as improvements to the web site, syllabi, slides, homework and exam questions, and
- so on, you can email cs6410-prof@cornell.edu

Emotional Help

Engineering Advising	https://www.engineering.cornell.edu/resources/advising	Academic advising for engineering students
Arts College Student	www.arts.cornell.edu/stu-adv/	Listing of general support services for a variety of concerns
Gannett	www.gannett.cornell.edu	Cornell University Health Service
CAPS	https://www.gannett.cornell.edu/services/counseling/caps	If you experience emotional distress, please contact Counseling and Psychological Services
Student Disability Services	sds.cornell.edu	Ensures that all aspects of student life are accessible, equitable, and inclusive of those with disabilities. Send accommodation letters to Veronica VanCleave-Seeley (vv48, Gates 401) by Sep 15.

Get help. Get documentation. The earlier the better.
Also, please look out for each other

Emotional Help

Engineering Advising	https://www.engineering.cornell.edu/students/undergraduate-students/advising	Academic advising for engineering students
Arts College Advising	https://as.cornell.edu/advising	Listing of general support services for a variety of concerns
Cornell Health	https://health.cornell.edu/	Cornell University Health Service
CAPS	https://health.cornell.edu/services/mental-health-care	If you experience emotional distress, please contact Counseling and Psychological Services
Student Disability Services	https://sds.cornell.edu	Ensures that all aspects of student life are accessible, equitable, and inclusive of those with disabilities. Send accommodation letters to Veronica VanCleave-Seeley (vv48, Gates 401) by Sep 15.

Get help. Get documentation. The earlier the better.
Also, please look out for each other

How this class is organized

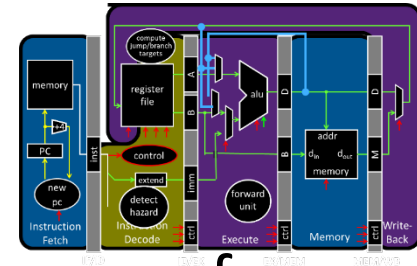
- **Who's Who**
- Before you take this class...
- Course Content
- Communication
- Grading
- Academic Integrity

Who's Who

- Where are you from?
- Where did you go to undergrad?
- Area of research? Who are you working with?
- Fun fact?

About Prof. Hakim Weatherspoon

- Professor, CS, Cornell
 - PhD from Berkeley
 - Thesis: Oceanstore, Distributed Storage System



- Focus in Computer Systems

- Fault-tolerance, reliability, security, and performance of internet-scale systems such as cloud and distributed systems
- Digital Ag, Datacenter Networks, and Cloud computing



- CEO -> Chief Scientist of a cloud-based startup company,
- Husband, father of four



How this class is organized

- Who's Who
- **Before you take this class...**
- Course Content
- Communication
- Grading
- Academic Integrity

Prerequisites

- CS 3410 or CS 3420 (ECE 3140) or equivalent
- CS 4410 or equivalent
- ***Otherwise: you must contact Prof. Hakim, explain your situation and request permission***

How this class is organized

- Who's Who
- Before you take this class...
- **Course Content**
- Communication
- Grading
- Academic Integrity

Objectives

- Learn about current research perspectives on computer systems, including advanced systems abstractions, principles, and artifacts that have had lasting value
- Understand attributes of systems research that is likely to have impact
- Become comfortable navigating the literature in this field
- Gain experience in thinking critically and analytically about systems research
- Acquire the background needed to work on research problems currently under study at Cornell and elsewhere

CS PhD program requirements

- As a Cornell CS PhD student, you are required to demonstrate (1) basic competency in systems at the CS 4410 level, (2) take a systems breadth course, and (3) take at least one 6xxx course in the "systems style"
- **CS6410 can satisfy all of these requirements** (assuming you get a B+ or better)
 - you won't need to take CS4410

Draft Topics

- Specification of Concurrent Systems
- Concurrent Systems
- Networked Systems
- Distributed Systems
- Storage Systems
- Cloud Systems
- IoT Systems
- ...

Lectures

- Mixture of lectures from me and lectures from PhD/MS students in the class
- I will present initial lectures
- Attendance taken for all lectures
- **Participate!!**

Reading

- Mixture of seminal papers and contemporary papers on systems research
- Approx. one mandatory and one or two recommended papers per class meeting
 - i.e., 2 per week, or about 25-50 pages per week
 - First read “How to read a paper” by S. Keshav
 - <https://dl.acm.org/doi/abs/10.1145/1273445.1273458> →
 - <https://dl-acm-org.proxy.library.cornell.edu/doi/abs/10.1145/1273445.1273458>
- Not all class meetings will have a new paper though

Paper reports

- Submit for each required paper (plus possible bonus points for related papers)
 1. a summary in a few lines (3-5 lines)
 2. one to three major strengths of the paper: what is really cool about this?
 3. most major weakness of the paper, including whether it is fixable or not
 4. discuss future work opportunities
- Report (half to one page) must be submitted before the class in which the paper is presented
 - free passes for 3 papers!

Paper reports, cont'd, and GenAI

- ***AFTER*** you have read the paper and written a review
- Use GenAI, <https://copilot.microsoft.com>
- Submit
 - your GenAI prompts and output,
 - And, write how GenAI compares to your summary

GenAI and Warnings

<https://copilot.microsoft.com>

- We have a GenAI Enterprise subscription
- NEVER use a non-licensed version of GenAI
 - You give your contributions to the world, unattributed
 - You received other contributions, unattributed
- ALWAYS *cite* **when** and **how** you used GenAI
- Best practice, read and write yourself, then compare to GenAI

Paper presentation and discussion

- Paper and related materials presented by PhD/MS student
 - who will receive the paper reports ahead of time
- Student also leads a discussion
 - participation in the discussion counts towards your grade!
- May also want to present the same paper afterward in systems lunch seminar!

Presentation Schedule

- By September 3rd, PhD and MS students must select a date and paper they want to present in class
- 2 ½ weeks before the presentation: outline of slides and list of related readings
- 1 ½ weeks before presentation: slide set
- 1 week before presentation: dry run

Paper choice

Either

- seminal system paper (I have a list)
- paper in a recent major systems conference

Programming Assignments

- A small network programming assignment in programming language of choice
- Must be done individually

Project

- Do a research project
- Encouraged to be done with a partner, preferably selected from your own "class"
 - PhD/MS
 - MEng
 - Senior undergrad
- Phases:
 - develop and present an idea
 - do literature survey
 - work out idea to some degree
 - evaluate the idea to some degree
 - write a 6-10 page paper
- Could be something you're working on already, but should be in state of infancy

Project Group Code of Conduct

- Each student should feel safe, welcome, respected
- *Participate, but don't dominate*
- Be patient
- Respect diverse talents and ways of learning
- Fight your implicit biases

A well-run study group benefits **all** participants

Peer Reviews

- Papers written in class will be peer reviewed by the students
 - 3 peer reviews per student
 - If there are N projects and M students, then each project will get $3M/N$ reviews. I hope M/N will be about 2, so approx. 5-6 reviews per project
 - Uses state-of-the-art conference submission system (HotCRP)
- Paper must be revised based on peer reviews

Project Schedule

- Sept. 12th: project proposal due (1 page)
 - what is interesting about the problem?
 - what is interesting about the solution?
 - why are you the right person(s) to do it?
- Sep 26th: survey due (2-3 pages)
 - what is the related work in the area?
 - what is different about what you're proposing?
- Oct 31st: draft report due (4-6 pages)
 - Design, implementation plan, evaluation plan

Project Schedule, cont'd

- Nov. 14th (peer reviews due)
 - reviews are similar to paper reports
 - can contain detailed comments for improvements
- Nov. 21st: revised draft paper due
 - distributed to all students
- Dec 5th: Final presentation, demo, and discussion in the final class meetings
- Final paper due on December 5th

Success stories

- Lorenzo Alvisi and Jon Kleinberg took CS6410 (well, CS614 to be precise) and published their projects

Collaboration

- You may discuss the questions for each discussion paper with other students, but you may not look at other student's answers. You must write your answer yourself.
- You may use any idea from any other person or group in the class or out, provided you *clearly state what you have borrowed and from whom*
- If you do not provide a citation---that is, you turn other people's work in as your own---that is cheating
 - We will be grading you on the ideas you have added, but you should always borrow as much as you can as a starting point as there is no point in reinventing the wheel
- See “GenAI and Warning” slide

How this class is organized

- Who's Who
- Before you take this class...
- Course Content
- **Communication**
- Grading
- Academic Integrity

Course Web Page

<https://www.cs.cornell.edu/courses/cs6410>

- Schedule
- Homework release and due dates
- Slides posted before each lecture

Lets have a look around the [web site](https://www.cs.cornell.edu/courses/cs6410)

CMSX

<https://cmsx.cs.cornell.edu>

- Assignments
 - reading review submissions
 - Project milestone submissions
- Presentation Schedule
 - Sign up to present
- Grades

Ed Discussion

<https://edstem.org>

- Discussion of project
- Questions in general

Lectures

- Live in-person participation is required
 - In-person in Ithaca or NYC
 - Remote participation possible if requested ahead of time

Office Hours

- Prof. Hakim
 - By appointment
 - Office: 427 Gates Hall
- Goktug
 - By appointment

Email

cs6410-staff@cornell.edu: **time sensitive** matters

- Goes to professor & TAs

cs6410-prof@cornell.edu: **sensitive** matters

- Goes to Prof. Hakim only

Please no emails to personal email accounts

How this class is organized

- Who's Who
- Before you take this class...
- Course Content
- Communication
- **Grading**
- Academic Integrity

Academic Integrity & Honor Code

All submitted work must be your own

- OK to discuss concepts with any other student
- Study group can submit the same code
- Different study groups are not allowed to share code
- Cite all everything that is not yours
- Cite when and how GenAI used, if at all

Violations will be prosecuted

Academic Integrity

Why not cheat?

- It hurts you in various ways:
 - It reduces the value of your Cornell degree
 - It stresses you out because you might get caught
 - You won't feel good about yourself afterward
 - The energy that goes into cheating is better used for learning
- It hurts other students:
 - It stresses them out
 - By far most students are honest

If you need help, get it early

Semester Grades

- CS6410 is not a competition
- Your grade reflects your learning objectives, not how well you did compared to others
- Goal is to give everyone an A
 - I will come after you if you're not on track...
- Grade based on (7 P's):
 1. Paper reports (approx. 10%)
 2. Programming problem (approx. 10%)
 3. Project (approx. 20%)
 4. Project paper (approx. 25%)
 5. Peer reviews (approx. 5%)
 6. Participation in discussions (approx. 10%)
 7. Presentations (approx. 20%)

Before Next Time

- Read first papers below and write review
 - ***Epidemic algorithms for replicated database***, A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, D. Terry. In ACM Symposium on Principles of distributed computing (PODC), August 1987, pages 1-12.
<https://dl.acm.org/doi/10.1145/41840.41841>
- Check website for updated schedule