

Course Information

§ Course Overview (or – What Will I Learn from this Course?)

This course will equip students with foundational basics needed to understand, work with, and build distributed systems such as cloud computing systems. Topics include, but are not limited to, MapReduce, distributed hash tables (peer-to-peer systems), failure detectors, synchronization, election, distributed agreement, inter-process communication, consensus, gossiping, concurrency control, replication, key-value stores, NoSQL, security, probabilistic protocols, measurements, etc. These topics are discussed in the context of real-life and deployed systems such as clouds and datacenters, databases, peer to peer systems, clusters, etc. This course does not deal with details of computer networking and routing and wireless/mobile computing – other courses in the department (e.g., CS 438/439) cover those.

§ Course Essentials

Lecture: Tuesday and Thursday, 2:00 PM - 3:15 PM, 1320 DCL (Digital Computer Laboratory).

Course Website: <http://courses.engr.illinois.edu/cs425/>

All lectures and some announcements will be posted on this website, so please check periodically.

Discussion Forum: We will be using Piazza (link is on the website). Important announcements will be posted on this forum, so please check periodically. This is your go-to place to ask clarification questions for HWs (Homeworks) and MPs (Machine Programs, or Programming Assignments). Before you post, check if someone has already asked that question.

The course staff will attempt to ensure your questions are answered within 24-48 hours of it being posted. This does not mean you should post questions at the last minute, e.g., before a deadline.

Prerequisites:

- For all students: CS 241 (Systems Programming) or ECE 391, or equivalent course on Operating Systems or Networking (approval of instructor required if you've taken only a Networking course but not an OS course).
- For 4 cr On-campus and 4 cr MCS-Online students: Prior experience with a major programming language (for the MPs). In the past, students have used Java, C++, Go, C, etc. (one of these suffices - you will have a choice).
- For MCS-DS students: Prior experience with C++ (for the online Coursera MPs).

Credits: 3 hours or 4 hours.

- For On-campus and MCS-Online students only 4 credit students do the MPs. If you want to do MPs, register/upgrade for/to 4 credits.
- All MCS-Students do the online Coursera MPs.

Main Textbook (Recommended, not Required): Coulouris, G., Dollimore, J., Kindberg, T., and Blair, G., *Distributed Systems: Concepts and Design*, Addison-Wesley, *Fifth Edition*, 2011, ISBN: 0132143011. [Recommended purchase – copies available at Illini Book Store. On reserve at Grainger Library]. We will refer to chapter, section, and problem numbers *ONLY* in the Fifth Edition. If you use an older edition, correct interpretation/translation of these numbers is solely the students' responsibility (no excuses).

Supplementary material are listed at the end of this handout.

Course Staff Contact Information and Office Hours: To get the fastest responses, please first post to:

1. Discussion Forum: All students (3 cr, 4 cr, MCS-DS) must use Piazza. See website for link.
2. If your question is not answered on the discussion forum, then email the staff mailing list. See website for address. DO NOT email individual TAs, as the mailing list will get you the fastest response.
3. Use office hours of the TA or professor. See website for dates and times.
4. For private matters, you can contact the professor.

For more details, please see course website.

§ Course Participation

Sections and Assignments:

There are 3 categories of students (i.e., sections) in this class. They are listed, and their assignments summarized in Table 1 below. More description is below.

The three categories (sections) of students and their assignments are:

1. On-campus and online Students 3 credits: Homeworks, Midterm Exam, Final Exam. We will refer to these as "**3 credit students**". These include all 3 credit on-campus and online students *except* MCS-DS students. Where necessary, we will say "on-campus" or "online" to distinguish between the subcategories (neither of these include MCS-DS students).
2. On-campus and online Students 4 cr: Homeworks, Midterm Exam, Final Exam, On-campus Programming Assignments (MPs). We will refer to these as "**4 credit students**". These include all 4 credit on-campus and online students *except* MCS-DS students. Where necessary, we will say "on-campus" or "online" to distinguish between the subcategories (neither of these include MCS-DS students).

	3 cr (On-campus and Online)	4 cr (On-campus and Online)	MCS-DS
HWs 1-4 On-campus	Y	Y	Y
MP On-campus	N	Y	N
Coursera Quizzes 5 + 5=10 Coursera Finals: Part 1 + Part 2	N	N	Y
Coursera Programming Projects Part 1 + Part 2	N	N	Y
On-campus Midterm Exam + On-campus Final Exam Coursera	Y	Y	Y

Table 1: Assignments for Sections in CS425/ECE428. Y = Yes. N = No.

3. All **MCS-DS Students** (MCS-Data Sciences) **Please read carefully.** Students (coming in via Coursera): All Quizzes (5+5=10) on Coursera, All Exams on Coursera (2 Finals), 2 Programming Assignments on Coursera, All homework sets (given to On-campus students), Midterm Exam and Final Exam (both given to On-campus students). MCS-DS students *do not* do the 3-4 MPs given on campus. For each of the 5+5 *Coursera quizzes* (but *not Coursera exams*), we will have a *Pass/Fail grade (cutoff is 70% of the quiz points for that quiz)*. Overall, the work for the MCS-DS students is *equivalent to* (though not equal to) the 4 credit on-campus version of the course.

Assignments - Detail:

- (For All Students) There will be four to five homework sets (HWs) , with about 1-2 weeks turnaround time per homework. **Your homework solution submissions are required to be typed and printed** (you may use any of your favorite word processors). We will not accept handwritten solutions. Figures and equations (if any) may be drawn by hand. Homeworks will be **due at the beginning of class on the day of the deadline**. All on-campus students must submit printed HWs. The only exception is for MCS-DS students who can upload their HWs via Coursera.
- (For all 4 credit Students, Except MCS-DS Students) Four programming assignments (Machine Programs, or MPs) will be given throughout the semester, each requiring 2-4 weeks of effort. All MPs are evaluated via 1) a **report**, and 2) a live **demo** (only exception: groups with only online 4 cr students can submit video demos). All groups will receive access to a cluster (typically the CS VM Farm cluster). All groups can choose their programming language of choice (past groups have used C++, Java, Go, C, Python, etc.).
- (For Only MCS-DS Students) 2 Programming Assignments Building Distributed Systems via a C++ Emulator. **Prior knowledge of C++ is required and essential.**
- (For Only MCS-DS Students) All Quizzes (5+5=10) on Coursera, All Exams on Coursera (2 Finals). The 5+5 quizzes will each be pass/fail (cutoff is 70% of points on that quiz). Coursera exams are not pass/fail.

5. (For All Students) In-class Midterm Exam and Final Exam. MCS-DS Students will need to use a proctor for the exams.
6. **Start** your HWs and MPs (and for MCS-DS students: quizzes, programming assignments, and exams) **early**. DO NOT start at the last minute – these assignments involve more time than you think. If you start an assignment the day before the deadline, you're already too late.

Grading (tentative splits):

For the 3 groups of students:

1. 3 credit students will have total course grades out of 90 points (this group will be curved and graded separately).
2. 4 credit students will have total course grades out of 120 points (this group will be curved and graded separately).
3. MCS-DS students will have total course grades out of 156.75 points (this group will be curved and graded separately).

Grading (tentative splits) - Detail:

- On-campus Homework sets (HWs) 40 points (**For all students:** 3 credit and 4 credit students, and MCS-DS students)
- On-campus Programming Assignments (MPs) 30 points (**Only for 4 credit students, not for 3 credit or MCS-DS**)
- On-campus Midterm Exam 15 points (**For all students:** 3 credit and 4 credit students, and MCS-DS students)
- On-campus Final Exam 35 points (**For all students:** 3 credit and 4 credit students, and MCS-DS students)

Additional work **for MCS-DS** students:

- Coursera Quizzes (5+5=10 quizzes): $67 + 60 = 127$ questions. With 0.25 points per question (on average), that is a total of 31.75 points. (The "+" indicates the split between C3 Part 1 and C3 Part 2. C3 = Cloud Computing Concepts course on Coursera.). Note that individual questions are not graded, but each Coursera quiz is pass/fail with cutoff of 70%. E.g., If a quiz has 10 questions, you either get 2.5 points if you get 7 or more questions correct, or 0 points otherwise.
- Coursera Final Exams: $30 + 30 = 60$ questions. With 0.25 points per question, that is a total of 15 points.
- Coursera Programming Assignments: $10 + 10 =$ total of 20 points.

- While the Coursera points for each of the above may be different (e.g., Programming Assignments are 90 points on Coursera), they will be scaled down and weighted according to the numbers above.

Grades will be available in Compass 2g for on-campus and online students (3 cr and 4 cr), and for MCS-DS students on Coursera.

Grading for undergraduate and graduate students will be separated. Grades will be assigned on a curve (relative grading). The fraction of students receiving A's is not fixed a priori, and depends on the overall class performance.

On-campus 3 credit vs 4 credit Students: The MPs are only for 4 credit students. We will be grading and evaluating only MPs submitted by 4 credit students. We will not be grading or giving extra credit to MPs by 3 credit students or MCS-DS students. If you're a 3 credit student and wish to do the MPs, please upgrade yourself to 4 credits.

MP Groups:

- All MCS-DS students must do their (Coursera) programming assignments individually (see "Academic Integrity Policy" on the next page of this document).
- All other 3 credit and 4 credit students: MPs must be done only in **groups of 2**. Groups of size 1, 3 (or more), or zero are not accepted. Here's your chance to practice pair programming! Please read the "Academic Integrity Policy" on the next page of this document.

HW Groups: There are no HW groups. All HWs must be done individually. You cannot discuss solutions or ideas with anyone else other than course staff (but you can discuss lecture concepts and the HW question itself). Please read the "Academic Integrity Policy" on the next page of this document.

Lecture Participation:

- For on-campus students: Attending the lectures is important. Some lectures may have a quiz or in-class homework solving (included in your grade) - if you miss the class, you've lost the points.
- For MCS-DS students: viewing all lectures is a must. Do not attempt quizzes, homeworks, or exams without viewing and studying lecture material carefully.

§ Course Policies

Academic Integrity Policy We adhere by the CS academic integrity policies outlined at the webpage <https://wiki.cites.illinois.edu/wiki/display/undergradProg/Honor+Code>. It is the course policy that all of the work you submit for grading, or in support of graded material, as an individual or project group, shall be your own product, from inception to completion. The only resources you can

avail of in your HWs and MPs are the provided course materials (slides, textbooks, etc.), and communication with instructor/TA via newsgroup and email. Please do not reveal solutions on any of these fora. Exams are closed-book, closed-notes, unless otherwise specified.

All homeworks are individual. Other than with course staff, you should discuss, or share, or work together on, **solutions with no one else** (whether in the class or outside). You can however discuss with other students **lecture concepts** or **the HW question itself** (but not the solution).

For MP groups, other than with the course staff, you can discuss, share the **solution** (or ideas or code or whatever), or work together on, the MP, **only with your group partners**. You cannot discuss, work on, or share, the MP solution with anyone else. You can however discuss with other students **lecture concepts** or **the MP spec itself** (but not the solution).

Please use the Piazza discussion forum for all questions/discussion on HWs and MPs. **All students, on-campus, online and MCS-DS must use Piazza for all course-related questions. MCS-DS students – if you use Coursera forums to post questions, your question may be answered very late (or never). So please use Piazza.**

Please note however, that on the Piazza forum, if you post a solution (code or write-up) to the forum, you will lose all points for that particular assignment. Email the staff only when you cannot use the discussion forum, e.g., if you have an urgent question, or if you have a personal matter to ask/discuss.

Result of Cheating Violations of this academic integrity policy will be treated seriously. A first violation of these academic integrity guidelines will result in a zero grade on that HW/MP. A second violation will result in an F grade for the entire course. (And yes, in the past, we have caught almost all cheating students. And these cheating students have suffered both above types of consequences.)

Just don't cheat – your time is not worth it. If you cheat in this course and get away with it miraculously, later in life, when you are working in a job and need to use these concepts, you will be at a loss (Karma comes around).

Policy on Late Submission: Unless otherwise specified, **all MP assignments**, or components thereof, that are to be electronically submitted are **due at 11:59 PM US Central Time** on the due date. Similarly, **homework sets**, or components thereof, that must be submitted by other means are **due at the beginning of class** on the due date.

All MP and HW submission deadlines are hard and will not be extended. No late homework sets or MPs will be accepted except under extremely rare non-academic circumstances (which usually require approval from the Dean's office). So please ensure you have submitted everything necessary by the deadline time.

§ Supplementary Textbooks

We will use many readings off the Web - these will be available under the Lectures link on the course website. The following textbooks may be used for supplementary course material. You are not re-

quired to own these books. They are on reserve at the Grainger Library (if available). You *may* be able to use a prior edition of the textbook if you own one.

1. Fourth edition of the Coulouris-Dollimore-Kindberg textbook.
2. "Distributed Systems: An Algorithmic Approach," Sukumar Ghosh, CRC Press, 2006, ISBN: 1584885645. (Available online free at the Illinois Library Catalog).
3. "Distributed systems: principles and paradigms," A. Tanenbaum and M. Steen, Prentice Hall, Second Edition, 2005, ISBN: 0132392275.
4. "Distributed algorithms: concepts and design," N. Lynch, Morgan-Kaufmann, 1ed, 1996, ISBN: 1558603484.

§ **Copyright** All material used in this version of the course is copyrighted by Indranil Gupta and the University of Illinois. This version of the course has evolved based on past lectures given by the following professors over the past 15 years: Jennifer Hou, Mehdi Harandi, Nitin Vaidya, Klara Nahrstedt, and Sayan Mitra.