
CS 425/ECE 428 Fall 2025 – Practice Midterm

(This practice midterm is structured in the same way, and for the same time as, the regular midterm exam.)

This midterm contains 5 questions, each worth 20 points.

1. On-campus students and MCS-Chicago students and City Scholars students will be given printed copies of the exam, and will have to write answers on the exam booklet itself. It is a closed-book, closed-notes exam (you are not allowed to bring extra pages, NO cheatsheets).
2. Coursera students (DSO section) have an online open-book open-notes unproctored exam. Coursera students must type their answers in a separate blank file (just like you do with Homework solutions that you submit), and submit this file as their solution. Don't forget to TAG each question with its corresponding pages in your solution doc (otherwise they will not be graded)!
3. On-campus students get 75 minutes (since exam is in-person)
4. Coursera students get 90 minutes (since exam is online), plus a small buffer time for submitting. Please start a timer for 90 minutes when you start the exam, and submit well before 90 minutes is up. We DO NOT accept emailed submissions. If you miss the window to submit, you will get a zero on the midterm. You can submit multiple times until the timer expires.
5. This exam is formatted as if for on-campus students. Coursera students will receive a similar exam copy (without all the white space to write, since you will type your answers on your own separate document).

1. Short Multiple Choice [10 X 2 = 20 points]

There are 10 questions below. For each question, **select the ONE BEST** answer among the choices. No justification required. No partial credit.

1. In an asynchronous system, which of the following is TRUE?
 - a. Messages are not delayed beyond a pre-fixed known latency
 - b. Clocks of different processes are synchronized to within a known bound
 - c. Failures are not allowed to occur
 - d. None of the above
2. A Mapreduce task is running on a cluster with 4 racks, each with 2 machines. The machines are named as S<rack number><machine number>: S11, S12, S21, S22, S31, S32, S41, S42. A Map task needs to access as input a block that has replicas on machines S21, S22, S41. Where will the task be scheduled, if the only free containers in the cluster are available at the following machines: S11 and S42
 - a. S11
 - b. S42
 - c. Somewhere
 - d. Nowhere – this task cannot be scheduled
3. A datacenter run by an upstart company called DataStones consumed about 100 kWh so far in 2015 AD. If only 80 kWh of this 100 kWh was used in running IT equipment (servers, routers, etc.), then the PUE of DataStones' datacenter is:
 - a. 0.8
 - b. 1.25
 - c. 1.33
 - d. 300
 - e. $(100-80)/80 = 0.25$
 - f. None of the above
4. A BitTorrent client C is trying to download a file with four blocks 0 through 3. C is talking to three neighboring nodes (clients) A, B, D. Client A currently is storing blocks 0, 2, 3; while B is storing blocks 0, 1, 2; while D is storing blocks 2, 3 only. Which blocks should the client C fetch FIRST?
 - a. 0
 - b. 1
 - c. 2
 - d. 3

-
5. In a system of three processes sending unicasts, an event e_1 has a vector timestamp of (1, 2, 30) while an event e_2 has a vector timestamp of (0, 2, 301). Which of the following statements is TRUE?
 - a. e_1 happened before e_2
 - b. e_2 happened before e_1
 - c. e_1 and e_2 are concurrent
 - d. We can't tell what the causality relation between e_1 and e_2 is
 6. In Hadoop, which of the following entities is responsible for scheduling decisions?
 - a. AM
 - b. RM
 - c. NM
 - d. M&M
 - e. All of the above
 7. A set of Cassandra clients uses the QUORUM consistency level for both reads and writes. Then, assuming no failures:
 - a. Once a client C has received an ack for a write W , any subsequent reads by C will see the value written by W or subsequent writes
 - b. Once a client C has received an ack for a write W , any subsequent reads by any client will see the value written by W or subsequent writes
 - c. Both above statements are true
 8. A Cassandra cluster is running a RackInferring snitch. You are reviewing a log that contains IP addresses 111.121.213.122 and 111.122.213.122 and 111.123.213.56. These IP addresses:
 - a. Are the same machine
 - b. Are in the same rack in the same datacenter, but are different machines
 - c. Are in different racks but in the same datacenter
 - d. Are in different datacenters
 9. Which of the following systems prefers availability over consistency, under partitions?
 - a. Cassandra
 - b. HBase
 - c. A traditional relational database (e.g., MySQL).
 - d. None of the above
 10. A modified gossip protocol uses a spanning tree to get the multicast to half of the processes. This takes $O(\log(N))$ time. Thereafter you have the choice of either pull or push gossip. Which will complete faster, i.e., reach a given fraction of infected processes with the same given probability, but earlier?


- a. Push
- b. Pull
- c. Both complete at about the same speed (or time)

(Go to next page)


2. Lamport Timestamps [10 + 10 = 20 points]

This question has two parts.

- (a) In the figure below, events include message sends and receipts (M^*), and local steps (E^*). Assign *Lamport timestamps* to each event. The initial timestamp of each process is 0. (Coursera students – please list the timestamp for the events at each process. Please include event names against timestamps! Use $\text{Send}(M_{xx})$ and $\text{Receive}(M_{xx})$ to distinguish those two events for a given message.) The initial timestamp of each process is 0.



- (b) Now mark vector timestamps on the same timeline (below).



3. Chord Routing [4 + 8 + 8 = 20 points]

This question has three parts.

In a Chord P2P system, there are 256 points on the ring. Six machines join, with the following peer ids: 1, 10, 100, 150, 200, 250.

(a) Indicate the successor entries of each node on the Chord ring.

(b) Show the finger table entries of machine with peer id 10.

- (c) Use the Chord routing algorithm to route a message for key (file id) 2 starting from peer id 150. There is no file replication, and each peer has exactly one successor and no predecessors.

4. MapReduce [20 points]

You are given a social network log that captures information about all the posts on a social network during the course of a day. Each line contains tuples (a, hh:mm:ss) where a is a user id, and hh:mm:ss is the timestamp of a post made by the user. If a user has multiple posts, they will appear as separate entries. The entries are not sorted in any order.

Write a MapReduce program to find those 3 hours of day when the maximum, minimum and median number of posts were made. Your output should be three integers, each in the range 0-23.

You should have at least some parallelism. Ensure that your output does not contain duplicates. You can set your key and value to arbitrary objects. You cannot retain data at any of the machines from a task (Map or Reduce) for use in a later task. Chaining MapReduces is allowed, as long as you don't over-use chaining (where parallelization could have helped instead). Each MapReduce in a chain can read the dataset, but there is no other persistent memory across the chain. Pseudocode is preferable, and it can be coarse-grained, e.g., you can say "get top k from this list by field x", "sort this list by key x", etc.). Be clear and concise and don't miss any steps.

5. Key-Value Stores [10+10 = 20 points]

This question contains two parts.

- (a) Someone tells you that Cassandra should really be using the finger tables from Chord as well. Give three reasons arguing against this proposal.

- (b) Instead of using the Bloom filter along with an SSTable why not simply use a list of the keys present in that SSTable?

(You can continue writing here)

(You can continue writing here)

SCRATCH SPACE