

COSC-211: DATA STRUCTURES MIDTERM 2 TOPICS AND LOGISTICS

1 Logistics

- This midterm will be asynchronous; you'll be able to complete it at your preferred time within the exam window.
- The midterm will be available on Gradescope from Thursday, April 29, 12:00am ET until Friday, April 30, 11:59pm ET (i.e., the 48-hour period including all of Thursday and all of Friday, ET).
- To ensure that you're familiar with the Gradescope setup and that you've practiced the process of scanning and uploading your solutions, we're again going to have a practice run. From Thursday, April 22, 12:00am ET until Tuesday, April 27, 11:59pm ET, there will be a sample midterm available on Gradescope with the same setup as next week's real midterm. This will give you an opportunity to see what the Gradescope system will look like, and to practice submitting your solutions. This time around, the practice run is **required**.
- This is a timed exam; you'll have 90 minutes to solve the problems and upload your work to Gradescope. We have written this exam with the expectation that it will take you about 50 minutes to solve the problems. We are giving you 40 extra minutes in case of technical difficulties, in case you want to print out the problems, etc. Please budget your time accordingly.
- Your 90 minutes will start as soon as you open the midterm on Gradescope. Gradescope won't let you submit your work after the 90 minutes are up. If you run into a technical problem (e.g., your internet goes out) and you can't submit your work on time, **don't panic**. Send your instructor (Prof. Gardner or Prof. Kaplan) an email right away, attaching your work, and we'll figure it out. Note, however, that the purpose of the required practice run is to minimize the chance that you will run into technical problems or run out of time to upload your work.
- If you use extended time for exams, please email your instructor as soon as possible—and no later than the morning of Wednesday, April 28—to let us know (a) that you use extended time, and (b) what % time you use. Send an email even if you think we already know! We need to manually configure this in Gradescope, and it will not happen unless we hear from you.

2 Intellectual Responsibility and Asking Questions

- You're responsible for doing and submitting your own work on this midterm. Talking with your classmates, friends, family, etc. about the exam problems is not permitted. Because students will be taking the exam at different times during the exam window, we highly recommend not discussing the problems at all, with anyone, until Saturday.

- The midterm is open notes, meaning that you're allowed to look at your notes and anything posted on the course Moodle page or Slack channel while you're taking the midterm.
- As is the case for all work in this class, do not go to the internet with the express purpose of looking for solutions.
- Unfortunately, we cannot be available during the entire 48-hour exam window to answer questions. To ensure equitable midterm conditions for all students, including those who are in distant time zones or who may opt to take the midterm late at night, we are not going to answer any exam-related questions during the exam window. We will try to make the questions as clear and unambiguous as possible. If you're truly unsure of what a question is asking you to do, you can write down the question that you want to ask and your best guess at how we'd answer it, and we may take this into account when grading.

3 Topics and Structure

The midterm will have five problems. You're responsible for all content up through splay trees. Topics include:

- ADT definitions: Dictionary
- Data structures: unbalanced binary search trees, hash tables (with chaining and with linear probing), bloom filters, red-black trees, AVL trees, splay trees

For unbalanced binary search trees, hash tables, and red-black trees, you should know how to do the main operations (add, remove, lookup, in-order traverse). For bloom filters, you should know how to add and lookup, and you should understand the one-sided error for lookup. For AVL trees, you do not need to be able to do the operations, but you should be able to recognize if a tree is an AVL tree. For all data structures, you should be familiar with the worst-case and average-case runtime results (and the amortized result, for splay trees).

The exam is not cumulative in the sense that the focus will be on the material we've covered since the first midterm. It is cumulative in the sense that some of the material builds on itself (e.g., you should still understand the definition of big- O and be able to determine the runtime of a new operation).

Want practice problems? Open Data Structures has lots of good exercises. Feel free to give some of these a try, and discuss them amongst yourselves on Slack!