

Final Exam Study Guide

May 19
Spring 2021

Your final will be an “open book” exam. The exam will be held on Tuesday, May 25 at 9:30 am. You will have 2.5 hours to complete the exam.

You are responsible for anything we covered in class or in lab, as well as any topics discussed in the lecture videos/slides, conferences, and handouts. The exam is cumulative, but it will heavily weight topics from the second half of the course. Note, however, that we used arrays, `Vectors`, and `Lists` to implement many data structures; we used Big-O notation to evaluate and compare data structures; we used recursion to traverse trees; we used induction to prove properties of data structures and algorithms; and so on. These concepts are unavoidable and the second half of the semester built heavily upon them.

The following non-exhaustive list may be helpful in reminding you about some of the key topics we have covered:

- **Pre-Midterm**

- Java syntax, as we have used it in our programming assignments.
- Classes, abstract classes, and interfaces and their respective roles.
- Information hiding (abstraction) and why it’s good.
- Extending classes with inheritance.
- Generic classes and their use.
- Pre- and post-conditions, and assertions.
- The meaning of `static` (and non-static) as applied to variables and methods.
- `Vector`, its implementation in the `structure5` package, and its methods.
- Complexity: Big “O” definition.
 - * Determining the asymptotic behavior of mathematical functions.
 - * Determining the time and space complexity for a given algorithm.
 - * Worst and best case analysis.
- Linear and binary search.
- Recursion and induction.
- Sorting.
 - * Bubble sort, selection sort, insertion sort, merge sort, quicksort (and heapsort post-midterm).
 - * Using `Comparator/Comparable` for sorting.
- Linked lists: Singly, Doubly, and Circularly linked lists

- **Post-midterm**

- Stacks (LIFO)
 - * List and Vector implementations
 - * Relationship with recursion and graph/tree traversals (DFS)
- Queues (FIFO)
 - * List, Vector, and fixed-size array implementations
 - * Relationship with graph/tree traversal strategies (BFS)

(continued on the flip side...)

- Priority queues
 - * OrderedVector Implementation
 - * Heap implementation
 - heap property
 - array representation and tree fullness/completeness
 - heap insert/remove
- Trees
 - * Recursively-defined, link-based representation
 - * Array/Vector-based representation
 - * Binary search trees
 - Binary search tree property
 - * Traversing trees: In-order, post-order (depth-first), pre-order; level-order (breadth-first)
 - * Tree balancing schemes (Red-Black, AVL)
- Iterators (including Iterators based on other iterators)
- Bitwise operations
- Hashtables
 - * Using a hashing function to find an object's "bin"
 - * Open Addressing vs External Chaining
 - * Managing collisions
 - * Hash table load factor
- Graphs
 - * Directed/undirected
 - * Weighted/unweighted
 - * Adjacency List representation
 - * Adjacency Matrix representation
 - * Reachability/traversal (Breadth-first, depth-first)
 - * Minimum cost spanning trees and Prim's algorithm
 - * Single source shortest paths and Dijkstra's algorithm

Tips on Preparing for the Exam

Our goal is to test concepts, so it is *not* important to memorize the exact code or method signatures for every data structure. Descriptive variable/method names are enough to demonstrate understanding. However, it *is* important to know the types of operations that different data structures do/do not support and how efficient (in the Big-O) sense these methods are for the various implementations that we have studied. For example, we cannot access arbitrary elements in a queue: we can only add to the back and remove from the front. And the time required to add to/remove from a queue depends on the implementation: if, for example, a queue is implemented using a Vector, dequeuing involves removing from the front of the Vector and so takes $O(n)$ time. Using a circularly linked list, we can add to the tail and remove from the head in $O(1)$ time.

If you understand how each data structure is implemented, then you should be in a good position to figure out the space required to store the structure as well as the time complexities of its methods. This is a much better approach than trying to memorize run-times of all of the methods of all of the data structures—it is also much more valuable to you in the future. Similarly, understanding how an algorithm works, for example Prim's algorithm, is much more valuable than trying to memorize the pseudo-code (or—gasp—Java code) for the algorithm. To test your understanding of an algorithm, see if you can describe what it does to a friend and try to execute it by hand on small examples.

And, as they say, practice makes perfect. We have posted a variety of sample questions on the course website. Try to work through each problem, first alone, then—if needed—with a classmate. Try to do a few problems from the text. Answers to odd-numbered book questions can be found in the appendix, so that you can check your solutions.

We've been impressed by your hard work and all you've learned this semester. Study hard but don't stress—and good luck!