

The Final exam is Tuesday May 1st from 8:00 - 9:50 AM in Wright 9. It will be closed-book and notes, except for three 8" x 11" sheets of paper containing any notes that you want. (Plus, the Python Summary Handout) About 75% of the test will cover the following topics (and maybe more) since the second mid-term test, and the remaining 25% will be comprehensive (mostly big-oh analysis and general questions about stacks, queues, priority queues/heaps, lists, and recursion).

Chapter 6: Trees - coding question only over Ch. 6
Terminology: node, edge, root, child, parent, siblings, leaf, interior node, branch, descendant, ancestor, path, path length, depth/level, height, subtree
General and binary tree recursive definitions
Tree shapes and their heights: full binary tree, balanced binary tree, complete binary tree
Applications: parse tree, heaps, binary search trees, expression trees
Traversals: inorder, preorder, postorder
Binary search tree ADT: interface, implementation, big-oh of operations
Balanced binary search trees: AVL tree ADT: interface, implementation, big-oh of operations

File Structures - Lecture 24 handout:

http://www.cs.uni.edu/~fienup/cs1520s18/lectures/lec24_questions.pdf

We talked about how the in memory data structures need to be adapted for slow disks.

From this discussion you should understand the general concepts of Magnetic disks:

- layout (surfaces, tracks/cylinders, sectors, R/W heads)
- access time components (seek time - moving the R/W heads over the correct track, rotational delay - disk spins to R/W head, data transfer time - reading/writing of sector as it spins under the R/W head)

Hash Table as a useful file structure

B+ trees as a useful file structure - see web resources:

<http://www.sci.unich.it/~acciaro/bpiutrees.pdf>

http://en.wikipedia.org/wiki/B%2B_tree

<http://www.ceng.metu.edu.tr/~karagoz/ceng302/302-B+tree-ind-hash.pdf>

Chapter 7: Graphs

Terminology: vertex/vertices, edge, path, cycle, directed graph, undirected graph

Graph implementations: adjacency matrix and adjacency list

Graph traversals/searches: Depth-First Search (DFS) and Breadth-First Search (BFS)


General Idea of the following algorithms: topological sort, Dijkstra's algorithm (single-source, shortest path), Prim's algorithm (determines the minimum-spanning tree), TSP (Traveling-Saleperson Problem)

Approximation algorithm to solve TSP, general idea of backtracking and best-first search branch-and-bound.


You should understand the graph implementations and algorithms listed above. You should be able to trace the algorithms on a given graph.

AVL add 50, 70, 90, 80, 30, 40, 25, 10, 5

pivot node \rightarrow 50 ~~2~~ ~~2~~ ($h+LST - h+RST$)




$1.44 \log_2 n$



9. A B+ Tree is a multi-way tree (typically in the order of 100s children per node) used primarily as a file-index structure to allow fast search (as well as insertions and deletions) for a target key on disk. Two types of pages (B+ tree "nodes") exist:

- Data pages - which always appear as leaves on the same level of a B+ tree (usually a doubly-linked list too)
- Index pages - the root and other interior nodes above the data page leaves. Index nodes contain some minimum and maximum number of keys and pointers bases on the B+ tree's *branching factor* (b) and *fill factor*. A 50% fill factor would be the minimum for any B+ tree. All index pages must have $\lceil b/2 \rceil \leq \# \text{ child} \leq b$, except the root which must have at least two children.

Consider an B+ tree example with $b = 5$.




a) How would you find 88?

b) The insert algorithm for a B+ tree is summarized by the below table. Where would you insert 50, 100, 105, 110, 180, 200, 210?

Situation		insertion Algorithm
Data Page Full?	Parent Index Page Full?	
No	No	Place record in sorted position in the appropriate data page.
Yes	No	<ol style="list-style-type: none"> 1. Split data page with records $<$ middle key going in left data page and records \geq middle key going in right data page. 2. Place middle key in index page in sorted order with the pointer immediately to its left pointing to the left data page and the pointer immediately to its right pointing to the right data page.
Yes	Yes	<ol style="list-style-type: none"> 1. Split data page with records $<$ middle key going in left data page and records \geq middle key going in right data page. 2. Adding middle key to parent index page causes it to split with keys $<$ middle key going into the left index page, keys $>$ middle key going in right index page, and the middle key inserted into the next higher level index page. If the next higher index page is full continue to splitting index pages up the B+ tree as necessary.

Stack implementation



push pop peek
 $O(n)$ $O(n)$ $O(1)$

