



READ THIS NOW!

- Print your name in the space provided below.
- There are 5 short-answer questions, priced as marked. The maximum score is 100.
- This examination is closed book and closed notes, aside from the permitted one-page fact sheet. Your fact sheet may contain definitions and examples, but it may not contain questions and/or answers taken from old tests or homework. However, you may include examples from the course notes.
- No calculators, cell phones, or other computing devices may be used. The use of any such device will be interpreted as an indication that you are finished with the test and your test form will be collected immediately.
- Until solutions are posted, you may not discuss this examination with any student who has not taken it.
- Failure to adhere to any of these restrictions is an Honor Code violation.
- When you have finished, sign the pledge at the bottom of this page and turn in the test and your signed fact sheet.

Name (Last, First) _____ **Solution** _____
printed

Pledge: On my honor, I have neither given nor received unauthorized aid on this examination.

_____ *signed*



xkcd.com

Answers to questions are in blue.

Commentary about the questions and answers is in green.

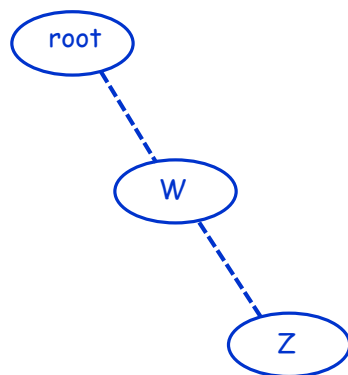
1. Suppose we have a BST and that two values, X and Y , have been inserted into the tree so that Y is the left child of X , and Y is in a leaf node. Could there be a value, Z , which was also inserted into the BST such that $Y < Z < X$? Let's consider some (but not all) cases; what can you conclude about each of the following scenarios? (The BST may contain other values as well, but not duplicates.)

So, what we know is that two values, X and Y , are in the BST, Y is the left child of X , and Y is a leaf. And, since Y is a child of X , X was inserted before Y . Those are absolute facts. The question is: could there be another value Z , that is somewhere in the BST, such that $Y < Z < X$?

- a) [10 points] Suppose that Z was inserted sometime before both X and Y were inserted.

If Z was inserted before both X and Y , then X would have gone right at (or above?) Z and Y would have gone left at (or above?) Z , so the arrangement described earlier could not have occurred.

But... must X (or Y) actually reach Z when they are inserted? Could X have taken a different direction during its insertion, and not be in Z 's right subtree? Suppose there is some value, say W , along the path from the root to Z such that X went in a different direction at W than Z did. Then X must have gone right at W and Z must have gone left at W :



Now, if X went right at W , then Y will go left at W , which contradicts the fact that Y is a child of X . Therefore X must have reached Z during its insertion (and so must Y).

- b) [10 points] Suppose that Z was inserted sometime after X was inserted, but before Y was inserted.

Now, we might ask: could Z have gone in a different direction and not have reached X at all? Interesting question... which I will leave for you to ponder.

If Z was inserted after X but before Y , then IF Z reached X , X could not already have a left child since Y is supposed to wind up being X 's left child. Therefore, Z would now become the left child of X , which is a contradiction.

2. Suppose we have a value, X that has been inserted into some data structure. Suppose we delete X from the structure, and then immediately reinsert X into the structure. Could the cost of finding X before it was deleted be different than the cost of finding X after it has been reinserted? If yes, would there always be a difference, or only in certain cases?

The big- O results for the data structures are not relevant here. The question is really simpler than that. You need to consider whether, when X is reinserted, it must wind up in the same spot it was before, or it might wind up in a "better" or "worse" spot.

The question was not, at all, about the cost of the reinsertion of X .

- a) [10 points] What conclusion follows if the data structure in question was a PR quadtree? Explain clearly.

There will be no difference:

The position of X in the quadtree is determined entirely by how much splitting must be done to separate X from any nearby points (even if buckets are used).

So, X will be reinserted at precisely the same position it was deleted from.

I was picky about this. You needed to offer an explanation of why the insertion process for a PR quadtree will place X back in the same (node) position. I didn't give full credit if you just said "PR quadtree structure is independent of the order of insertion". The question did not specify bucket size. If you took that into account, your conclusion should have been that the search cost might be a little higher since X might move to the end of the bucket.

- b) [10 points] What conclusion follows if the data structure in question was a BST? Explain clearly.

There may or may not be a difference:

If X was in an internal node of the BST, then X will be in a leaf node when it's reinserted.

On the other hand, if X was in a leaf node, reinserting X will put in a leaf node in precisely the same position.

- e) [10 points] What conclusion follows if the data structure in question was a hash table that uses probing to resolve collisions? Explain clearly.

There may or may not be a difference:

When X is deleted, it's position will become a tombstone.

When X is reinserted, the probing logic could put X into an earlier slot in the probe sequence (if tombstones are recycled, and there is an earlier tombstone in the probe sequence), or into the same slot it occupied before (if tombstones are recycled, but there's no earlier tombstone in the probe sequence), or into a later slot in the probe sequence (if tombstones are not recycled, although that's unlikely).

Tombstones **MUST** be used on deletion operations if the hash table uses probing, no matter what form of probing is used. And, tombstones will always be recycled because doing so costs you an insignificant amount of extra work, but also makes some future searches more efficient.

3. [10 points] Suppose we have a PR quadtree that is storing data objects with coordinates in the range [0, 128], and that subtree currently looks like Figure 3A below. The node labelled **Parent** is an internal node; it may or may not be the root of the tree. Now, suppose that the data object **Y** is deleted from the tree, and the implementation results in a branch contraction, making the subtree look like Figure 3B below.

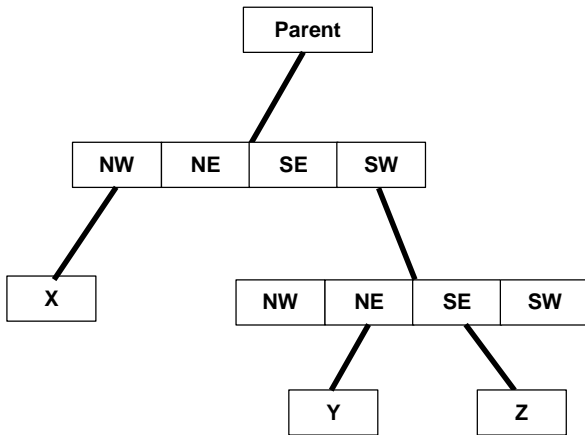


Figure 3A


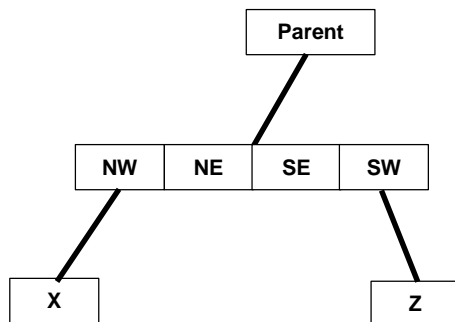


Figure 3B

Is this correct? If not, explain why it is wrong, and draw the subtree (below) as it should look after the deletion. If this is correct, carefully explain why. Either way, your explanation should refer to regions and partitioning.

This is incorrect. Z lies in the SE quadrant of the SW quadrant of the parent of X. Therefore, Z lies in the SW quadrant of the parent of X. The resulting tree should look like this:



4. Suppose you have a large collection of data values. You are considering storing the data values in a data structure, and are trying to decide between a hash table and a minimum-height BST (which you could build since you have all the data values). You are convinced that no matter which data structure you choose, you can achieve its theoretically optimal performance. You can also implement any additional public methods that seem helpful, no matter which data structure you choose.

There were a couple of common misconceptions in the answers:

Some answers said things like "hash tables are not good for strings" or "hash tables are not good for integers". I have no idea where such ideas came from.

Many answers assumed that the data for the search you were performing was available before the relevant data structure was constructed. This was often implicit, as in "let's use a hash table of dimension K " in part a). That doesn't make sense. If you had that information in advance, and you were only going to perform one such search, there would be no reason to use a data structure at all. You'd just traverse the data elements linearly, $O(M)$ cost, which would be cheaper than building either data structure.

Another issue was to only consider the theoretical optimum costs for search in the data structure, and not take into account the nature of the search problem. If the big- O bounds were always sufficient to make the best choice, life would be much simpler. And, considering the specific problem we have can lead to clever ways of employing the structure.

- a) [10 points] Suppose your data values are integers, and that your primary concern is to support determining whether, for two given positive integers, N and K , there is an integer, M , in the collection such that $M \% K$ equals $N \% K$. Which data structure would be the better choice? Explain why.

Suppose there are I integers in the collection.

It seems the only way to search the BST or the hash table for a suitable M would be to perform a possibly full traversal of the tree ($\Theta(I)$) or of the hash table's array (at least $\Theta(I)$, and worse unless the table is of minimum size).

That seems to favor the BST slightly.

Unless N , or at least K , were known in advance, there's no better way.

- b) [10 points] Now suppose your data values are strings, and your primary concern is to support determining whether, for a given string S , there are any strings in the collection that are anagrams of S . For example, "cinema" is an anagram of "iceman". Which data structure would be the better choice? Explain why.

Suppose there are M strings in the collection.

Given S , we can compute all the rearrangements of the characters in S .

Given those, we can search for them, one by one, seeking a match. Each of those searches would take $\Theta(\log M)$ in the BST and $\Theta(1)$ in the hash table.

Therefore, the hash table is the better choice.

An optimization would be to compute one rearrangement, then search the structure for it, and then generate the next rearrangement, and so forth. Another would be to use the `sumhash()` function, which is otherwise lousy. In this case, that would put all the strings that contained the same characters into the same table slot (assuming we use chaining). So, all the anagrams of S , if any, would be in the home slot of S .

5. Haskell Hoo IV is attempting to implement a BST generic that conforms to the specification for Project 2. Being concerned that his logic may be untrustworthy, he has written the private member function shown below. He plans to call it at the end of each insertion and deletion operation while he is testing his BST.

```
// Pre:      root is null or points to the root node of a BST
// Post:     the BST is unchanged
// Returns:  true iff the BST ordering property is NOT violated somewhere
//           in the tree
//
private boolean hasBSTProperty( BinaryNode root ) {

    if ( root == null ) return true;

    if ( root.left != null &&
        root.element.compareTo( root.left.element ) <= 0 )
        return false;

    if ( root.right != null &&
        root.element.compareTo( root.right.element ) >= 0 )
        return false;

    return hasBSTProperty( root.left ) && hasBSTProperty( root.right );
}
```

One strange error was to say the function is wrong because it would classify an empty tree as a BST; an empty tree is a BST, so is a one-node tree.

Another common error was to worry about duplicate entries; the BST from Project 2 did not allow duplicate entries.

Another common error was to simply misinterpret the Java code. For instance, the test in the first if makes sure there IS a left child, and then checks whether the element in the parent node is less than or equal to the element in its left child. If so, you don't have a BST. The logic there is correct. And, there are appropriate NULL tests everywhere one is needed. And, there are no syntax errors in the code.

- a) [10 points] Could Hoo's function ever return false even though his tree did have the BST ordering property? That is, could Hoo's function ever report a false negative? Justify your answer; if yes, use a specific example of a BST to illustrate; if no, explain exactly why.

The only way that the function will ever return false would be to find a node holding a value, X, such that the value in the left child was larger than X or the value in the right child was smaller than X. Either way, we would not have a valid BST.

Therefore, no, the function will never report a false negative.

- b) **[10 points]** Could Hoo's function ever return true even though his tree did not have the BST ordering property? That is, could Hoo's function ever report a false positive? Justify your answer; if yes, use a specific example of a BST to illustrate; if no, explain exactly why.

The flaw in the function is that the checks are entirely "local"; that is, each node is only checked against its immediate children. The function would return true for the following tree, even though the BST property is violated:

