UNIVERSITY OF CALIFORNIA
Department of Electrical Engineering
and Computer Sciences
Computer Science Division

CS61B   P. N. Hilfinger and Josh Hug
Fall 2014

Test #2

READ THIS PAGE FIRST. *Please do not discuss this exam with people who haven't taken it.*
Your exam should contain 11 problems on 11 pages. Officially, it is worth 17 points (out of a total
of 200).

   This is an open-book test. You have 110 minutes to complete it. You may consult any books,
notes, or other inanimate objects available to you. You may use any program text supplied in
lectures, problem sets, or solutions. Please write your answers in the spaces provided in the test.
Make sure to put your name, login, and lab section in the space provided below. Put your login
and initials *clearly* on each page of this test and on any additional sheets of paper you use for your
answers.

   Be warned: our tests are known to cause panic. Fortunately, this reputation is entirely unjus-
tified. Just read all the questions carefully to begin with, and first try to answer those parts about
which you feel most confident. Do not be alarmed if some of the answers are obvious. Should
you feel an attack of anxiety coming on, feel free to jump up and run around the outside of the
building once or twice.

Your name: _____   Login: _____

Login of person to your Left: _____   Right: _____

Discussion TA: _____

   1. ____/1   2. ____/1   3. ____/3   4. ____/2   5. ____/2

   6. ____/2   7. ____/   8. ____/1   9. ____/2   10. ____/2   11. ____/1

1.   [1 point] We know that with a complete binary search tree with $2^n - 1$ nodes, we can perform searches in $\Theta(\lg n)$ time. All nodes in such trees have either zero or two non-null children. In contrast, in a binary search tree where every node has either zero or one children, searches take $\Theta(n)$ time. Suppose that you have a general binary search tree in which every node has either zero or two children. What is the worst-case search time in that case? What do worst-case trees look like?

2.   [1 point] Suppose we know that a certain type of object is going to be used as a key in a HashMap and looked up many times. Then one way to improve performance is to store the hash code of each such object in the object itself. Consider the following implementation of this idea for a wrapper class for the String class:

```java
public class EfficientString {

    private int _hash;
    private String _val;

    public EfficientString(String val) {
        if (val == null) {
            throw new IllegalArgumentException();
        }
        _val = val;
        _hash = val.hashCode();
    }

    @Override
    public String toString() {
        return _val;
    }

    @Override
    public int hashCode() {
        return _hash;
    }

    public void setVal(String val) {
        _val = val;
    }

    @Override
    public boolean equals(Object obj) {
        return obj instanceof EfficientString && _val.equals(obj.toString());
    }
}
```

a. Unit testing shows that in some cases, the return value of `hashCode` is wrong. What is the cause of this error and how can it be fixed?

b. Even after problem (a) is fixed (so that the hash value is always correct,) some users report problems using `EfficientStrings` in hash tables. What is the problem?

3. [3 points] In the table below, there are three pairs of sequences—a, b, and c. The first of each pair is the input to a sorting algorithm. The second is an intermediate sequence that might occur during the execution of the algorithm. Fill in each box of the table below with 'Y' or 'N' depending on whether the second sequence could or could not occur at any point during the execution of the indicated algorithms on the first sequence.

Note 1. For insertion sort, assume that items are swapped pairwise one-by-one towards the left of the array until they reach their final position, i.e. that we use the same swapping procedure that we used in lecture and in the optional HW8 solutions.

Note 2. For mergesort, assume that we merge from the "top down." For example, items 0 through 3 will be merged before items 4 and 5.

Note 3. For quicksort, we mean *any* valid partitioning method that divides the data into portions less than, equal to, and greater than a pivot, not necessarily the version from the homework (where elements stay in the same relative order after partitioning). Assume that the leftmost item is always used as the pivot.

Note 4. For heapsort, assume that the initial heapification process is performed by inserting all items in the array into a fresh (initially empty) heap. An intermediate sequence is considered valid even if it represents the state of an array in the middle of the re-heapification process after a delete (also called "sinking" or "heapify-down").

	Sequence sort	Insertion sort	Merge sort	Quicksort sort	Heap sort
a.	12 17 13 7 16 23 22 24 7 12 13 17 16 23 22 24				
b.	23 22 16 11 15 3 4 10 10 22 16 11 15 3 4 23				
c.	23 23 21 14 1 1 21 12 21 12 21 14 1 1 23 23				

Put 'Y' or 'N' in each box: 'Y' for "could occur" or 'N' for "could not occur."

4. [2 points] The *diameter* of a tree (for this problem, a binary tree) is the maximum number of edges on any path connecting two nodes of the tree. For example, here are two sample trees and their diameters. In each case the nodes along a longest path are starred (there can be more than one longest path).

```
           A*                        A
          / \                       / \
        B*   C*                   B*   C
        / \  / \                  / \
       D  E* F*                  D*  E*
              \                  /     \
              G*                F*      G*
                                 \       \
    Diameter: 5                  H*       I*

                             Diameter: 6
```

The diameter of an empty tree or a tree with one node is 0.

Fill in the method `diameter` to compute the diameter of a binary tree. Assume that `null` represents a missing child (empty tree), and that `.left` and `.right` give the left and right children of a node, and that the `height` method has been implemented correctly.

```java
/** Return the height of T, where the height of an empty tree is -1 and
 *  that of a tree with one node is 0. */
public static int height(BinTree T) {
    if (T == null) {
        return -1;
    } else {
        return Math.max(height(T.left), height(T.right)) + 1;
    }
}


/** Return the diameter of T. */
public static int diameter(BinTree T) {

    if (_____) {
        return 0;
    } else {



        _____



        _____


    return _____

        // You need not use the two blank lines above the return.

}
```

5. [2 points] The following classes represent nodes in a 2-4 tree (aka 2-3-4 tress). In order to avoid having to handle leaf nodes specially (all of which are empty in a 2-4 tree), this representation uses a single leaf object (Node2_4.EMPTY) to represent all leaves (instead of the usual null value). Fill in the blanks in InnerNode2_4.contains to look up keys in a tree.

```
/** Represents a node in a 2-4 (aka 2-3-4) tree. There is only one
 *  instance of this base class itself, Node2_4.EMPTY, representing
 *  the empty tree. */
class Node2_4 {

    /** The unique empty node. */
    final static Node2_4 EMPTY = new Node2_4();

    /** Return my Kth child (numbering from 0). */
    Node2_4 kid(int k) { /* Implementation not shown. */ }

    /** Return the number of my children (which is one more than the
     *  the number of my keys). */
    int arity() { /* Implementation not shown. */ }

    /** Return my Kth key (numbering from 0). */
    String key(int k) { /* Implementation not shown. */ }

    /** Return true iff KEY is a key in the tree rooted at me. */
    boolean contains(String key) {
        return false;
    }
}

/** Represents non-empty nodes. Exam note: Java short-circuits
 *  conditionals, e.g. 'if (true || b)' does not evaluate b. */
class InnerNode2_4 extends Node2_4 {
    @Override
    boolean contains(String key) {
        for (_____•_____) {

            if (_____) {
                return true;

            } else if (_____) {

                return kid(_____).contains(key);
            }
        }

        return _____
    }
}
```

6. [2 points] Suppose you are given an array A with n elements such that no element is more than k slots away from its position in the sorted array. Assume that $k > 0$ and that k, while not constant, is much less than n ($k \ll n$).

    a. Fill in the blanks such that the array A is sorted after execution of this method. The important operations on a `PriorityQueue` are `add(x)`, `remove()` (remove smallest), and `isEmpty()`. Your solution should be as fast as possible.

```
public static void zorkSort(int[] A, int k) {
    int i;
    int n = A.length;
    i = 0;
    PriorityQueue<Integer> pq = new PriorityQueue<>();
    while (i < k) {


        _____

        i += 1;
    }

    while (_____) {

        A[i - k] = _____


        _____

        i += 1
    }

    while (_____) {


        _____

        i += 1;
    }
}
```

    b. What is the running time of this algorithm, as a function of n and k? Justify your answer.

7. [1 point] During an oral dissertation defense, the great mathematician David Hilbert supposedly asked the student a single question: "In all my years, I have never seen such beautiful evening clothes; pray, who is the candidate's tailor?" Who was this PhD candidate?

8. [1 point] Consider a binary search tree whose labels are each one capital letter, and assume that a level-order traversal yields the sequence "GENIUS". What is the preorder traversal of this tree? Draw the tree itself.

9. [2 points] The following questions are about tries.

a. How can you use a trie to do a range query (e.g., "What are all strings between "alpha" and "beta" in the dictionary?").

b. Henry Hacker has decided, for some reason, to replace a 2-4 tree he is using to implement a type of SortedSet of unbounded positive integers by a trie. (In the Java library, a SortedSet is an extension of Set whose iterator delivers all items in order—numeric order in this case.) He wants to be sure that his new implementation, like his 2-4 trees, can iterate through its values or a bounded subset of them in linear time. He says, "It's easy; I'll just treat my numbers as strings of digits." Why doesn't this work?

c. If searches through a certain trie for two different strings, X and Y, traverse the same particular node in the trie that is at depth k in the trie (where the root is depth 0), what can you say about X and Y?

d. Bernice Bitwiddle has successfully defined a new implementation of the Java Set<String> interface (the type HashSet<String> is an existing type that implements the same interface). Her new class uses a trie in which she stores the strings *backwards;* that is, with the top of trie corresponding to the *last,* rather than the first letter of the word. Could she easily implement an iterator() method for this representation that returns (in constant time) an iterator that then allows one to iterate through all the elements in sorted order in $\Theta(N)$ time? If not, why not, and if so, how?

10.   [2 points] Suppose we have a complete binary tree, X, that is not a heap, and we wish to heapify it. The obvious thing to do is just insert all nodes of X into a new binary heap Y. This works, but it doubles the space requirement. Suppose we'd prefer to heapify in place. Which of the following procedures will convert X into a heap (containing all of the original values)? To each, either answer "yes" or give a counter-example.
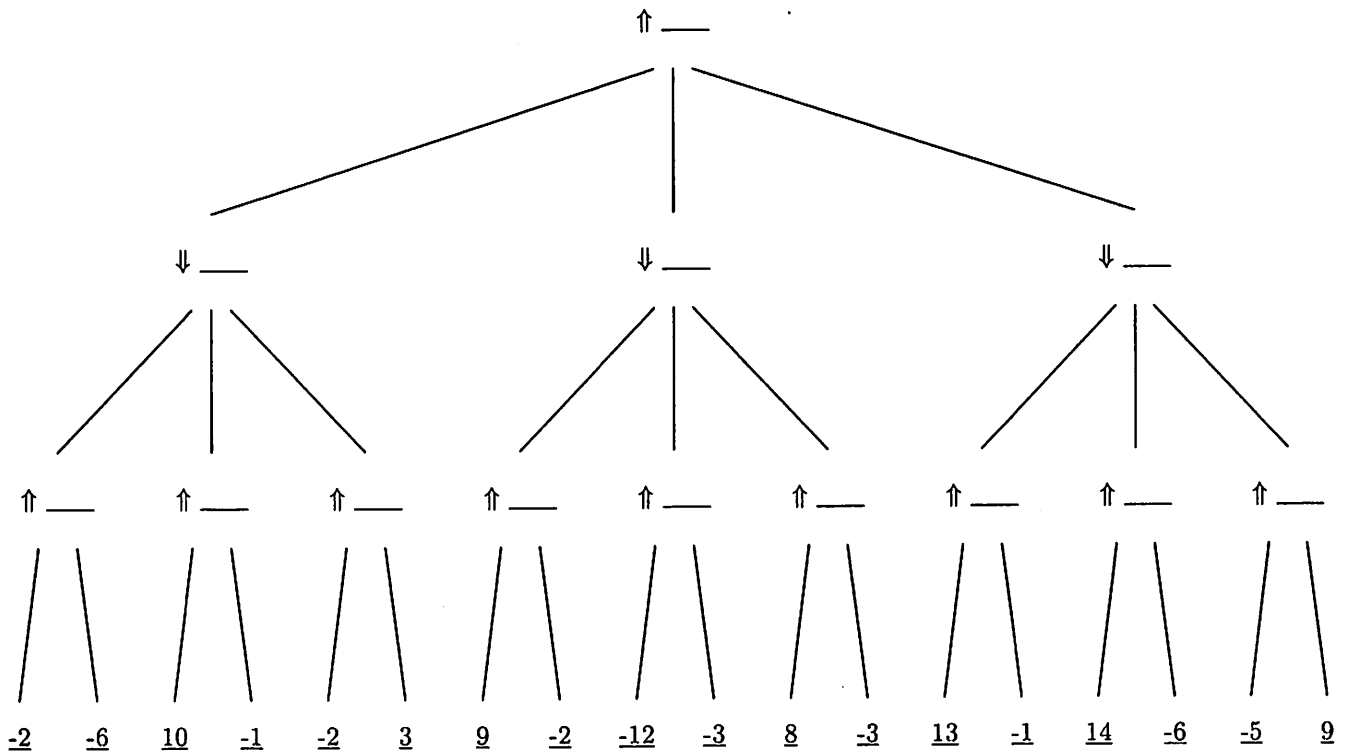
   a. Sink (heapify down) all nodes in level order (first the root, then its left child, then the right child of the root, etc).

   b. Swim (heapify up) all nodes in level order.

   c. Swim all nodes in reverse level order.

   d. Sink all nodes in reverse level order.

11.   [1 point] Consider the following game tree (the top node represents a position in which it is the "max" player's turn).

⇑ ___

⇓ ___   ⇓ ___   ⇓ ___

⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___   ⇑ ___

-2   -6   10   -1   -2   3   9   -2   -12   -3   8   -3   13   -1   14   -6   -5   9

a. Fill in the blanks in the first three rows of nodes with the values computed by the minimax algorithm.

b. Show which of these nodes can be skipped by alpha-beta pruning by crossing them out.

c. True/False: Assume that the min player plays sub-optimally at every turn, but MAX does not know this. Then the outcome of the game for the max player could be larger than predicted via minimax. Justify your answer.