

Solution to CSE 250 Final Exam

Fall 2013
Time: 3 hours.

December 13, 2013

Total points: 100	14 pages
-------------------	----------

Please use the space provided for each question, and the back of the page if you need to. Please do not use any extra paper. The space given per question is **a lot** more than sufficient to answer the question. Please be brief. Longer answers do not get more points!

- **No electronic devices of any kind. You can open your textbook and notes**
- **Please leave your UB ID card on the table**
- **This booklet must not be torn or mutilated in any way and must not be taken from the exam room**
- **Please stop writing when you are told to do so. We will not accept your submission otherwise.**
- **If you wanted to, you can answer the extra credit question without answering all of the other questions**

Your name:	_____
Your UB ID:	_____

The rest of this page is for official use only. Do not write on the page beyond this point.

Problem name and id	Score	Problem	Score	Problem	Score
(2 max)		Problem 1 (10 max)		Problem 2 (4 max)	
Problem 3 (4 max)		Problem 4 (8 max)		Problem 5 (8 max)	
Problem 6 (8 max)		Problem 7 (8 max)		Problem 8 (4 max)	
Problem 9 (4 max)		Problem 10 (8 max)		Problem 11 (8 max)	
Problem 12 (8 max)		Problem 13 (8 max)		Problem 14 (8 max)	
Total Score: (100 max)					

Problem 1 (10 points). Mark the correct choice(s) or give a brief answer. Each question is worth 1 point. All codes are in C++.

1. Of the following C++ statements, which ones are both declarations *and* definitions? (Mark all that apply.)

int foo(int b) { return b*b;} long i; char a='h';
 bool easy; typedef mytype string; struct BSTNode;

2. Consider the following fragment

```
char name[] = {'D','a','v','i','d','\0','B','l','a','i','n','e','\0' };  
cout << name;
```

What does cout print?

David Blaine David Blaine None of the above

3. Consider the following fragment

```
int a[] = {1, 2, 3, 4}; *(a+2) += 3;
```

What is the value of a[2] after the fragment is executed?

1 2 3 4 5 6

4. Consider the following fragment

```
int a[] = {1, 2, 3, 4}; int* p = a+1; p++; (*p)++;
```

What is the value of a[2] after the fragment is executed?

1 2 3 4 5 6

5. Will the following program compile without error? YES NO

```
#include <iostream>  
using namespace std;  
int foo(int i) { return i; }  
int main() {  
    int a = 2;  
    int *b = &a;  
    cout << foo(*b) << endl;  
    return 0;  
}
```

6. Suppose you wanted to make use of terminal control routines and the Lexer class I gave and all your codes are put in `yourprog.cpp`. Write down one `g++` compilation line that will compile `term_control.cpp` and produce the object file `term_control.o`

```
g++ -c term_control.cpp
```

7. Write one line of C++ that defines a new type, named `mytype`, which represents a pointer to a function that takes two `char`s and returns nothing.

```
typedef void (*mytype)(char, char);
```

8. Write *one* C++ statement that defines a variable `mymap`, where `mymap` is a map from `string` to pointer to function that has two `int` arguments and returns an `int`.

```
std::map<string, int (*)(int, int)> mymap;
```

9. To print all keys of a BST in decreasing order, we use the following traversal order

In- Post- Pre- Reverse In- Reverse Post- Reverse Pre-

10. In the B-Tree data structure we discussed in class, the last pointer of each leaf node points to its sibling node on the right (except for the rightmost leaf node, whose last pointer points to `NULL`). Explain why those pointers point that way?

So range queries can be answered quicker.

Problem 2 (4 points). Order the following functions in increasing order of asymptotic growth rate. You don't have to explain how you get the order.

$$\frac{n}{(\log n)^3}, \quad \sqrt{n}, \quad 2^{n+10}, \quad n^2(\log n)^3, \quad n^{3/4}, \quad n \cdot 2^n, \quad n\sqrt{n}.$$

Solution.

$$\sqrt{n}, \quad n^{3/4}, \quad \frac{n}{(\log n)^3}, \quad n\sqrt{n}, \quad n^2(\log n)^3, \quad 2^{n+10}, \quad n \cdot 2^n.$$

□

Problem 3 (4 points). Use the recurrence tree method to solve the recurrence $T(n) = 6T(n/2) + n$. As usual, you can assume $T(k) = O(1)$ for $k \leq 10$. Draw the tree. Show your work.

Solution. (You should draw the figure.)

$$\begin{aligned} T(n) &= 6^k T\left(\frac{n}{2^k}\right) + \sum_{i=0}^{k-1} 6^i \left(\frac{n}{2^i}\right) \\ &= 6^k T\left(\frac{n}{2^k}\right) + \sum_{i=0}^{k-1} n \cdot 3^i \\ &= 6^k T\left(\frac{n}{2^k}\right) + n \sum_{i=0}^{k-1} 3^i \\ &= 6^k T\left(\frac{n}{2^k}\right) + n \cdot \frac{3^k - 1}{2} \\ &= 6^k T\left(\frac{n}{2^k}\right) + \Theta(3^k n). \end{aligned}$$

Select k such that $n/2^k = 1$ or $k = \log_2 n$. This means

$$\begin{aligned} 3^k &= n^{\log_2 3} \\ 6^k &= n^{\log_2 6}. \end{aligned}$$

Thus,

$$T(n) = n^{\log_2 6} T(1) + \Theta(n^{\log_2 3} \cdot n) = \Theta(n^{\log_2 6}).$$

□

Problem 4 (8 points). You can assume that `using namespace std;` is at the top of the file.

1. (4) Write a C++ function `iterative_range_count()` that takes a stack `st` of `int` and an integer `a` as arguments, and returns the number of integers in `st` that are at least `a`.

```
int iterative_range_count(stack<int> st, int a)
{
    int count = 0;
    while (!st.empty()) {
        if (st.top() >= a) count++;
        st.pop();
    }
    return count;
}
```

2. (4) Write a C++ function `recursive_range_count()` that solves the same problem recursively.

```
int recursive_range_count(stack<int> st, int a)
{
    if (st.empty()) return 0;

    int tmp = st.top() >= a? 1 : 0;
    st.pop();

    return tmp + recursive_range_count(st, a);
}
```

Problems 5-7 make use of the following linked list Node structure:

```
struct Node {
    int key;
    Node* next;
    Node(int k=0, Node* n=NULL) : key(k), next(n) {};
};
```

Problem 5 (8 points). Write a function `median_element` that takes a head pointer to a NULL-terminated singly linked list whose keys are sorted in increasing order. All keys are distinct. The function returns a pointer to the node which stores the median key of all keys. If there are n keys then the median key is the $\lceil n/2 \rceil$ -smallest key. (If the list is empty then NULL is returned.) For example, if the input list is

$a.1 \rightarrow b.5 \rightarrow c.7 \rightarrow d.8 \rightarrow \text{NULL}$

then a pointer to b is returned ($n = 4$ in this case). As another sample, if the input list is

$a.1 \rightarrow b.5 \rightarrow c.7 \rightarrow d.8 \rightarrow e.10 \rightarrow \text{NULL}$

then a pointer to c is returned ($n = 5$ in this case). **Most importantly**, you can only use **one** while loop in your function, and that's the only looping structure you can use. In particular, the easy solution of looping through to find n , and traverse the second time to report the median is not valid. (**Hint:** recall the linked-list cycle detection problem I discussed in class.)

```
Node* median_element(Node* head) {
    if (head == NULL) return NULL;

    Node *slow = head, *fast = head;

    while (fast->next != NULL && fast->next->next != NULL) {
        slow = slow->next;
        fast = fast->next->next;
    }
    return slow;
}
```

Problem 6 (8 points). Write a function `fold_list()` that takes a head pointer to a NULL-terminated singly linked list consisting of Nodes. The function modifies the list in the following way: it cuts the list in half, moves the first half to the end, and returns the head pointer to the new list. **Only pointer manipulation is allowed.** If the list has an odd number of elements, then the middle element belongs to the first half. If the list is empty, NULL is returned. Feel free to call the `median_element()` function from Problem 5.

For example, if the input list is

$$a.4 \rightarrow b.1 \rightarrow c.3 \rightarrow d.6 \rightarrow e.19 \rightarrow f.14 \rightarrow \text{NULL}$$

then the output list is

$$d.6 \rightarrow e.19 \rightarrow f.14 \rightarrow a.4 \rightarrow b.1 \rightarrow c.3 \rightarrow \text{NULL},$$

and a pointer to *d* is returned. And, if the input list is

$$a.4 \rightarrow b.1 \rightarrow c.3 \rightarrow d.6 \rightarrow e.19 \rightarrow \text{NULL}$$

then the output list is

$$d.6 \rightarrow e.19 \rightarrow a.4 \rightarrow b.1 \rightarrow c.3 \rightarrow \text{NULL},$$

and a pointer to *d* is returned.

```
Node* fold_list(Node* head) {
    if (head == NULL || head->next == NULL) return head;

    Node* first_tail = median_element(head);
    Node* second_tail = head;
    while (second_tail->next != NULL) second_tail = second_tail->next;

    Node* new_head = first_tail->next;
    first_tail->next = NULL;
    second_tail->next = head;

    return new_head;
}
```

Problem 7 (8 points). Write a function `difference()` that takes *two* head pointers to `Node` as arguments. Each head pointer points to the head element of a `NULL`-terminated singly linked list which contains *distinct* keys sorted in increasing order. (The keys within each list are distinct, but keys in different lists aren't necessarily different.) The function `difference()` returns a *new sorted* linked list of `Nodes` that contain all keys of the two input linked lists that are **not** in common between the two lists. You have to keep the input linked lists intact. For example, if the two input lists are

$$a.1 \rightarrow b.4 \rightarrow c.7 \rightarrow d.16 \rightarrow \text{NULL}$$

$$e.2 \rightarrow f.4 \rightarrow g.9 \rightarrow h.16 \rightarrow i.17 \rightarrow \text{NULL}$$

Then, the output list is

$$a'.1 \rightarrow e'.2 \rightarrow c'.7 \rightarrow g'.9 \rightarrow i'.17 \rightarrow \text{NULL}$$


and a pointer to a' is returned. (I used a' , e' , c' , g' , i' because they are not the same nodes as the ones from the input lists.) Write the function so that its asymptotic running time linear in the total input list size.

```
Node* difference(Node* head1, Node* head2) {
    Node* head = NULL;
    while (head1 != NULL && head2 != NULL) {
        if (head1->key < head2->key) {
            head = new Node(head1->key, head); head1 = head1->next;
        } else if (head1->key > head2->key) {
            head = new Node(head2->key, head); head2 = head2->next;
        } else { // they are equal
            head1 = head1->next; head2 = head2->next;
        }
    }
    // at most one of the following while loops will ever execute
    while (head1 != NULL) {
        head = new Node(head1->key, head); head1 = head1->next;
    }
    while (head2 != NULL) {
        head = new Node(head2->key, head); head2 = head2->next;
    }
    return reverse_sll(head);
}

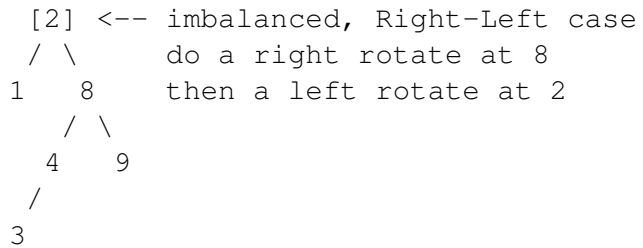

// copy from lecture 13's slides
Node* reverse_sll(Node* head) {
    Node *prev = NULL, *temp;
    while (head != NULL) {
        temp = head->next; head->next = prev;
        prev = head;      head = temp;
    }
    return prev;
}
```


Problem 8 (4 points). Draw the binary tree that has the following inorder and postorder sequences


Inorder : 4 15 6 10 8 7 5
 Postorder: 4 6 15 7 8 5 10




Problem 9 (4 points). Suppose we insert node 3 into the following AVL tree; Show where 3 is inserted at first, then circle the node that becomes imbalanced and explain how the re-balance step works to keep the AVL property. Draw the resulting AVL tree.




after right rotate at 8




after left rotate at 2




Problem 10 (8 points). (a) Suppose we splay node 6 in the following tree, what does the resulting tree look like? Draw all the intermediate trees after each zig-zig, zig-zag, or zig baby step and label the step with the name zig-zig, zig-zag, or zig.



(b) Consider the following Binary Search Tree. Is it an AVL Tree? Why or why not? Put a color red (R) or black (B) next to each node so that it is a red-black tree. (The x are NULL nodes.)



it is an AVL tree, because at every node the height of the left differs from the height of the right subtree by at most 1



Problems 10 to 14 make use of the following binary search tree node structure:

```
struct BSTNode {
    int key;
    BSTNode* left;
    BSTNode* right;
    BSTNode* parent;
};
```

Problem 11 (8 points). Write a function that takes a pointer to the root of a binary search tree with the above BSTNode structure, and that removes the node with the *minimum* key from the tree. Note that the root pointer is passed by reference. If the minimum node is the root node then the root pointer has to be modified properly. All your code has to be contained within the function, do not write any auxiliary function.

```
// this is virtually the same as a question in the sample exam
void delete_min(BSTNode*& root)
{
    if (root == NULL)
        return;
    BSTNode* min = root;
    while (min->left != NULL)
        min = min->left;

    BSTNode* p = min->parent; // NULL if min == root
    if (p == NULL) {
        root = root->right; // new root
        root->parent = NULL;
    } else {
        p->left = min->right;
        if (min->right != NULL)
            min->right->parent = p;
    }
    delete min; // don't forget to release memory
}
```

Problem 12 (8 points). Write a function `common_ancestors()` that takes two (not necessarily distinct) pointers to `BSTNodes` in a `BST` and returns a vector of pointers to nodes in the tree which are the common ancestors of the two given nodes. (Note that by definition a node is an ancestor of itself.)

```
vector<BSTNode*> common_ancestors(BSTNode* node1, BSTNode* node2)
{
    stack<BSTNode*> s1, s2;
    while (node1 != NULL) {
        s1.push(node1);
        node1 = node1->parent;
    }
    while (node2 != NULL) {
        s2.push(node2);
        node2 = node2->parent;
    }
    vector<BSTNode*> ret;
    while (!s1.empty() && !s2.empty()) {
        if (s1.top() == s2.top()) {
            ret.push(s1.top());
            s1.pop(); s2.pop();
        } else {
            break;
        }
    }
    return ret;
}
```

Problem 13 (8 points). (a) Write a function `size()` that takes the root pointer of a binary tree with the `BSTNode` structure and returns the size of the tree, which is the number of non-NULL nodes in the tree.

```
size_t size(BSTNode* root) {
    if (root == NULL) return 0;
    return 1 + size(root->left) + size(root->right);
}
```

(b) Write a function `rank()` that takes a pointer `node` to a node (not necessarily the root) of a binary search tree with the `BSTNode` structure and returns the *rank* of that node. The rank of a node is k if the node's key is the k th smallest key among all keys. You can assume that all keys in the tree are distinct. You should make use of the `size()` function from part (a). If `node` is `NULL` then 0 is returned.

```
size_t rank(BSTNode* node) {
    if (node == NULL) return 0;
    size_t r = 1 + size(node->left);
    while (node->parent != NULL) {
        if (node->parent->right == node) {
            r += 1 + size(node->parent->left);
        }
        node = node->parent;
    }
    return r;
}
```

Problem 14 (8 points). Write a function `last_node_depth_k()` that takes the root pointer of a binary tree with the `BSTNode` structure, a non-negative integer k and returns a pointer to the last node at depth k of the tree. (Last as we go from left to right.) The root node is at depth 0, the root's children are at depth 1, and so on. If k is greater than the largest depth, then `NULL` is returned.

For example, consider the tree drawn in Problem 10b. The last node at depth 1 is node 50, at depth 2 is node 60, at depth 3 is 15, at any depth > 3 is `NULL`.

```
// this is a slight variation of a question from the sample exam
BSTNode* last_node_depth_k(BSTNode* root, size_t k)
{
    if (root == NULL) return NULL; // boundary case

    size_t current_depth = 0;
    deque<BSTNode*> q;
    q.push_front(root);
    size_t c=1; // c = # of nodes at current depth still in q
    size_t n=0; // n = # of nodes at next depth already in q
    while (!q.empty()) {
        BSTNode* cur = q.back(); // head of q
        q.pop_back();
        c--;
        if (cur->left != NULL) { q.push_back(cur->left); n++; }
        if (cur->right != NULL) { q.push_back(cur->right); n++; }
        if (c == 0) {
            if (current_depth == k)
                return cur;
            c = n; n = 0;
            current_depth++;
        }
    }
    return NULL;
}
```