

CS F211

Data Structures and Algorithms

Practice Problems

Hashmaps, AVL Trees, and Red-Black Trees

Problem 1

Use hashmaps and prefix sums to write a program that, given an array of integers and a target sum, finds the number of subarrays of the array such that its elements sum up to the target.

```
1 #include <iostream>
2 #include <vector>
3 #include <list>
4 #include <string>
5
6 using namespace std;
7
8 class HashMap {
9 private:
10     static const int TABLE_SIZE = 10;
11     vector< list< pair< int, int > > > table;
12
13     int hashFunction(const int &key) {
14         hash<int> hashFunc;
15         return hashFunc(key) % TABLE_SIZE;
16     }
17
18 public:
19     HashMap() {
20         table.resize(TABLE_SIZE);
21     }
22
23     void set(const int &key, int value) {
24         int index = hashFunction(key);
25         for (auto &p : table[index]) {
26             if (p.first == key) {
27                 p.second = value;
28                 return;
29             }
30         }
31         table[index].emplace_back(key, value);
32     }
33
34     int get(const int &key) {
35         int index = hashFunction(key);
36         for (auto &p : table[index]) {
37             if (p.first == key)
38                 return p.second;
```

```
39     }
40     return 0;
41 }
42 }
43
44 int main() {
45     /* ??? */
46 }
```

```
1 >> 8 7
2 >> 3 4 7 2 -3 1 4 2
3 4
```

Next task: Modify the code to output one such subarray with the least length.

Problem 2

Write code that behaves as an interpreter for a simple programming language. In this simple language, each line is an assignment consisting of a variable on the LHS and a postfix expression on the RHS which contains both integers and previously defined variables. Use stacks for evaluating expressions, and a hashmap (which we call symbol table) for storing the values assigned to variables.

```
1 #include <iostream>
2 #include <vector>
3 #include <list>
4 #include <string>
5 #include <stack>
6
7 using namespace std;
8
9 class HashMap {
10 private:
11     static const int TABLE_SIZE = 10;
12     vector< list< pair< string, int > > > table;
13
14     int hashFunction(const string &key) {
15         hash<string> hashFunc;
16         return hashFunc(key) % TABLE_SIZE;
17     }
18
19 public:
20     HashMap() {
21         table.resize(TABLE_SIZE);
22     }
23
24     void set(const string &key, int value) {
25         int index = hashFunction(key);
26         for (auto &p : table[index]) {
27             if (p.first == key) {
28                 p.second = value;
29                 return;
30             }
31         }
32         table[index].emplace_back(key, value);
33     }
34 }
```

```

35     int get(const string &key) {
36         int index = hashFunction(key);
37         for (auto &p : table[index]) {
38             if (p.first == key)
39                 return p.second;
40         }
41     return 0;
42 }
43 };
44
45 int evaluate(string expr, HashMap* m) {
46     /* ??? */
47 }
48
49 void assign(string token, string expr, HashMap* m) {
50     int val = evaluate(expr, m);
51     m->set(token, val);
52     cout << token << " = " << val << '\n';
53 }
54
55 int main() {
56     HashMap m;;
57     assign("a", "5 3 +", &m);
58     assign("b", "a 5 * 8 -", &m);
59     assign("result", "b a /", &m);
60 }
```

```

1 a = 8
2 b = 32
3 result = 4
```

Next task: Use the BT expression evaluating code you have written in a previous problem sheet to allow the program to instead evaluate regular arithmetic expressions.

Problem 3

You are given a list of n integers and a target sum. Write a program to find if there exists a subset of the integers such that they sum up to the target. The naive way is to explore all 2^n possibilities, but a more efficient way involves first dividing the array into two halves, finding the two sets of $2^{\frac{n}{2}}$ possible sums for either half, and then checking if there exists a way to add a sum from the first half and a sum from the second half to obtain the target sum.

The recursive function that generates all possible sums for a half stores the sums in a set implemented using an AVL tree. This automatically handles duplication and ensures that all possible sums can be retrieved in the end in sorted order, which is important for the merge step. Use two pointers on the two sorted arrays of possible sums to efficiently check if the target sum is obtainable from two sums belonging to the two different halves.

```
1 #include <iostream>
2 #include <vector>
3
4 using namespace std;
5
6 struct Node {
7     int key, height;
8     Node *left, *right;
9
10    Node(int k) : key(k), height(1), left(nullptr), right(
11        nullptr) {}
12};
13
14 class AVLTree {
15 private:
16     Node* root;
17
18     int getHeight(Node* node) {
19         return node ? node->height : 0;
20     }
21
22     int getBalance(Node* node) {
23         return node ? getHeight(node->left) - getHeight(node-
24             >right) : 0;
25     }
26
27     Node* rotateRight(Node* y) {
```



```

62         node->right = rotateRight(node->right);
63         return rotateLeft(node);
64     }
65
66     return node;
67 }
68
69 void inorder(Node* node, vector<int>& v) {
70     if (!node) return;
71     inorder(node->left, v);
72     v.push_back(node->key);
73     inorder(node->right, v);
74 }
75
76 public:
77     AVLTree() : root(nullptr) {}
78
79     void insert(int key) {
80         root = insert(root, key);
81     }
82
83     void getElems(vector<int>& v) {
84         inorder(root, v);
85     }
86 };
87
88 void genSumsRec(int idx, int sum, vector<int>& v, AVLTree* s)
89 {
90     /* ??? */
91 }
92
93 void genSums(vector<int>& v, AVLTree* s) {
94     /* ??? */
95 }
96
97 int main() {
98     int n, sum;
99     cin >> n >> sum;
100    vector<int> v1, v2;
101    for (int i = 0; i < n; i++) {
102        int a;
103        cin >> a;
104        if ((i % 2) == 0) v1.push_back(a);
105        else v2.push_back(a);
106    }

```

```
106     AVLTree* s1 = new AVLTree(), *s2 = new AVLTree();
107     genSums(v1, s1);
108     genSums(v2, s2);
109     vector<int> b1, b2;
110     s1->getElems(b1), s2->getElems(b2);
111     /* ??? */
112     return 0;
113 }
```

```
1 >> 7 16
2 >> 1 3 4 2 5 6 8
3 Yes
```

Next task: As is clear in the given code, there is a method that converts the elements of the AVLT into a vector, and this vector is then later used in the merge step. We could simplify this by instead implementing direct methods in the AVLT to retrieve next and previous elements starting from a particular state. Implement these methods. Of course, the methods only work if each node has a pointer to its parent, so add a parent attribute and modify the other methods to handle changes to parent pointers.

Problem 4

Modify the AVLT code given in the previous problem to let the tree find the median of its elements in $O(\log n)$. Formally, your code must handle two types of queries, which are described below.

1. 1 x – add integer x to the set.
2. 2 – return the median of all the integers in the tree.

To achieve this, as a hint, consider adding a size attribute to node that keeps track of size of the subtree rooted at that node.

```
1 >> 5
2 >> 1 6
3 >> 1 2
4 >> 2
5 2
6 >> 1 9
7 >> 1 4
8 >> 1 5
9 >> 2
10 5
```

Next task: Modify the code to answer general k th-smallest integer queries.

Problem 5

Red-black trees are special binary search trees which are capable of balancing themselves, like AVL trees. In this problem, you have to examine the efficiency of RBTs using a metric called *balance*, which we define as the sum of differences of the heights of left and right subtrees of all nodes. Add a recursive method to the RBT that calculates this metric.

```
1 #include <iostream>
2 #include <string>
3 #include <cstdlib>
4 #include <ctime>
5
6 using namespace std;
7
8 enum Color { RED, BLACK };
9
10 struct Node {
11     int data;
12     Color color;
13     Node *left, *right, *parent;
14     int height, balance;
15
16     Node(int data) {
17         this->data = data;
18         left = right = parent = nullptr;
19         color = RED;
20     }
21 };
22
23 class RedBlackTree {
24 private:
25     Node* root;
26
27     void rotateLeft(Node* &root, Node* &pt) {
28         Node* rightChild = pt->right;
29         pt->right = rightChild->left;
30
31         if (pt->right)
32             pt->right->parent = pt;
33
34         rightChild->parent = pt->parent;
35     }
```

```

36     if (!pt->parent)
37         root = rightChild;
38     else if (pt == pt->parent->left)
39         pt->parent->left = rightChild;
40     else
41         pt->parent->right = rightChild;
42
43     rightChild->left = pt;
44     pt->parent = rightChild;
45 }
46
47 void rotateRight(Node* &root, Node* &pt) {
48     Node* leftChild = pt->left;
49     pt->left = leftChild->right;
50
51     if (pt->left)
52         pt->left->parent = pt;
53
54     leftChild->parent = pt->parent;
55
56     if (!pt->parent)
57         root = leftChild;
58     else if (pt == pt->parent->left)
59         pt->parent->left = leftChild;
60     else
61         pt->parent->right = leftChild;
62
63     leftChild->right = pt;
64     pt->parent = leftChild;
65 }
66
67 void fixViolation(Node* &root, Node* &pt) {
68     Node* parent = nullptr;
69     Node* grandparent = nullptr;
70
71     while ((pt != root) && (pt->color != BLACK) && (pt->
72         parent->color == RED)) {
73         parent = pt->parent;
74         grandparent = parent->parent;
75
76         if (parent == grandparent->left) {
77             Node* uncle = grandparent->right;
78
79             if (uncle && uncle->color == RED) {

```

```

80         parent->color = BLACK;
81         uncle->color = BLACK;
82         pt = grandparent;
83     } else {
84         if (pt == parent->right) {
85             rotateLeft(root, parent);
86             pt = parent;
87             parent = pt->parent;
88         }
89
90         rotateRight(root, grandparent);
91         swap(parent->color, grandparent->color);
92         pt = parent;
93     }
94 } else {
95     Node* uncle = grandparent->left;
96
97     if (uncle && uncle->color == RED) {
98         grandparent->color = RED;
99         parent->color = BLACK;
100        uncle->color = BLACK;
101        pt = grandparent;
102    } else {
103        if (pt == parent->left) {
104            rotateRight(root, parent);
105            pt = parent;
106            parent = pt->parent;
107        }
108
109        rotateLeft(root, grandparent);
110        swap(parent->color, grandparent->color);
111        pt = parent;
112    }
113 }
114
115 root->color = BLACK;
116 }
117
118
119 public:
120     RedBlackTree() {
121         root = nullptr;
122     }
123
124     void insert(const int &data) {

```

```

125     Node* newNode = new Node(data);
126
127     Node* parent = nullptr;
128     Node* current = root;
129
130     while (current) {
131         parent = current;
132         if (newNode->data < current->data)
133             current = current->left;
134         else
135             current = current->right;
136     }
137
138     newNode->parent = parent;
139
140     if (!parent)
141         root = newNode;
142     else if (newNode->data < parent->data)
143         parent->left = newNode;
144     else
145         parent->right = newNode;
146
147     fixViolation(root, newNode);
148 }
149
150 void inorderRec(Node* root) {
151     /* ??? */
152 }
153
154 void inorder() {
155     inorderRec(root);
156     cout << endl;
157 }
158
159 void preorderRec(Node* root) {
160     /* ??? */
161 }
162
163 void preorder() {
164     preorderRec(root);
165     cout << endl;
166 }
167
168 void getBalanceRec(Node* root) {
169     /* ??? */

```

```

170     }
171
172     int getBalance() {
173         /* ??? */
174     }
175 };
176
177 class Random {
178 private:
179     int a, c, m, x;
180
181 public:
182     Random(int _m) {
183         a = 1664525;
184         c = 1013904223;
185         m = _m;
186         x = static_cast<int>(time(0));
187     }
188
189     int generate() {
190         x = (int)((long long)a * (long long)x + c) % m;
191         return x + 1;
192     }
193 };
194
195 int main() {
196     RedBlackTree rbt;
197     rbt.insert(10);
198     rbt.insert(3);
199     rbt.insert(7);
200     rbt.insert(2);
201     rbt.insert(8);
202     rbt.insert(1);
203     rbt.inorder();
204     rbt.preorder();
205     cout << rbt.getBalance() << '\n';
206     return 0;
207 }
```

```

1 1 2 3 7 8 10
2 7 2 1 3 10 8
3 1
```

Next task: Use the random generation class given in the code to generate a

large number of random values (like 1000 or 10000) to insert into the RBT. Insert these values into the tree and calculate the balance at regular intervals of insertion. Examine how the balance of the tree changes with size. Try the same experiment on other kinds of binary tree.