

ECE 2400 Computer Systems Programming, Fall 2017

Prelim 2 Prep

revision: 2017-11-04-22-42

These problems are not meant to be exactly like the problems that will be on the prelim. These problems are instead meant to represent the kind of understanding you should be able to demonstrate on the prelim.

In the following problems, we will explore two different data structures to track the directory hierarchy in an operating system. The data hierarchy could be used to store metadata about each directory in the system (e.g., number of files in each directory, access control for the directory), and the operating system would provide an *application level interface* (API) to enable system-level software to read and write the directory hierarchy. For example, the following commands will make four directories and then display the directory hierarchy.

```
1 % mkdir foo
2 % mkdir foo/bar0
3 % mkdir foo/bar1
4 % mkdir foo/bar0/baz
5 % tree
```

The directory hierarchy is made up of directories where each *parent directory* can have one or more *child directories*. So in the above example, `foo` is the parent of `bar0` and `bar1`. Or equivalently, `bar0` and `bar1` are the children of `foo`.

We will explore two data structures to store the directory hierarchy. The data structures should support making new directories and printing the full directory hierarchy. For this problem, you can assume that the maximum number of directories in the system is 16 and that there are no more than four child directories in any given parent directory. There are two common errors that we might want our data structure to check:

- **Directory Already Exists:** The data structure should cause an error if the user attempts to make a directory that already exists. *For this problem you are required to correctly detect this error and throw an `std::invalid_argument` exception.*
- **Parent of Directory Does Not Exist:** Ideally, the data structure should cause an error if the user attempts to create a new directory where one or more of the parents of that directory do not exist. *To simplify the problem, you can assume the user never attempts to create a new directory unless all of the parents of that directory already exist.*

Problem 1. Array-Based Directory Hierarchy Data Structure

Our first data structure simply uses an array to store all of the directories in the system. The class declaration is as follows:

```
1  class DirList
2  {
3  public:
4
5      DirList();
6      void mkdir( const std::string& path );
7      void print() const;
8
9  private:
10     size_t      m_paths_sz;
11     std::string m_paths[16];
12 };
```

The `m_paths_sz` variable is used to store how many directories are presently in the system.

Part 1.A Implementing DirList

Implement the constructor, `mkdir`, and `print` member functions in the `DirList` class. Your implementation of `mkdir` should throw an `std::invalid_argument` exception if the user attempts to make a directory that already exists. You can assume you have the following private helper member function

```
1  DirList::DirList()
2      : m_paths_sz(0)
3  { }
4
5  void DirList::mkdir( const std::string& path )
6  {
7      for ( size_t i = 0; i < m_paths_sz; i++ ) {
8          if ( m_paths[i] == path )
9              throw std::invalid_argument("dir already exists!");
10     }
11
12     m_paths[m_paths_sz] = path;
13     m_paths_sz++;
14 }
```

```
1 void DirList::print() const
2 {
3     for ( size_t i = 0; i < m_paths_sz; i++ )
4         std::cout << m_paths[i] << std::endl;
5 }
```

Part 1.B Algorithm Analysis for DirList

For this analysis assume we wish to store N paths in the directory hierarchy. **What is the best- and worst-case time complexity for `mkdir` as a function of N ? What is the space complexity of this data structure as a function of N ?**

Each call to `mkdir` will need to do a linear search through all of the paths in the data structure to verify that the path being added does not exist. This search will be $O(N)$. The data structure will have N entries in the `m_paths` array, so the space complexity is $O(N)$. Note that this analysis ignores the actual size in characters of each directory name which may or may not be a reasonable assumption.

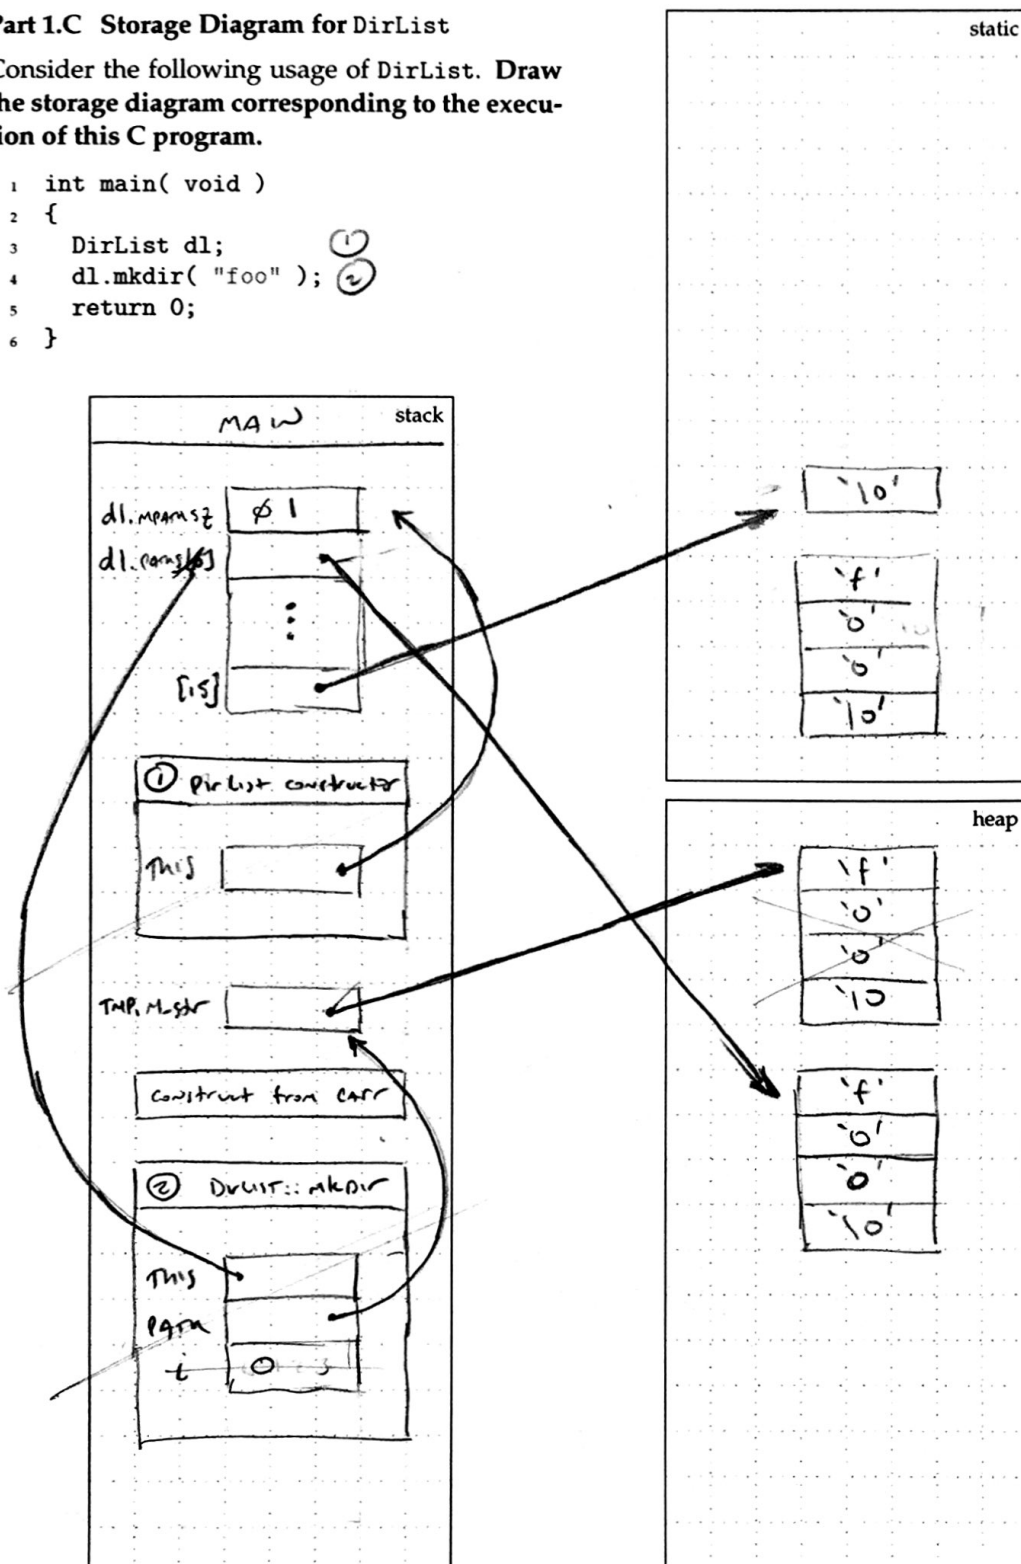
Part 1.C Storage Diagram for DirList

Consider the following usage of DirList. Draw the storage diagram corresponding to the execution of this C program.

```

1 int main( void )
2 {
3   DirList dl;
4   dl.mkdir( "foo" );
5   return 0;
6 }

```



Problem 2. Tree-Based Director Hierarchy Data Storage

Our second data structure uses a *tree*. A tree is like a linked list, except every node can have more than one “next node”. The class declaration is as follows:

```

1  class DirTree
2  {
3  public:
4
5      DirTree() : m_root(Node()) { }
6
7      void mkdir( std::string path );
8
9      void print()
10     {
11         for ( size_t i = 0; i < m_root.children_size; i++ )
12             print_h( "", m_root.children[i] );
13     }
14
15 private:
16
17     class Node {
18     public:
19
20         Node( const std::string dir_ = "" )
21             : dir(dir_), children_size(0)
22         { }
23
24         std::string dir;
25         size_t      children_size;
26         Node*       children[4];
27     };
28
29     size_t split( std::string dirs[], std::string path )
30     {
31         size_t i = 0;
32         auto pos = path.find("/");
33         while ( pos != std::string::npos ) {
34             dirs[i] = path.substr(0,pos);
35             i++;
36             path.erase( 0, pos+1 );
37             pos = path.find("/");
38         }
39         dirs[i] = path.substr(0,pos);
40         i++;
41         return i;
42     }
43
44     void print_h( const std::string& prefix, Node* node_p );
45
46     Node m_root;
47
48 };

```

Part 2.A Implementing DirTree::mkdir

Implement the mkdir member function. This function should iterate across the directories in the path. For each directory, the function should check to see if that directory already exists and if not, the function should dynamically allocate a new Node to represent that directory. The end result should be a tree which represents the directory hierarchy. We have provided a private helper function to split a path into an array of strings, with one string per directory in the path.

```

1  void DirTree::mkdir( std::string path )
2  {
3      std::string dirs[8];
4      size_t dirs_size = split( dirs, path );

5      // Search through parents
6      Node* node_p = &m_root;
7      for ( size_t i = 0; i < dirs_size-1; i++ ) {
8          for ( size_t j = 0; j < node_p->children_size; j++ ) {
9              if ( node_p->children[j]->dir == dirs[i] )
10                 node_p = node_p->children[j];
11            }
12        }

13        // Make sure directory does not exist
14        for ( size_t j = 0; j < node_p->children_size; j++ ) {
15            if ( node_p->children[j]->dir == dirs[dirs_size-1] )
16                throw std::invalid_argument("dir already exists!");
17        }

18        // Create new node
19
20        Node* new_node_p = new Node(dirs[dirs_size-1]);
21        node_p->children[node_p->children_size] = new_node_p;
22        node_p->children_size += 1;
23    }

```

Part 2.B Implementing DirTree::print_h

Implement the `print_h` private helper member function. This should be a recursive function which prints out the current path corresponding to the given `Node` pointer and then recursively processes the current node's children.

```

1 void DirTree::print_h( const std::string& prefix, Node* node_p )
2 {
3     std::string new_prefix = prefix + "/" + node_p->name;
4
5     std::cout << new_prefix << std::endl;
6     for ( size_t i = 0; i < node_p->children_size; i++ ) {
7         print_h( new_prefix, node_p->children[i] );
8     }
9 }
```

Part 2.C Algorithm Analysis of DirTree

For this analysis assume we wish to store N paths in the directory hierarchy. **What is the best-case and worst-case time complexity for `mkdir` as a function of N ? What is the space complexity of this data structure as a function of N ?**

Each call to `mkdir` will need to search all the way through the tree from the root to one of the leaves. We will need to do one comparison per node; we check to make sure the parents exist as we go through the tree, and when we get to the end of the tree we check to make sure that the new directory does not exist. In the worst case, every directory has exactly one child, which means our directory “tree” is really a chain of nodes, and thus this search requires $O(N)$ comparisons. In the best case, every directory has exactly four children, and the search requires only $O(\log_4 N)$ comparisons.

Every new path adds one node to the tree, so our tree has $O(N)$ nodes and thus the space complexity is also $O(N)$. Note that this analysis ignores the actual size in characters of each directory name which may or may not be a reasonable assumption. The tree-based approach does exploit some redundancy in the way it stores paths, so this might warrant deeper analysis.