

EECS 281, Week 8: Wednesday

Hash Tables

1. Describe hashing as a two-step process.

2. Consider the hash function for a `string` below:

```
.....  
int hash(const string &s) {  
    return toupper(s[0])) - 'A';  
}  
.....
```

Critique this hash function, identifying and explaining its downsides.

3. What is a collision?

4. When can collisions occur?

5. How are collisions different from clusters?

6. What are the properties of a good hash function?

7. What is a good way to generate hash codes for a string?

8. What is load factor α ?

EECS 281, Week 8: Wednesday

9. What is dynamic hashing?
10. What is simple uniform hashing?
11. Suppose we use a hash function h to hash n distinct keys into an array of length m . Assuming simple uniform hashing, what is the expected number of collisions?
12. Demonstrate what happens when we insert the keys 5, 28, 19, 15, 20, 33, 12, 17, 10 into a hash table with collisions resolved by chaining. Let the table have 9 slots, and let the hash function be $h(k) = k \bmod 9$.
13. Suppose we have a hash table that stores integer keys and uses chaining as its collision resolution technique. Assume that the hash code for an integer is that integer itself.
- | | | |
|---|----------------------|------|
| 0 | <input type="text"/> | → 8 |
| 1 | <input type="text"/> | → 25 |
| 2 | <input type="text"/> | → 10 |
| 3 | <input type="text"/> | → 15 |
- Demonstrate what happens when we insert the keys 18 and 5. Be sure to resize the hash table by doubling the array when the load factor reaches 1.5.
14. When you modify a key that has been inserted into a hash table, will you be able to retrieve that entry again? Explain.
- A. Always
 - B. Sometimes
 - C. Never
15. When you modify a value that has been inserted into a hash table, will you be able to retrieve that entry again? Explain.
- A. Always
 - B. Sometimes
 - C. Never

EECS 281, Week 8: Wednesday

16. (T/F) While a good hash table that contains N elements and uses separate chaining has average-case constant time for look-up, the worst case running time for a lookup is $\Theta(\log N)$.
17. (T/F) `std::vector` can be used as a key in `std::unordered_map`.
18. What do we need to do in order to use `std::unordered_map` with custom types?
19. (T/F) A hash table can be used to store a dictionary in a spell-checker program.
20. Consider this definition of a hash table that stores English words and uses separate chaining as its collision resolution technique:

```
struct HashTableNode {
    string word;
    HashTableNode* next = nullptr;
};
```

```
class HashTable {
public:
    size_t size() const;
private:
    vector<HashTableNode*> slots;
}
```

Suppose that `HashTable` has been implemented without an instance variable to keep track of the number elements. Complete the implementation of `size` method below in such a way that it returns the number of words stored in the hash table. Feel free to add other helper (private) functions.

```
size_t HashTable::size() const {
```

EECS 281, Week 8: Wednesday

21. Consider the definition of a TicTacToeGame class.

```
class TicTacToeGame {
public:
    // always 3-by-3
    const size_t size = 3;
    // hash code to use for std::hash()
    int hashCode() const;
private:
    // either 'X', 'O' or ' '
    char tiles[size][size];
    // keep track of turns
    bool isXTurn = true;
};
```

Implement the `hashCode` methods that computes the hash code to specialize `std::hash` for the Tic-Tac-Toe game. Ensure that different board configurations have different hash codes.

```
int TicTacToeGame::hashCode() const {
```

22. How many different configurations of the Tic-Tac-Toe board are there?

EECS 281, Week 8: Wednesday

Tries

23. (T/F) A trie is a binary tree.
24. To use objects of type T as keys in a trie, what property must the type T have?
25. Suppose we have a trie that stores n English words, and each of trie's nodes contains an array of p pointers to other nodes. What is the time complexity of looking up a word of size k in this trie?
26. (T/F) Tries can be used to store valid words (i.e., if they are in the dictionary), but cannot be used to store the words' definition.
27. (T/F) Unlike hash tables, we cannot implement tries in a way to store duplicate keys.
28. (T/F) Let T be a trie that stores N strings. We can traverse T in sorted order in $\Theta(N)$ time.
29. (T/F) Let T be a trie that stores strings backwards, with the top of trie corresponding to the last, rather than the first character of the string. We have traverse T in sorted order in $\Theta(N)$ time.
30. (T/F) We can use a trie to do a range query, e.g., to get all strings between "maxim" and "pizza" in the dictionary in $\Theta(N)$ time.
31. If we traverse the same node while searching for two strings, S_1 and S_2 , in a trie and that node is at depth k in the trie (where the root is depth 0), what can we say about S_1 and S_2 ?
32. Describe one disadvantage of using a trie to store a dictionary of English words?

EECS 281, Week 8: Wednesday

33. Suppose we have a trie that stores n English words, and each of trie's nodes contains an array of p pointers to other nodes. Describe two ways to improve memory usage of this trie.

34. Consider this definition of a case-insensitive trie for English words:

```
struct TrieNode {  
    bool isWord;  
    TrieNode* children[26];  
};  
  
class Trie {  
public:  
    size_t size() const;  
private:  
    TrieNode* root;  
}
```

Suppose that `Trie` has been implemented without an instance variable to keep track of the number of words in trie. Complete the implementation of `size` method below in such a way that it returns the number of words stored in the trie. Feel free to add other helper (private) functions.

```
size_t Trie::size() const {
```

35. When should we use a hash table over a trie? Trie over a hash table?

EECS 281, Week 8: Wednesday

Trees

36. What is the depth of a node?

37. Implement `depth`, a function that takes a `BinaryTreeNode*` and returns the depth of that node. Assume that `BinaryTreeNode` struct has a `parent` pointer.

```
int depth(BinaryTreeNode* node) {
```

38. What is the height of a node?

39. Implement `height`, a function that takes a `BinaryTreeNode*` and returns the height of the tree rooted at that node. Assume that `BinaryTreeNode` struct has a `left` pointer and a `right` pointer, and that the height of an empty tree is -1 and that the height of a tree with one node is 0.



```
int height(BinaryTreeNode* node) {
```

40. Complete the table below with amortized time complexities and the data structure generally used by the STL for these containers.

Operation	set	map	unordered_set	unordered_map
insert				
find				
remove				
Data structure				

EECS 281, Week 8: Wednesday

41. The *diameter* of a 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 longest path is shown in blue with stars next to all nodes in the path. Note that there can be more than one longest path.




Implement the function `diameter` that computes the diameter of a binary tree represented by a pointer to an object of `BinaryTreeNode` class.

```
int diameter(BinaryTreeNode* node) {
```

42. (T/F) The height of a binary tree with N elements is $\log N$.

43. What are pre-order, in-order, post-order and level-order traversals of this tree?



44. (T/F) All four traversals can be performed on all trees.

45. (T/F) For any non-empty binary tree with A leaves and B nodes of degree 2, $A = B + 1$.

46. (T/F) For any non-empty binary tree with A nodes and B nodes of degree 2, $A = B + 1$.

EECS 281, Week 8: Wednesday

47. (T/F) In-order traversal of a binary search tree will visit the nodes in sorted order.
48. A binary tree has a post-order traversal D A B E C and an in-order traversal D E B A C.
Draw the binary tree and give its pre-order traversal.
49. (T/F) BSTs are balanced trees.
50. (T/F) Searching for an element in a binary search tree with N nodes where each node has either 0, 1 or 2 children takes $\Theta(N)$ time in the worst case.
51. (T/F) Searching for an element in a binary search tree with N nodes where each node has either 0 or 2 children takes $\Theta(\log N)$ time in the worst case.
52. (T/F) In a binary search tree, all internal nodes, except the ones at the last level, have two children.
53. What is the purpose of AVL trees?
54. Insert the following values into an empty AVL tree: 9, 21, 64, 12, 1, 19, 32. Then remove 1 and 21.
55. Insert the following values into an empty AVL tree: 21, 32, 64, 72, 17. Then remove 21, 32, 64, 72 and 17.

EECS 281, Week 8: Wednesday


Graphs and Graph Algorithms

56. Describe the differences between trees and graphs.

57. Breadth-first search traversal of a graph is equivalent to which traversal of a tree?

58. Describe an algorithm for finding connected components in a graph.

Consider the following graph for the next three questions:



59. Is the graph weighted? Directed? Cyclic?

60. Give the adjacency matrix representation for the graph shown above.

61. Give the adjacency list representation for the graph shown above.


62. Give the (pre-order) DFS traversal for the graph shown above, starting at node 0.

63. Give the post-order DFS traversal for the graph shown above.

64. Give the BFS traversal for the graph shown above.

EECS 281, Week 8: Wednesday

Consider the following undirected graph for the next five questions:



65. Give the sequence of vertices added to the minimum spanning tree for the above graph when running Prim's algorithm. Start with vertex C.
66. Give the sequence of vertices added to the maximum spanning tree for the above graph when running Prim's algorithm. Start with vertex C.
67. Give the sequence of edges added to the minimum spanning tree for the above graph when running Kruskal's algorithm.
68. Give the sequence of edges added to the maximum spanning tree for the above graph when running Kruskal's algorithm.
69. Give the sequence of edges that would be "explored" when running Dijkstra's algorithm to find the shortest path from node D to node H.
70. Prim's, Kruskal's and Dijkstra's are considered _____ algorithms because they solve the problems by progressively making the locally optimal choice at each stage or iteration.
71. (T/F) Dijkstra's algorithm always finds the shortest path on any graph.
72. An airline is attempting to restructure its services so that that every city in its network can connect to any other city while minimizing the total length of its routes. What problem is the airline trying to solve?
73. An airline's network is set up so that most cities have flights to only a few destinations, and a few cities are hubs, with connections to many destinations. Given this setup, which algorithm should the airline use to create its new network?

EECS 281, Week 8: Wednesday

74. An airline wants to generate a minimum spanning tree of its network and represents it with the following distance matrix:

Could this be a valid solution to the airline's MST? Explain.

75. Suppose that a directed weighted graph is represented with an adjacency matrix. Implement `reverse`, a function that would replace all edges (v, w) with (w, v) .

```
void reverse(vector<vector<int>> &graph) {
```

Disjoint Sets

76. What operations can disjoint sets perform efficiently?

77. How can we represent disjoint set with an array?

78. What optimizations can we perform on disjoint sets?

79. (T/F) Disjoint sets can be used with Prim's algorithm.

80. (T/F) Disjoint sets can be used with Kruskal's algorithm.

EECS 281, Week 8: Wednesday

Algorithm Design

81. What do C, L, R and S stand for in CLRS?
82. (T/F) Any divide-and-conquer algorithm is an example of a dynamic programming algorithm.
83. (T/F) Merge sort and Quicksort are considered combine-/divide-and-conquer algorithms rather than dynamic programming because they split the work into overlapping subproblems.
84. Top-down dynamic programming trades _____ for _____.
85. Suppose you wanted to know the number of possible subsets of N tasks that cost less than a certain amount in total. What type of algorithm should you use? Explain.
86. What would be the worst case asymptotic time complexity of the algorithm in the previous question?
87. Suppose you wanted to know the subset of N tasks that contains the most tasks while costing less than a certain amount. What type of algorithm should you use? Explain.
88. What would be the worst case asymptotic time complexity of the algorithm in the previous question?
89. What is the difference between backtracking and branch-and-bound?

EECS 281, Week 8: Wednesday

90. Implement `countDuplicates`, a function that counts the total number of duplicate numbers in a vector. For example, in a vector with elements {183, 490, 381, 281, 381, 281, 280, 281}, the total number of duplicates is 3 (2 for 281 and 1 for 381).

```
int countDuplicates(const vector<int> &numbers) {
```

91. Write an efficient function to find the sum of contiguous elements in a subarray within a one-dimensional vector of integers which has the largest sum. What is the time complexity? What is the space complexity? What kind of algorithm is it?

```
int findMaxSubarraySum(const vector<int> &numbers) {
```

EECS 281, Week 8: Wednesday

92. State the N-Queens problem. How to solve it?

93. State the Knapsack problem. Then solve it with backtracking and branch-and-bound.

94. Implement `isSubsetSum`, a function that solves the Subset Sum problem: given a set of n integers S and a value W , determine if there is a subset of integers that sum to W .

```
bool isSubsetSum(const vector<int> &numbers, int sum) {
```

EECS 281, Week 8: Wednesday

95. Given a text and a wildcard pattern, implement wildcard pattern matching algorithm that finds if wildcard pattern is matched with text. The matching should cover the entire text (not partial text). The wildcard pattern can include the characters ? (matches any single character) and * (matches any sequence of characters, including the empty sequence).

Text = "baaabab",

Pattern = "*****ba*****ab", output : true

Pattern = "baaa?ab", output : true

Pattern = "ba*a?", output : true

Pattern = "a*ab", output : false

```
int match(const string &text, const string &pattern) {
```

EECS 281, Week 8: Wednesday

96. Given a weighted graph (whose edges can have negative weights) represented with an adjacency matrix, find the lengths of the shortest paths between all pairs of vertices. Print the lengths of all shortest paths at the end of the function.

Hint: this is Floyd-Warshall algorithm.

```
void shortestPaths(const vector<vector<int>> &graph) {
```