Question 1 (basic): complexity

Below are some operations parametrized by a variable n.

A. Merge sorting an array of size n.

B. Adding n^2 elements to an empty dynamic array.

C. Binary search in a sorted array of size 2^n.

Answer

For each operation, state its asymptotic time complexity in terms of n.
(For example, you can use O-notation with a simple, but exact bound.)

Question 2 (basic): heaps

Here is a minimum-heap with integer values:

0	1	2	3	4	5	6	7	8
2	12	5	16	20	8	10	33	17

We remove the minimum.

Answer

Draw the resulting heap in tree representation.

Question 3 (basic): searching and sorting

Here are my pets:

- **A**ndy, 5-year old ant
- **B**aba, 2-year old bee
- **C**indy, 18-year-old cat
- **D**avid, 9-year-old duck
- **E**lon, 3-year-old elephant
- **K**eke, 8-year old koala
- **Z**ack, 14-year-old zebra

I put them in an array alphabetically:

0	1	2	3	4	5	6
A	B	C	D	E	K	Z

But now I want to sort them **by age** instead (young before old). I heard quicksort can do that.

Can you show me how the partitioning works? Andy has agreed to be the pivot.
You only have to do the initial (first) partitioning, not the rest of quicksort.

Answer

State the swaps performed by the partitioning:

Show the array after partitioning:

0	1	2	3	4	5	6

If you used a partitioning algorithm different from that of your course, explain it:

Question 4 (basic): graphs

Here is an undirected graph with integer weights:

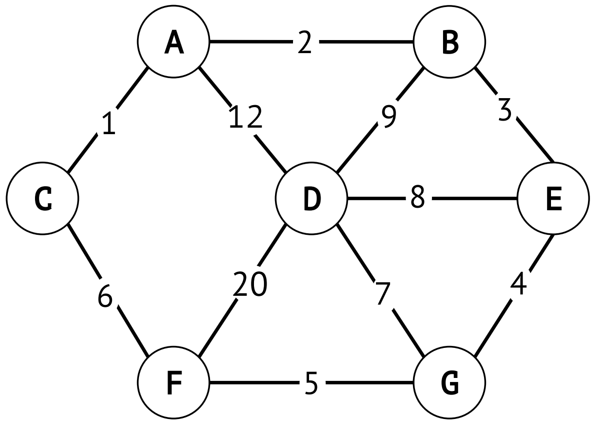

We run **Prim's algorithm** with starting node **D**.

Answer

Draw the resulting **minimum spanning tree** and **state the order** in which its edges are added.

Question 5 (basic): abstract data types

Here are some real-life situations:

A. You want to store the points each student gets in this exam. If they complain about their grade, you need to check their points to see if you calculated the grade incorrectly.

B. You don't have time to read all your emails. So you store the unread ones away. Later when you have time, you go through them by order of recency (most recent one first).

C. Someone gives you a map of Europe. You want to represent the information of which countries border each other.

For each situation, pick the abstract data type that models the situation most appropriately.
You can choose from the following:

- graph
- map
- priority queue
- queue
- set
- stack

No justification is necessary.

Answer

Question 6 (basic): hash tables

We use a separate-chaining hash table to store integer values, using a **sorted list** in each slot. The hash function is modular compression (the remainder of dividing by the table length).

The following hash table of length 5 has gotten quite full:

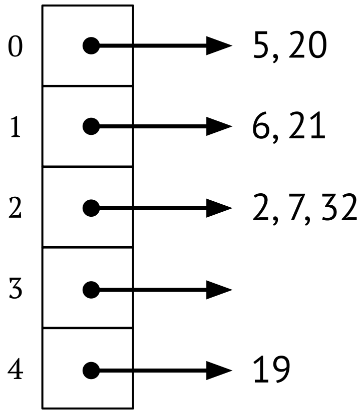

We resize it to table length 8.

Answer

Draw the resulting hash table. What is the new **load factor**?

Question 7 (basic): search trees

Here is a 2-3 tree (null nodes not drawn) with integer values:

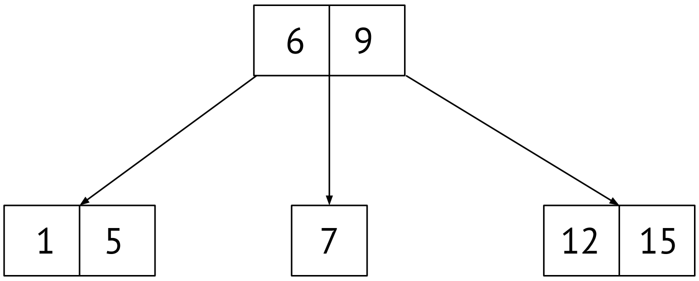

We insert value 4.

Answer

Draw the resulting 2-3 tree.

Question 8 (basic): order of growth

Here are some mathematical functions in a variable n:

- $f(n) = 495n^3 + 182 \cdot 3^n$
- $g(n) = 2024\, log(n + 12)$
- $h(n) = 1000 + 40\, n^2 + 64\, n^4$
- $k(n) = 9001$

Answer

Order the functions *f*, *g*, *h*, *k* according to their **growth rate**.

The smallest growth rate (i.e., slowest growing) should come first.

Question 9 (advanced): shortest cycle

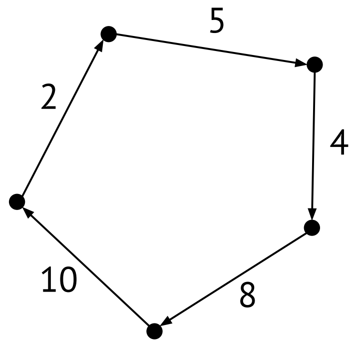

We consider directed graphs with natural number weights.

The *length* of a cycle is the sum of the weights of its edges.

Describe an algorithm that takes a graph and finds the smallest length of a cycle in that graph (returning ∞ if no cycle exists).

The graph is represented using adjacency lists, for example:

- graph.nodes(): Set<V>
- graph.outgoingEdges(*v*: V): List<E>

Your algorithm should run in time O(|*V*| |*E*| log(|*E*|)) where |*V*| is the number of nodes and |*E*| is the number of edges of the given graph. You can assume that:

- the above graph API methods are O(1),
- iterating over a set or list does not have hidden costs.

You can freely use data structures and algorithms from the course – you do not have to explain how they work.

Answer

function *length_of_shortest_cycle*(*graph*: **AdjacencyListGraph**) → natural number or ∞:

Question 10 (advanced): complexity

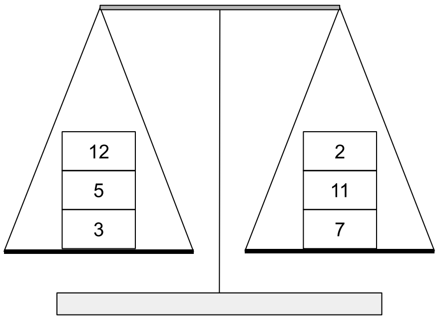

The following function balances an array of weights:

```
function balance(weights):
    left = new stack (using linked list)
    right = new stack (using linked list)

    function helper(k: int, left_total: int, right_total: int) → bool:
        if k equals length(weights):
            return left_total equals right_total

        left.push(weights[k])
        if helper(k + 1, left_total + weights[k], right_total):
            return true
        left.pop()

        right.push(weights[k])
        if helper(k + 1, left_total, right_total + weights[k]):
            return true
        right.pop()

        return false

    if helper(0, 0, 0):
        print "Solution found. Left side:"
        until left is empty: print left.pop()
        print "Right side:"
        until right is empty: print left.pop()
    else:
        print "Impossible to balance."
```

Determine the asymptotic **time complexity** and **space complexity** in the number *N* of weights.

- Printing a weight takes constant time and space.
- Printing a string takes time and space that is linear in the length of the string.
- Remember that your answer needs to be justified.

Question 11 (advanced): tree rotations

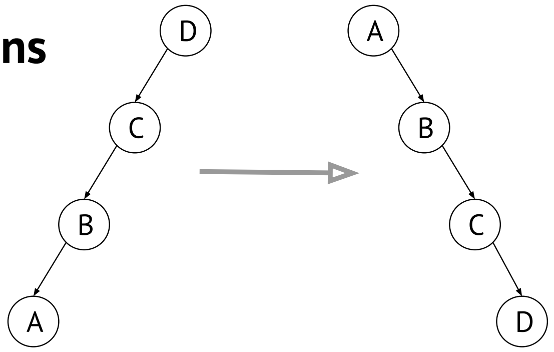

This question is about binary search trees (BSTs).
The following class represents non-empty nodes:

class *Node*:
  value
  left: Node
  right: Node

Recall that binary search trees can be unbalanced. Here, we look at two special cases:

- A binary tree is *left extreme* if the right child of any node is empty.
- A binary tree is *right extreme* if the left child of any node is empty.

We wish to turn a left extreme BST into a right extreme BST **using only rotations**.
Fortunately, I have already implemented the rotation functions for you to use.
They take a node and return the node that replaces it after the rotation.

- *rotate_left*(*node*: Node) → Node
- *rotate_right*(*node*: Node) → Node

Design an algorithm *left_to_right* that takes a left extreme BST and returns a right extreme BST storing the same values. Constraints:

- You may not create any new nodes or change node values.
- You may only change a child pointer by replacing it with the result of applying a rotation.

Answer

function *left_to_right*(*node*: Node) → Node:

Question 12 (advanced): heaps

In a binary minimum-heap, we rely on two operations to restore the heap invariant:

- *swim up*: while node is smaller than parent, swap it with parent,
- *sink down*: while node is larger than child, swap it with smaller child.

We have an array of *N* elements (representing a complete binary tree) and wish to turn it into a minimum heap. Instead of inserting the elements one-by-one, we want to run swim-up and sink-down operations in the array in some order to make the heap invariant satisfied.

Here are two possible strategies:

A. Go over the array from start to end and run sink-down at each position.
B. Go over the array from end to start (reverse order) and run sink-down at each position.

Answer

Determine which of the above strategies work correctly.

- For the ones that do, explain why.
- For the ones that do not, show a small counterexample.