Section 4: Balanced Trees Solutions

0. The ABC's of AVL Trees

What are the constraints on the data types you can store in an AVL tree? When is an AVL tree preferred over another dictionary implementation, such as a HashMap?

1. Let's Plant an AVL Tree

Insert `10, 4, 5, 8, 9, 6, 11, 3, 2, 1, 14` into an initially empty AVL Tree.

2. MinVL Trees

Draw an AVL tree of height 4 that contains the minimum possible number of nodes.

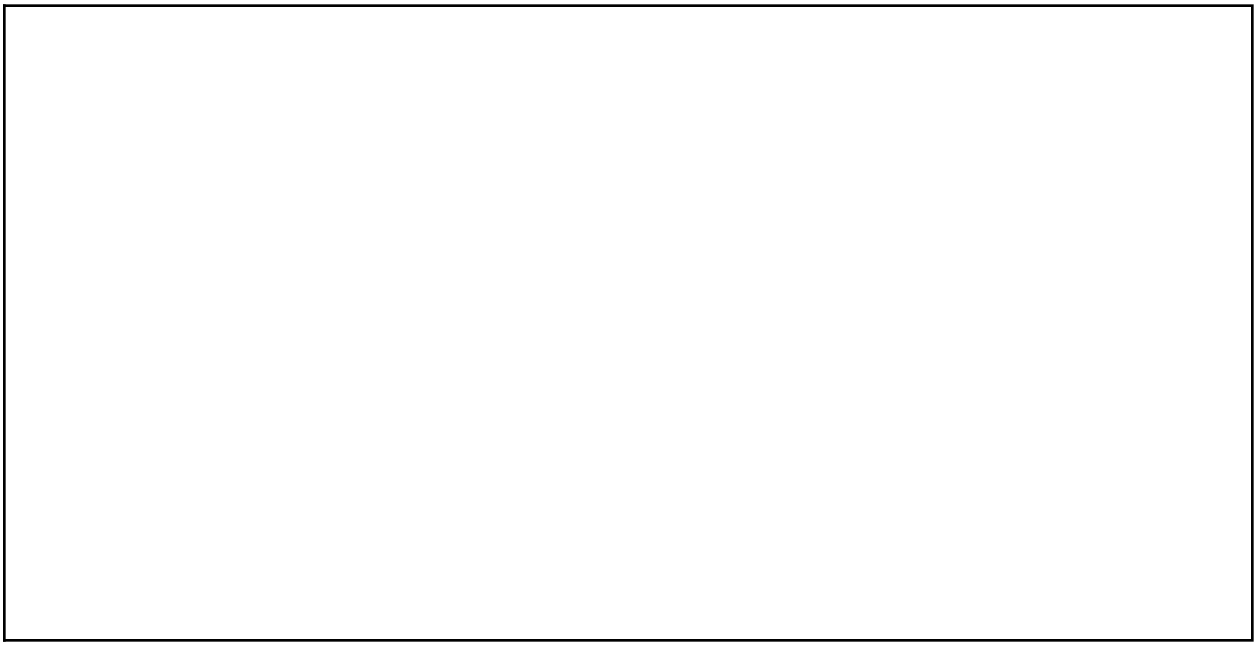

3. AVL Trees

Insert `6, 5, 4, 3, 2, 1, 10, 9, 8, 7` into an initially empty AVL Tree.

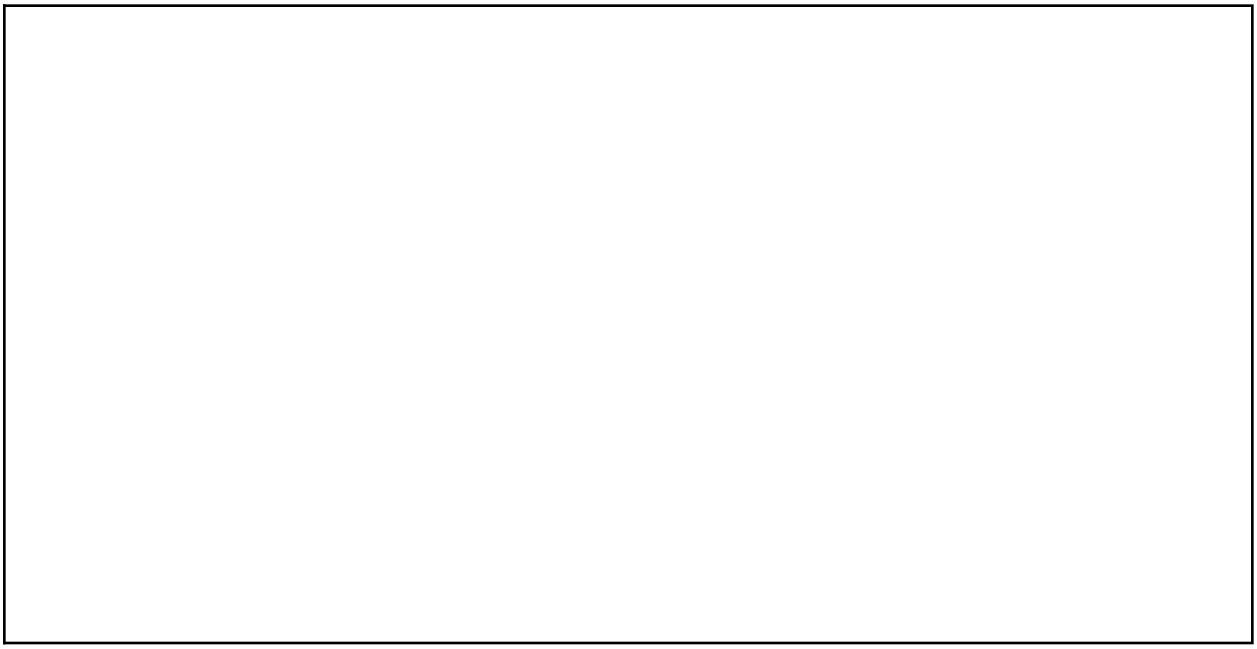

4. The ABC's of B-Trees

a) What properties must a B-tree of `n` values have with given values for M and L?

b) Give an example of a situation that would be a good job for a B-tree. Furthermore, are there any constraints on the data that B-trees can store?

5. Implement a B-Tree? Nah, Let's Analyze!

Given the following parameters for a B-Tree with a page size of 256 bytes:

- Key Size = 8 bytes
- Pointer Size = 2 bytes
- Data Size = 14 bytes per record (includes the key)

Assuming that M and L were chosen appropriately, what are M and L? Recall that M is defined as the maximum number of pointers in an internal node, and L is defined as the maximum number of values in a leaf node. Give a numeric answer and a short justification based on two equations using the parameter values above.

6. Oh, B-Trees

Find a tight upper bound on the worst case runtime of these operations on a B-tree. Your answers should be in terms of M, L, and n.

a) Looking up the value of a key

b) Inserting and deleting a key-value pair

7. B-Trees

a) Insert the following into an empty B-Tree with $M = 3$ and $L = 3$:
`12, 24, 36, 17, 18, 5, 22, 20`
(Assume these numbers are both the keys and the values.)

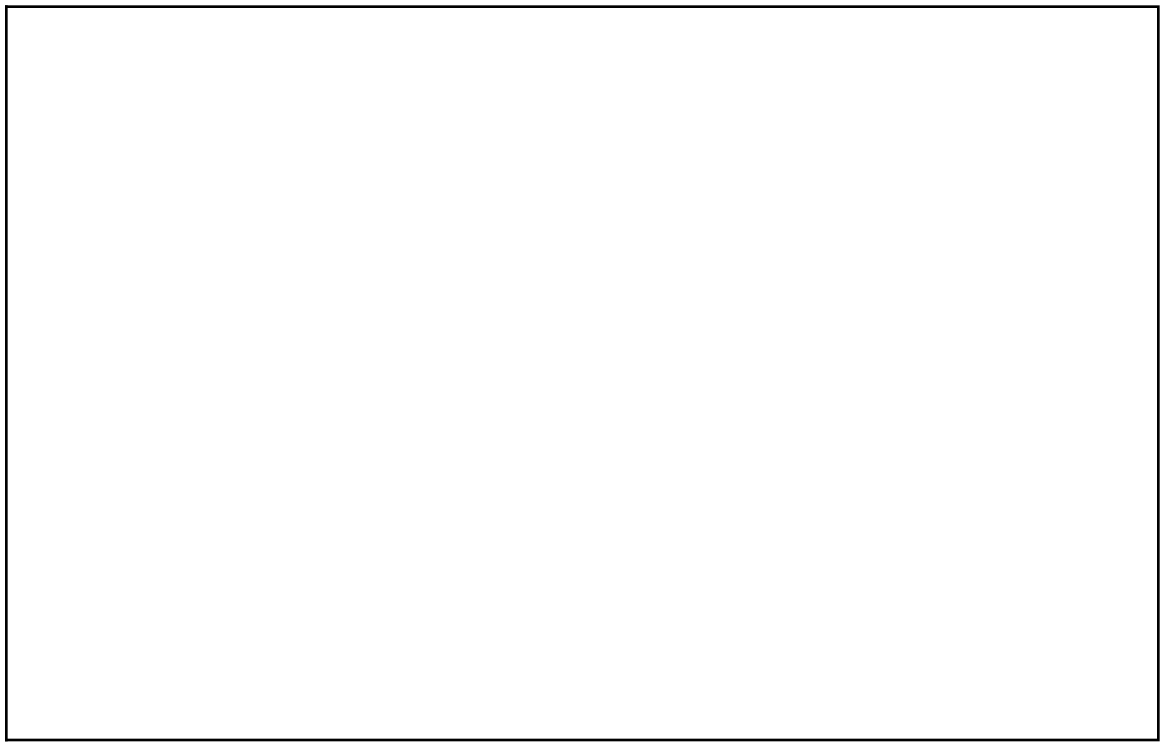

b) Delete `17, 12, 22, 5, 36.` (Assume these numbers are both the keys and the values.)

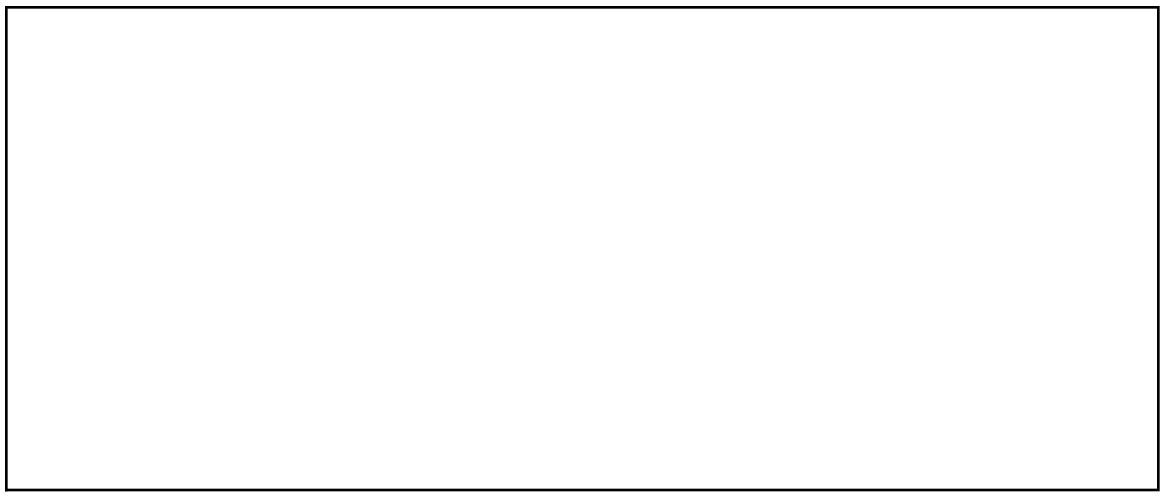

c) Given the following parameters for a B-Tree with $M = 11$ and $L = 8$
   - Key Size = 10 bytes
   - Pointer Size = 2 bytes
   - Data Size = 16 bytes per record (includes the key)

   Assuming that M and L were chosen appropriately, what is the likely page size on the machine where this implementation will be deployed? Give a numeric answer and a short justification based on two equations using the parameter values above.

8. It's Fun to B-Trees!

a) Insert the following into an empty B-Tree with $M = 3$ and $L = 3$:
`3, 32, 9, 26, 6, 21, 8, 4, 5, 30, 31`
(Assume these numbers are both the keys and the values.)

b) Delete `4, 5, 21, 9, 31, 3, 26, 8.` (Assume these numbers are both the keys and the values.)