Improving Map/Dictionary

- unsorted leads to O(1) insert/delete but O(n) search for both arrays and linked lists
- sorted leads to differences between arrays and linked lists
  - O(log n) search and O(n) delete for arrays
  - O(n) search and O(1) delete for linked lists

Can we do better?

- O(n) delete in arrays is due to shifting – can't do much about that
  - circular arrays worked for queues because insert/delete was only at the ends
- can we exploit the sorted order to improve searching in linked lists?
  - like binary search in arrays, perhaps...

Improving an Implementation – Map

Binary search exploits the sorted order – but it requires efficient random access.

Or does it?

- the first iteration of binary search requires knowing the middle element
- successive iterations require knowing the middle element of one of the halves

Finding the middle element is achieved in arrays by arithmetic involving array indexes, but what if we just stored the necessary info instead?

- store instead of computing…

Doing Better

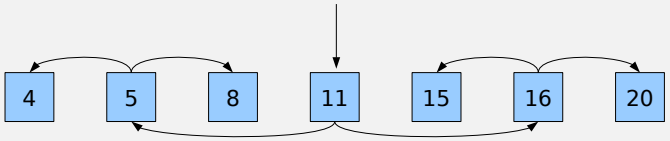

- each middle element only needs to store the location of two other middle elements → binary tree structure
- overall the elements are ordered, so the "other middle elements" are smaller and larger than the "middle element", respectively → binary search tree

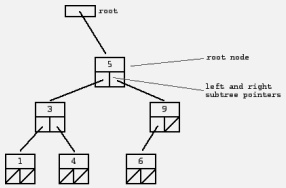

Binary Search Trees

- a binary tree with an ordering property for the elements
  - for every node, all of the elements in the left subtree are less than or equal to the node's element and all of the elements in the right subtree are greater than the node's element
- operations
  - find
  - insert
  - remove
  - visit all elements (traverse) in order

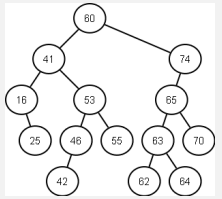

(dummy leaves not shown)

Binary Search Trees

- find
  - start at root
  - go to left child or right child depending on the target value and the current node's value
  - moving down, one child pattern → loop

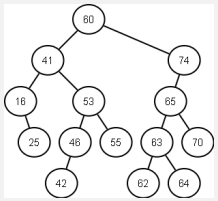

(dummy leaves not shown)

  - observation: if the element isn't there, search ends at a (dummy) leaf

Binary Search Trees

- insert
  - can only insert at a leaf
  - the correct insertion point is the leaf where an unsuccessful search for the element ends up

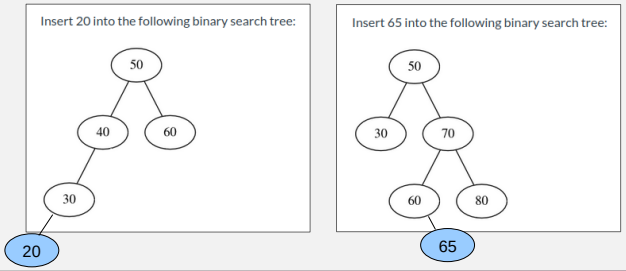

Binary Search Trees

- remove
  - can only remove above a leaf
  - if the element to remove does not have at least one leaf child, swap it with a safe element which does has at least one leaf child
    - i.e. the next element larger or smaller than the one to remove

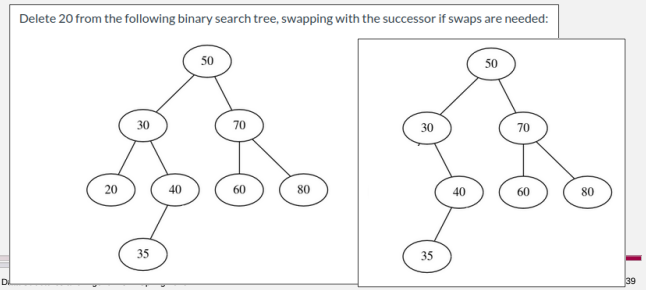

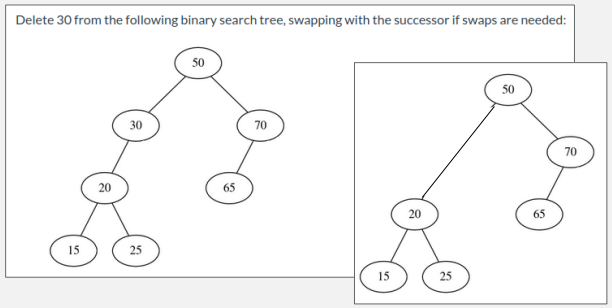

Delete 50 from the following binary search tree, swapping with the successor if swaps are needed:

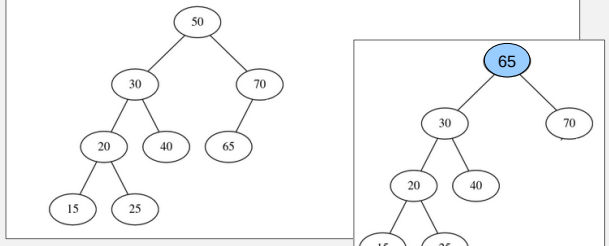

Binary Search Trees

- visit all elements in order
  - moving down, both children pattern → recursion
  - need to visit smaller elements before the current node's element before the larger elements → inorder traversal

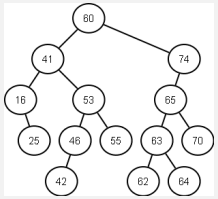

(dummy leaves not shown)

Implementing Map

- can store (key,value) pairs in a binary search tree ordered by key
  - let *h* be the height of the tree
  - all operations are O(h) as it may be necessary to go from the root all the way down to a leaf

BST Height

- height of a binary search tree
  - best case is O(log n)
  - worst case is O(n)

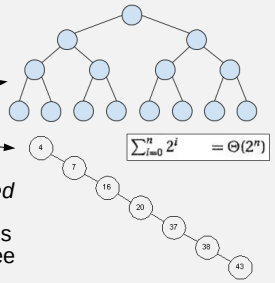

- whether a BST of a given size is *balanced* (O(log n) height) or unbalanced (O(n) height) depends on the order of insertions and removals, not the elements in the tree
- can we do better?
  - try to keep the tree balanced...