CS062

DATA STRUCTURES AND ADVANCED PROGRAMMING

18: Dictionaries and Binary Search Trees

Alexandra Papoutsaki
she/her/hers

Lecture 18: Dictionaries and Binary Search Trees

- Dictionaries
- Binary Search Trees

Dictionaries

- Also known as: symbol tables, maps, indices, associative arrays.
- Key-value pair abstractions that support two operations:
  - Insert a key-value pair.
  - Given a key, search for the corresponding value.
- Supported either with built-in or external libraries by the majority of programming languages.

Basic dictionary API

- `public class Dictionary <Key extends Comparable<Key>, Value>`
- Dictionary(): create an empty dictionary. By convention, values are **not** `null`.
- `void put(Key key, Value val)`: insert key-value pair.
  - Overwrites old value with new value if key already exists.
- `Value get(Key key)`: return value associated with key.
  - Returns `null` if key not present.
- `boolean contains(Key key)`: is there a value associated with key?
- `Iterable keys()`: all the keys in the dictionary.
- `void delete(Key key)`: delete key and associated value.
- `boolean isEmpty()`: is the dictionary empty?
- `int size()`: number of key-value pairs.

Ordered dictionaries

operation	keys	values
min() →	09:00:00	Chicago
	09:00:03	Phoenix
get(09:00:13) →	09:00:13	Houston
	09:00:59	Chicago
	09:01:10	Houston
floor(09:05:00) →	09:03:13	Chicago
	09:10:11	Seattle
select(7) →	09:10:25	Seattle
	09:14:25	Phoenix
keys(09:15:00, 09:25:00) →	09:19:32	Chicago
	09:19:46	Chicago
	09:21:05	Chicago
	09:22:43	Seattle
	09:22:54	Seattle
	09:25:52	Chicago
ceiling(09:30:00) →	09:35:21	Chicago
	09:36:14	Seattle
max() →	09:37:44	Phoenix

size(09:15:00, 09:25:00) *is* 5
rank(09:10:25) *is* 7

Ordered dictionary API

- `Key min()`: smallest key.
- `Key max()`: largest key.
- `Key floor(Key key)`: largest key less than or equal to given key.
- `Key ceiling(Key key)`: smallest key greater than or equal to given key.
- `int rank(Key key)`: number of keys less that given key.
- `Key select(int k)`: key with rank `k`.
- `Iterable keys()`: all keys in dictionary in sorted order.
- `Iterable keys(int lo, int hi)`: keys in `[lo, …, hi]` in sorted order.

Printed dictionaries are all around us

- Dictionary: key = word, value = definition.
- Encyclopedia: key = term, value = article.
- Phonebook: key = name, value = phone number.
- Math table: key = math functions and input, value = function output.
- Unsupported operations:
  - Add a new key and associated value.
  - Remove a given key and associated value.
  - Change value associated with a given key.

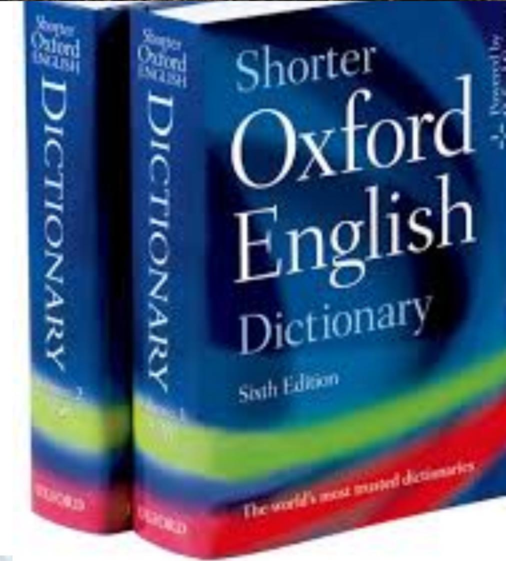

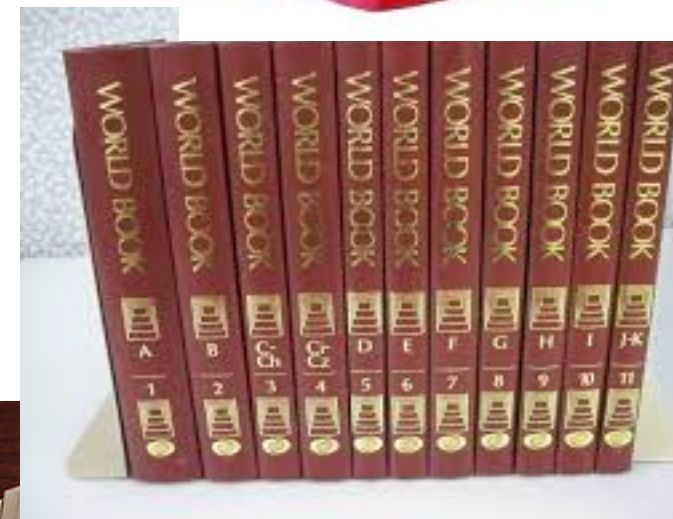

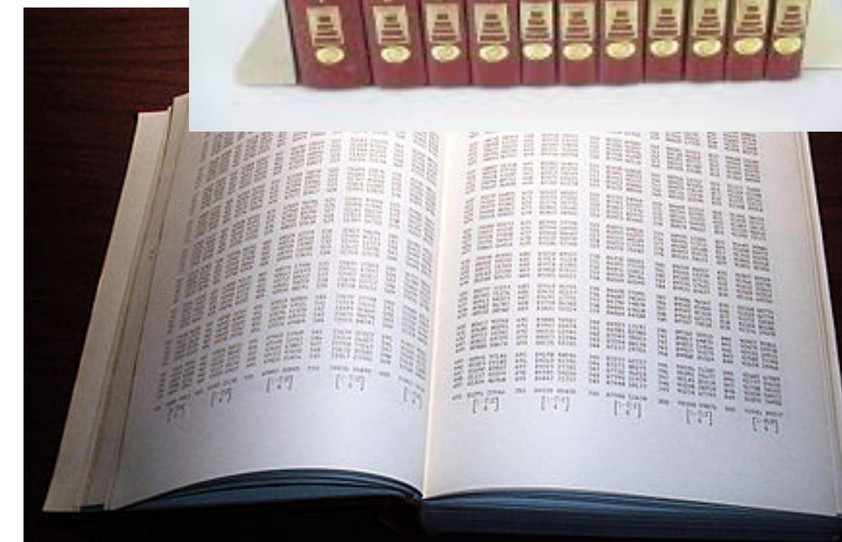

Lecture 18: Dictionaries and Binary Search Trees

- Dictionaries
- Binary search Trees

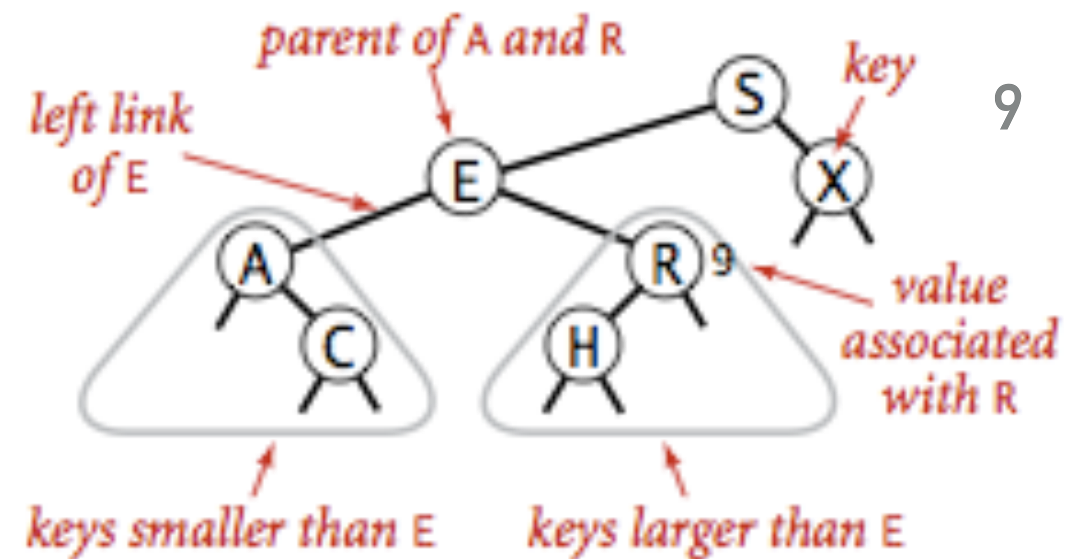

Definitions

- Binary Search Tree: A binary tree in symmetric order.
- Symmetric order: Each node has a key, and every node's key is:
  - Larger than all keys in its left subtree.
  - Smaller than all keys in its right subtree.
- Our textbook uses BSTs to implement dictionaries, therefore each node holds a key-value pair. Other implementations hold only a key.

Differences between heaps and BSTs

	Heap	BST
Used to implement	Priority queues	Dictionaries
Supported operations	Insert, delete max	insert, search, delete, ordered operations
What is inserted	Keys	Key-value pairs
Underlying data structure	(Resizing) array	Linked nodes
Tree shape	Complete binary tree	Depends on data
Ordering of keys	Heap-ordered	Symmetrically-ordered
Duplicate keys allowed?	Yes	No*

*: when BSTs used to implement dictionaries.

BST representation of dictionaries

- We will use an inner class Node that is composed by:
  - A Key that is comparable and a `Value`
  - A reference to the root nodes of the left (smaller keys) and right (larger keys) subtrees.
  - Potentially, the total number of nodes in the subtree that has root this node.
- A BST has a reference to a Node `root`.

BST and Node implementation

```
public class BST<Key extends Comparable<Key>, Value> {
   private Node root;                  // root of BST

   private class Node {
        private Key key;               // sorted by key
        private Value val;             // associated value
        private Node left, right;      // roots of left and right subtrees
        private int size;              // #nodes in subtree rooted at this

        public Node(Key key, Value val, int size) {
            this.key = key;
            this.val = val;
            this.size = size;
        }
    }
```

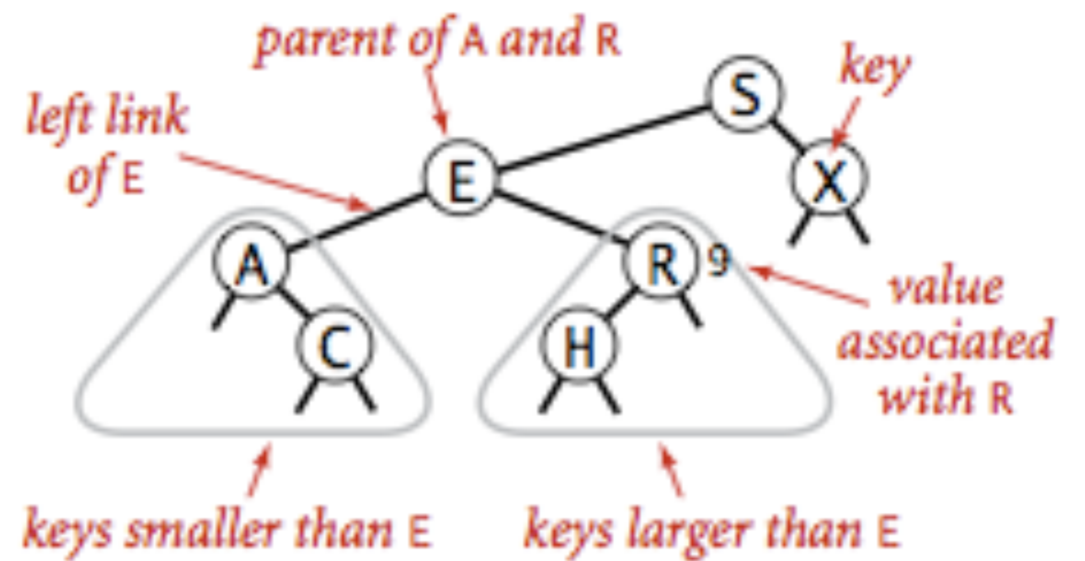

Search for a key

- If less than key in node go to left subtree.
- If greater than key in node go to right subtree.
- If given key and key at examined node are equal, search hit.
- Return value corresponding to given key, or `null` if no such key.
  - In other implementations, you return the last node you reached.
- Number of compares is equal to the depth of the node + 1.

Search example

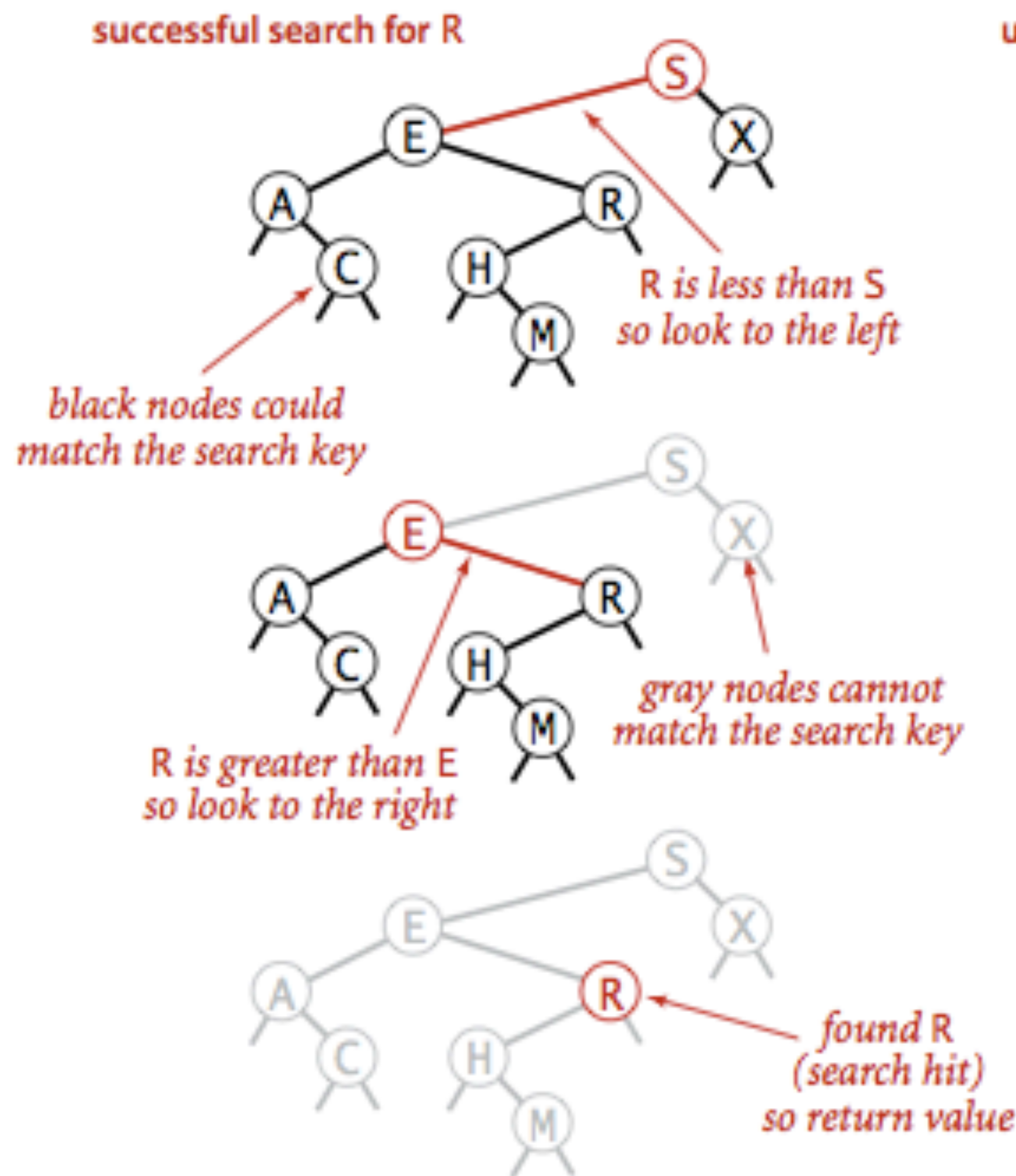

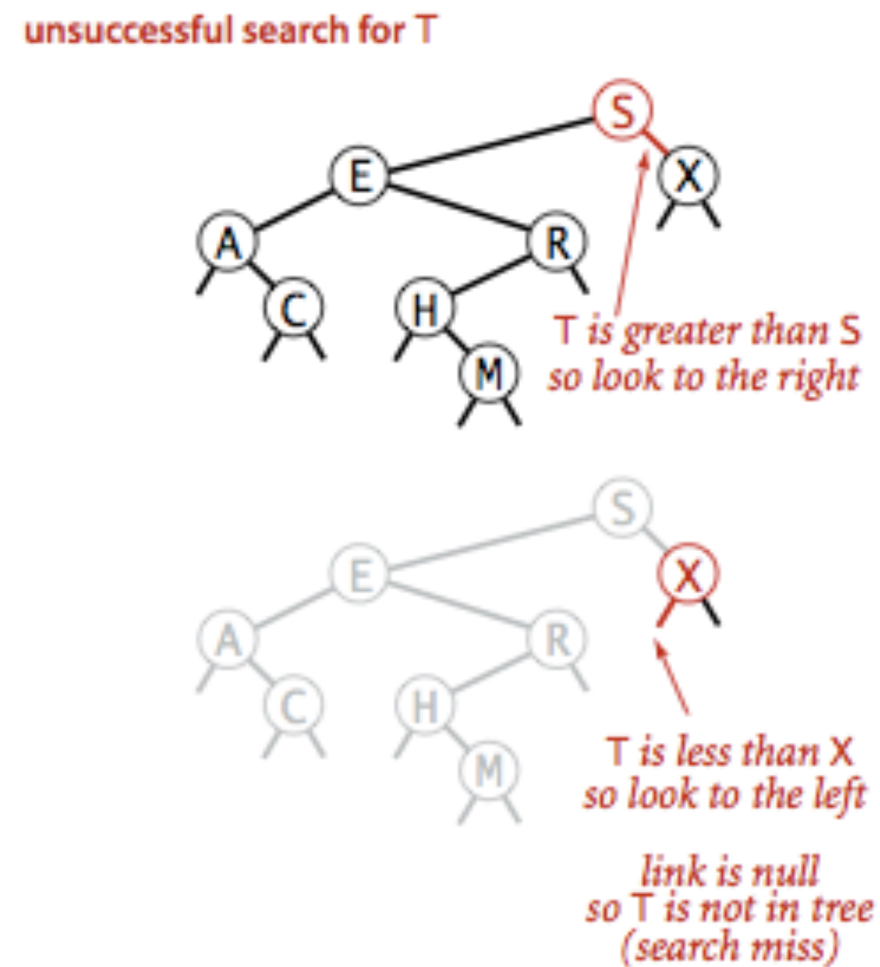

Successful (left) and unsuccessful (right) search in a BST

Search - iterative implementation

```
public Value get(Key key) {
        Node x = root;
        while (x != null) {
                int cmp = key.compareTo(x.key);
                if (cmp < 0)
                        x = x.left;
                else if (cmp > 0)
                        x = x.right;
                else if (cmp == 0)
                        return x.val;
         }
         return null;
}
```

Search - recursive implementation

- ```
  public Value get(Key key) {
       return get(root, key);
  }
  ```
- ```
  private Value get(Node x, Key key) {
       if (x == null)
             return null;
       int cmp = key.compareTo(x.key);
       if (cmp < 0)
           return get(x.left, key);
       else if (cmp > 0)
           return get(x.right, key);
       else
           return x.val;
  }
  ```

Practice Time - Problem 1 Worksheet #18

- Search for the keys 4 and 9 in the following BST:

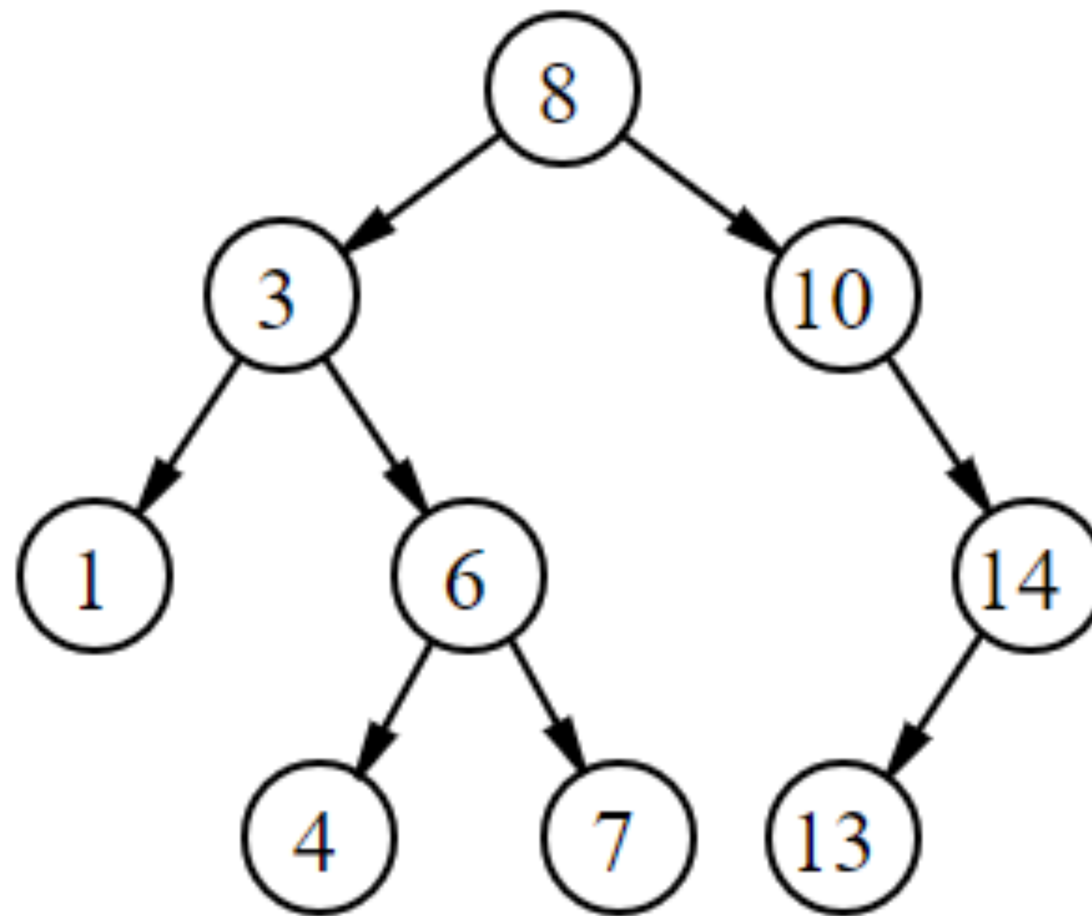

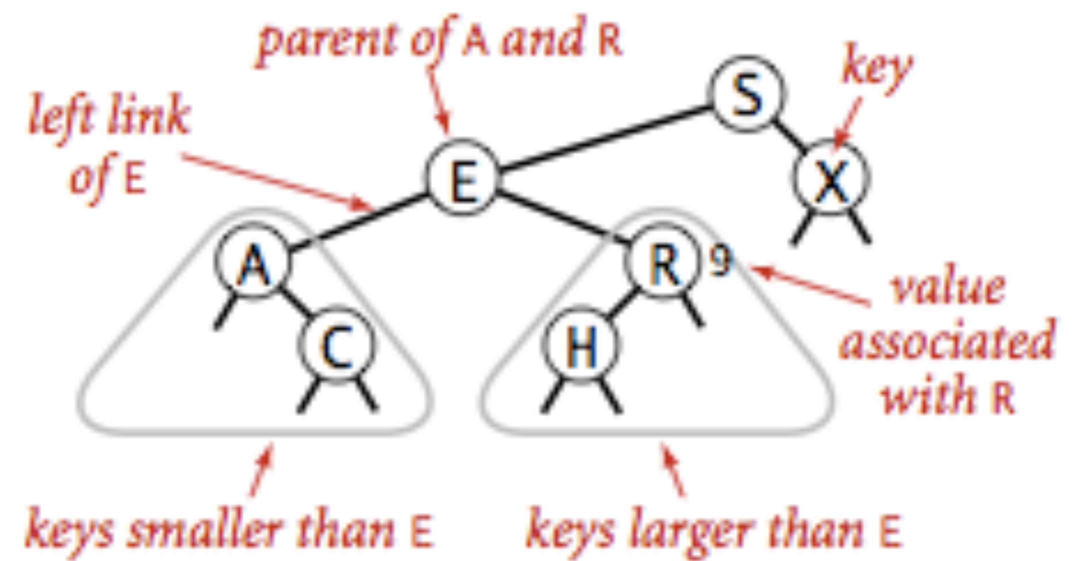

Insert

- If less than key in node go left.
- If greater than key in node go right.
- If null, insert.
- If already exists, update value.
- Number of compares is equal to the depth of the node + 1.

Insert example

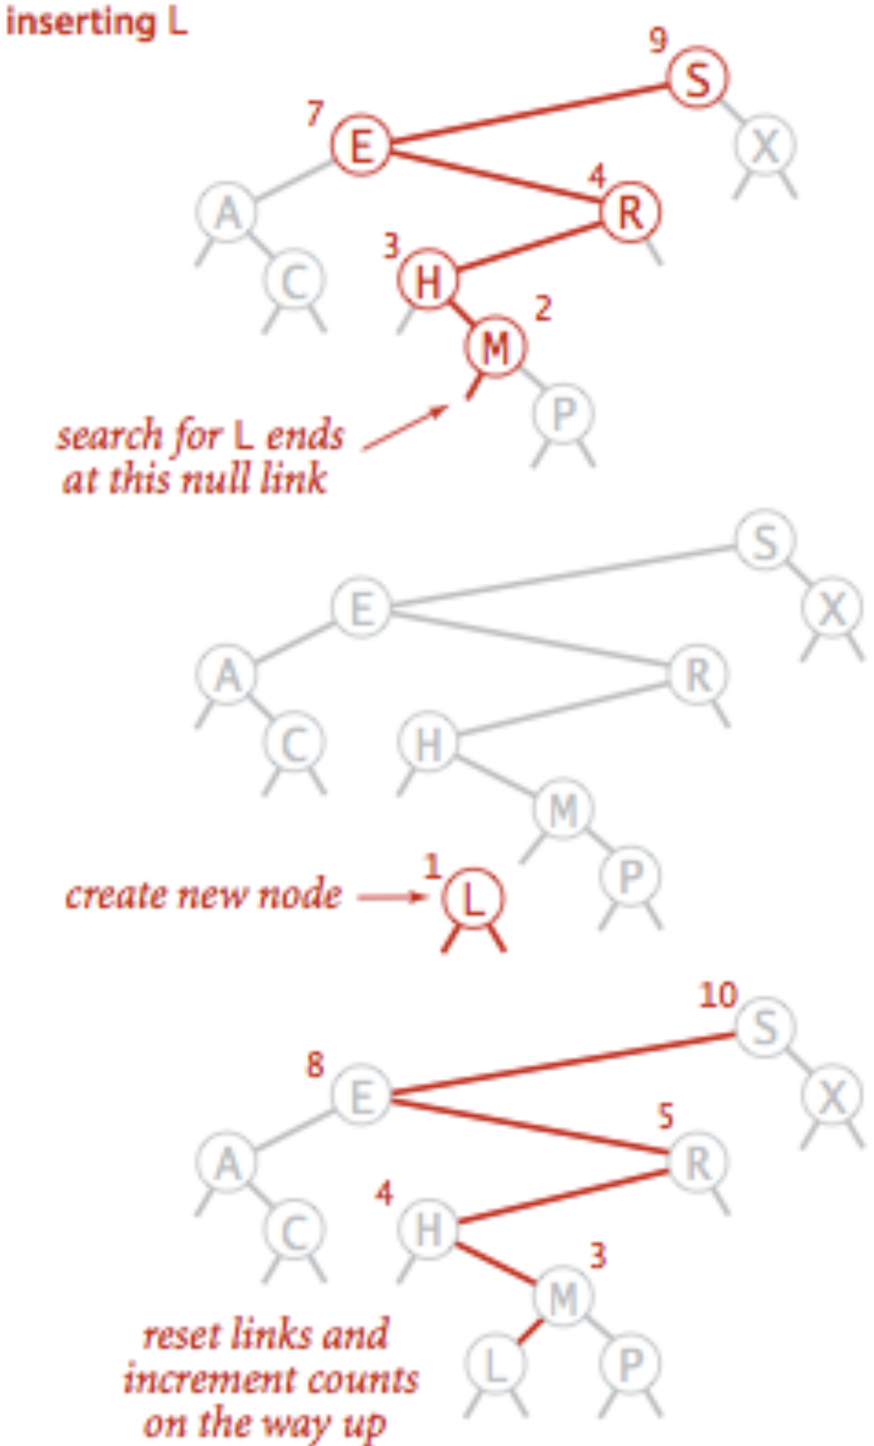

Insertion into a BST

Insert

- ```
  public void put(Key key, Value val) {
      root = put(root, key, val);
  }
  private Node put(Node x, Key key, Value val) {
      if (x == null)
          return new Node(key, val, 1);
      int cmp = key.compareTo(x.key);
      if (cmp < 0)
          x.left = put(x.left, key, val);
      else if (cmp > 0)
          x.right = put(x.right, key, val);
      else
          x.val = val;
      x.size = 1 + size(x.left) + size(x.right);
      return x;
  }
  ```
  ```

Practice Time - Problem 2 Worksheet #18

- Add the keys 4 and then the key 9 in the following BST:

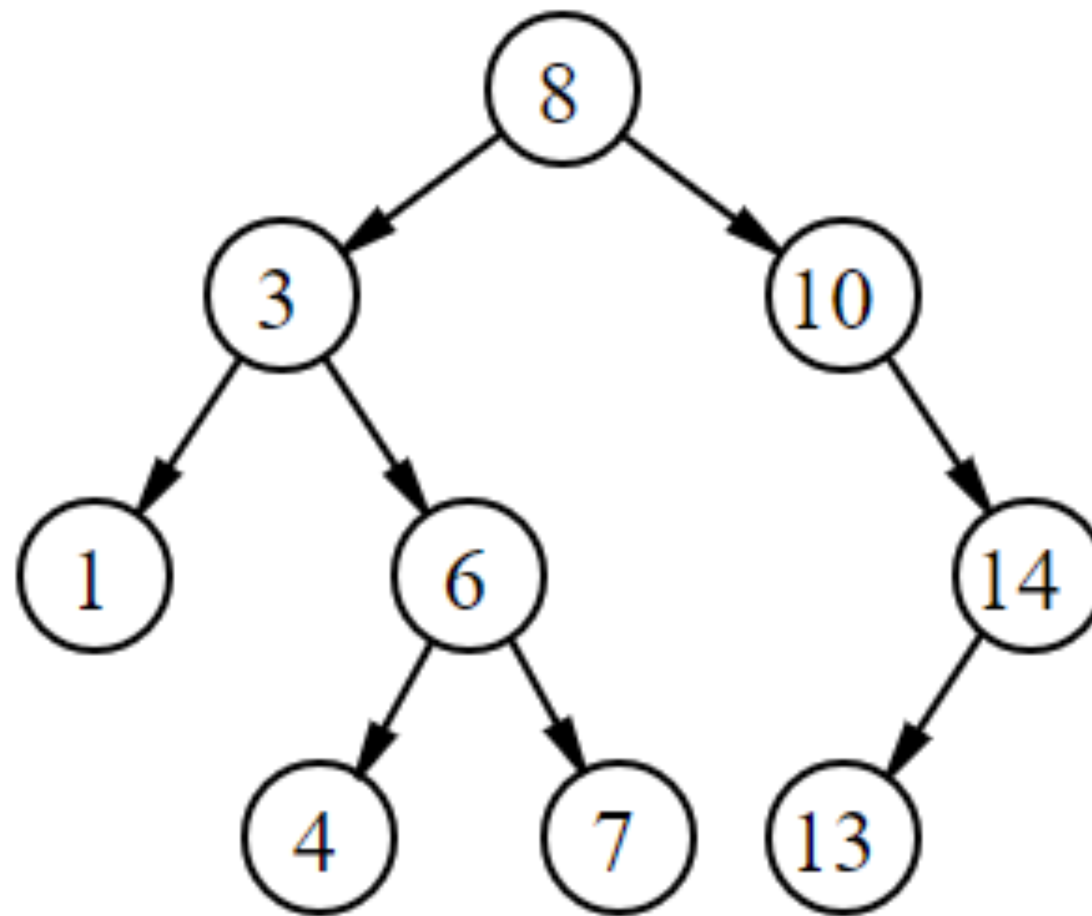

http://algs4.cs.princeton.edu

3.2 Binary Search Tree Demo

Tree shape

- The same set of keys can result to different BSTs based on their order of insertion.
- Number of compares for search/insert is equal to depth of node +1.

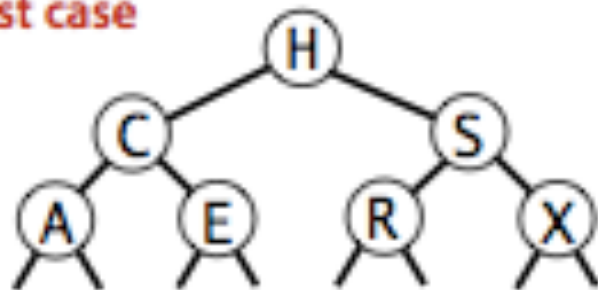

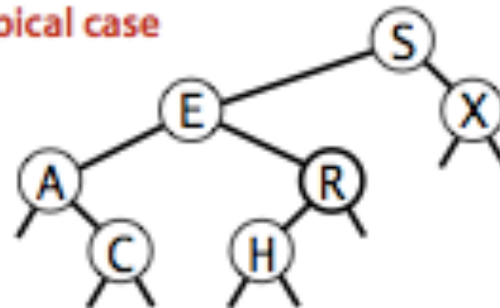

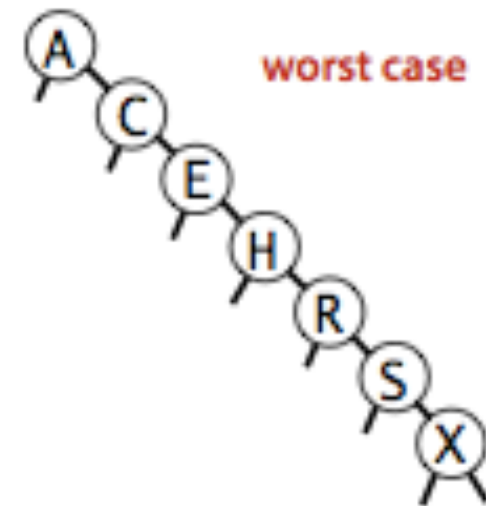

BSTs mathematical analysis

- If n distinct keys are inserted into a BST in random order, the expected number of compares of search/insert is $O(\log n)$.
  - If n distinct keys are inserted into a BST in random order, the expected height of tree is $O(\log n)$. [Reed, 2003].
- Worst case height is n but highly unlikely.
  - Keys would have to come (reversely) sorted!
- All ordered operations in a dictionary implemented with a BST depend on the height of the BST.

Hibbard deletion: Delete node which is a leaf

- Simply delete node.
- Example: delete 52 locates a node which is a leaf and removes it.

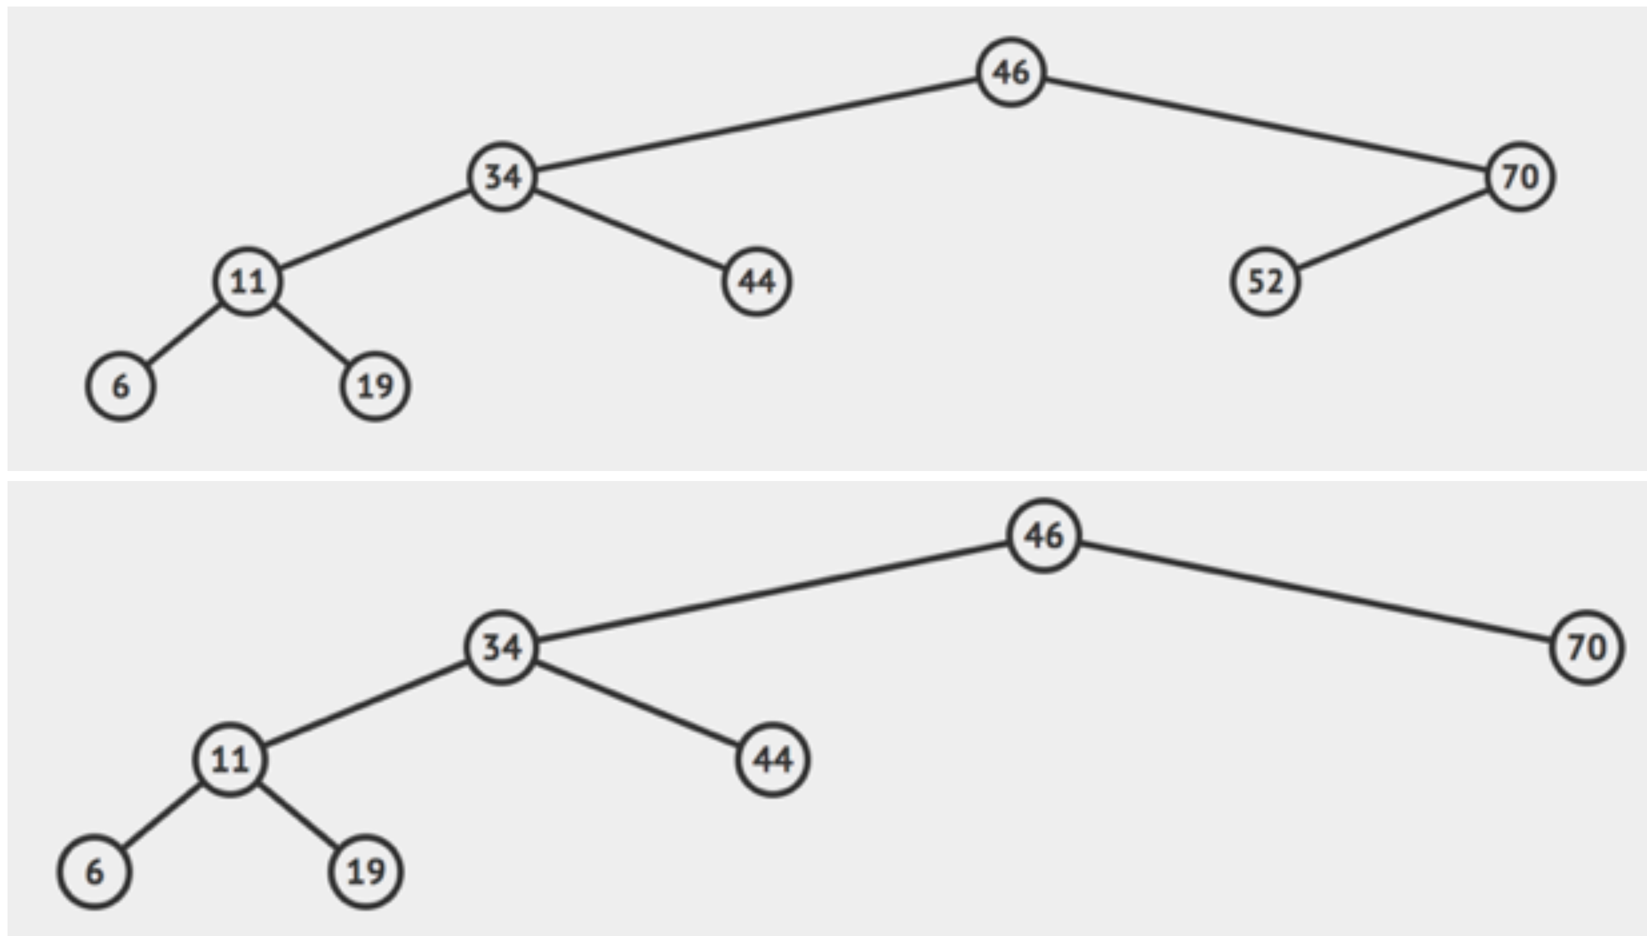

Hibbard deletion: Delete node with one child

- Delete node and replace it with its child.
- Example: delete 70 locates a node which has one child and replaces it with the child.

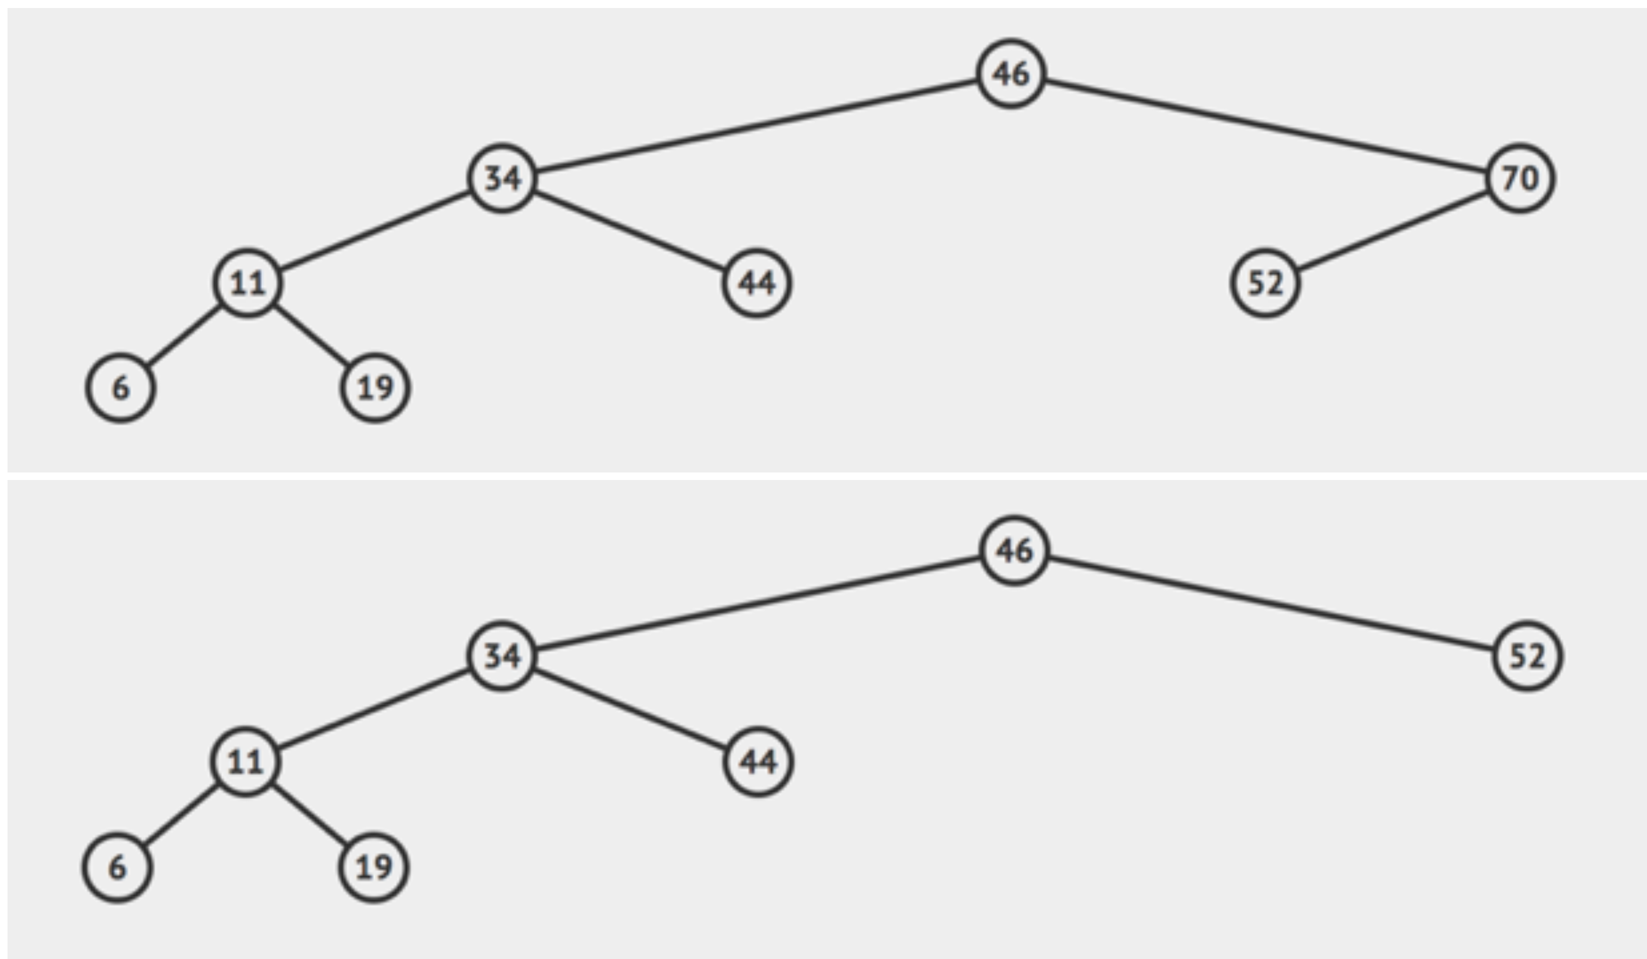

Hibbard deletion: Delete node with two children

- Delete node and replace it with successor (node with smallest of the larger keys). Move successor's child (if any) where successor was.

- Example: delete 50 locates a node which has two children. Successor is 51.

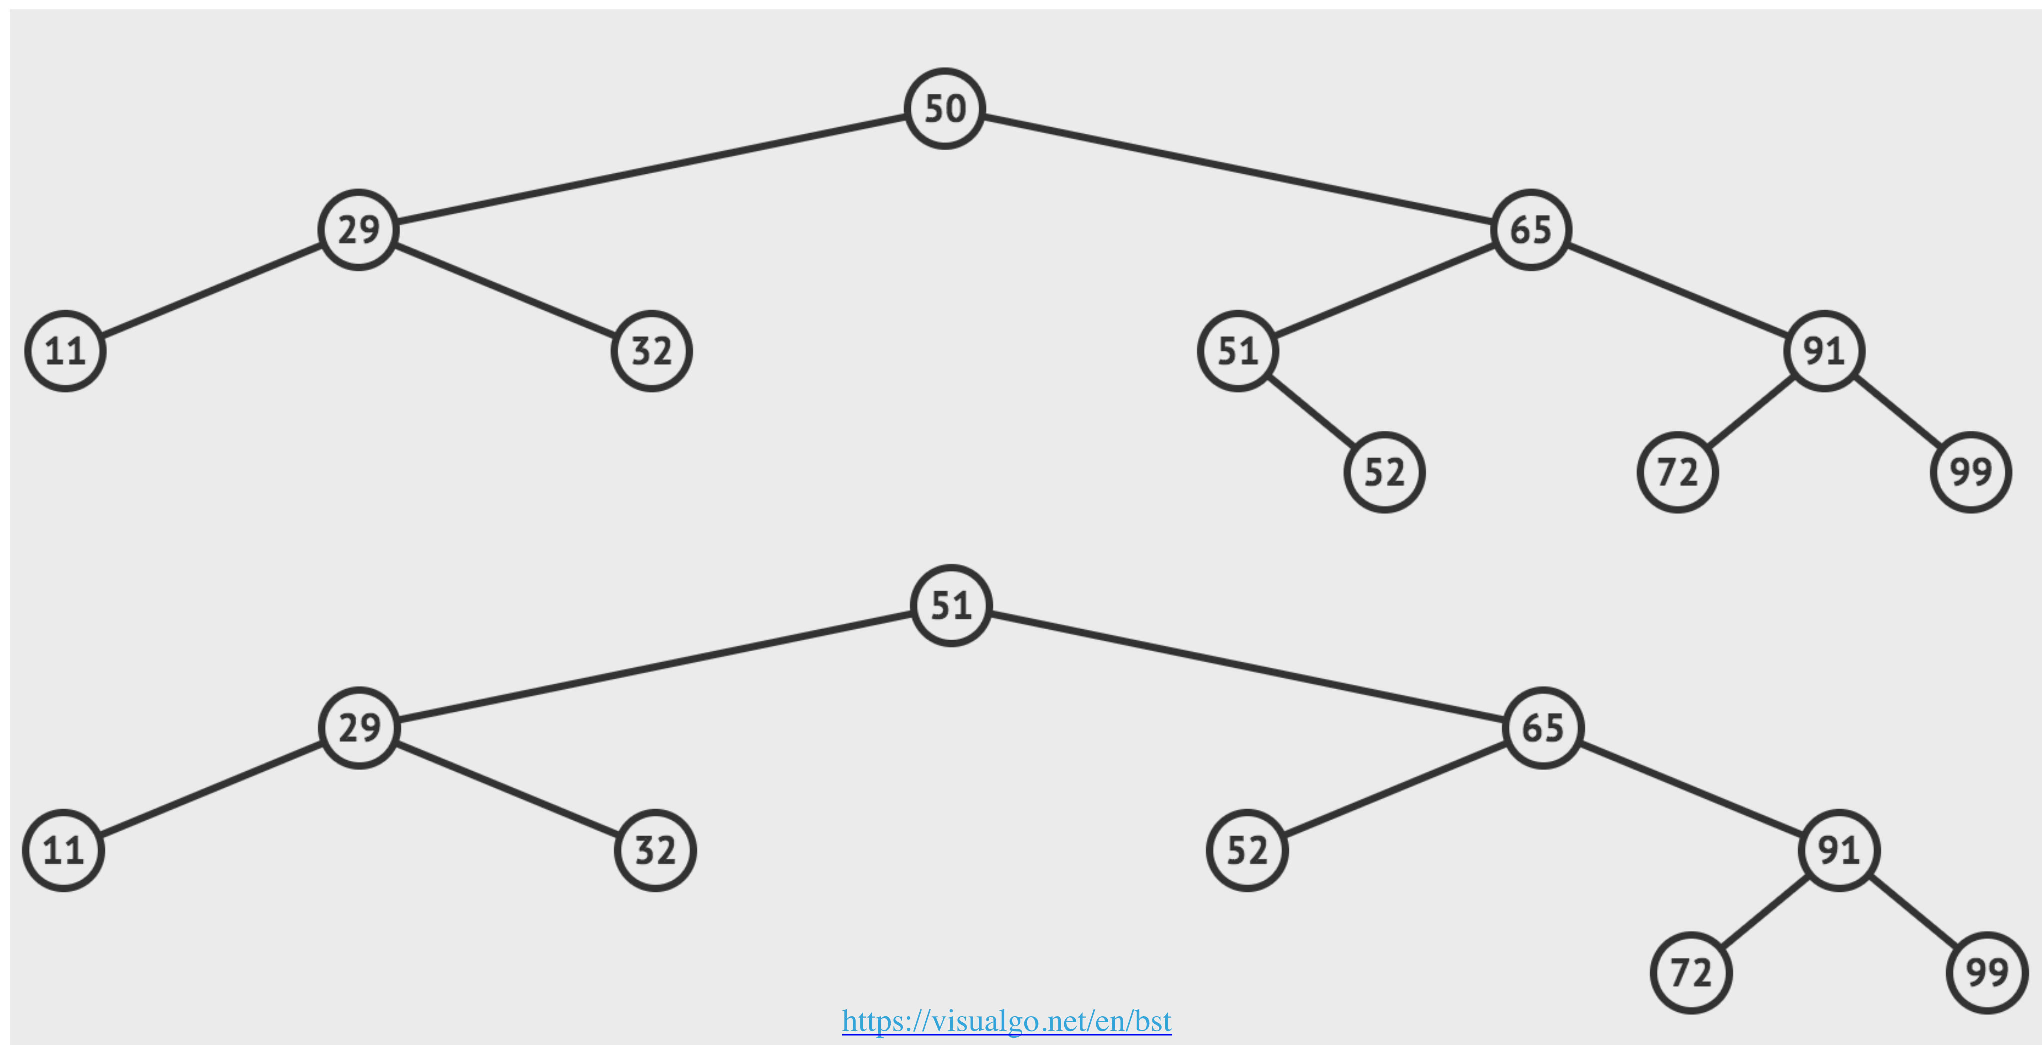

```
public void delete(Key key) {
    root = delete(root, key);
}

 private Node delete(Node x, Key key) {
     if (x == null) return null;

     int cmp = key.compareTo(x.key);
     if (cmp < 0)
         x.left  = delete(x.left,  key);
     else if (cmp > 0)
         x.right = delete(x.right, key);
     else {
         if (x.right == null)
             return x.left;
         if (x.left  == null)
             return x.right;
         Node t = x; //replace with successor
         x = min(t.right);
         x.right = deleteMin(t.right);
         x.left = t.left;
     }
     x.size = size(x.left) + size(x.right) + 1;
     return x;
 }
```

Practice Time - Problem 3 Worksheet #18

- Delete the node 21 following Hibbard's deletion

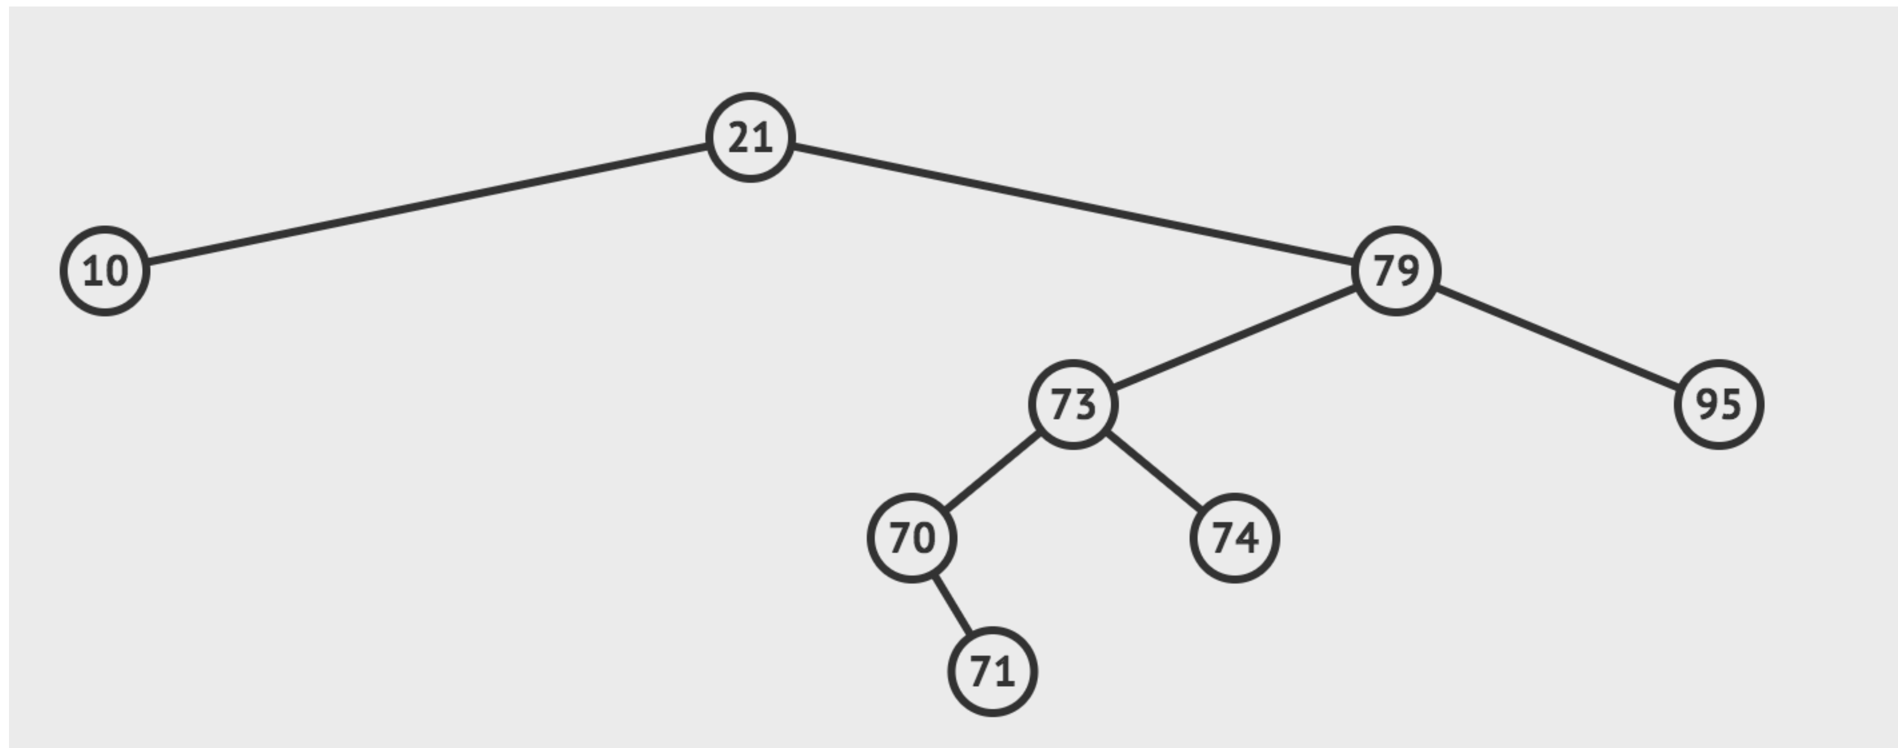

Answer

- Delete the node 21 following Hibbard's deletion

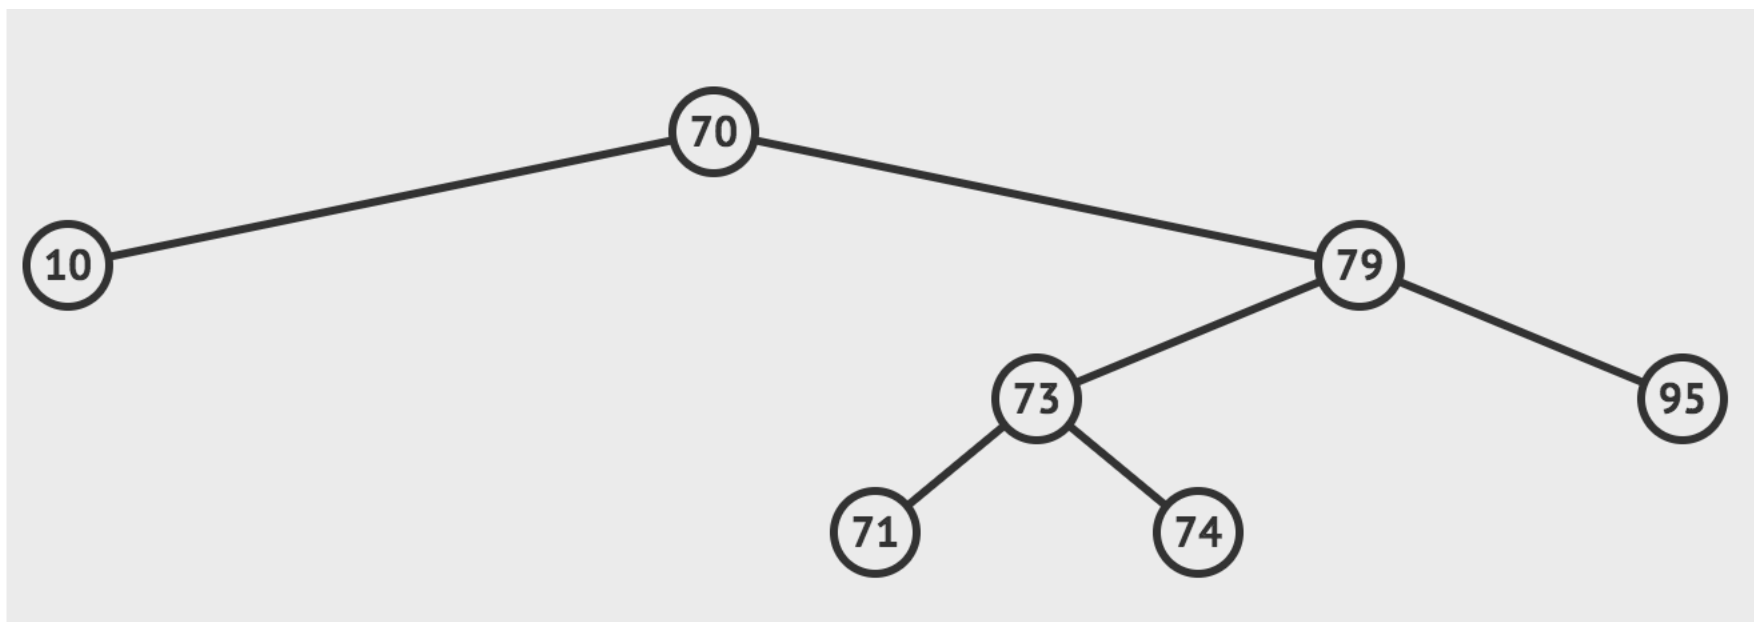

Hibbard's deletion

- Unsatisfactory solution. If we were to perform many insertions and deletions the BST ends up being not symmetric and skewed to the left.
  - Extremely complicated analysis, but average cost of deletion ends up being $\sqrt{n}$. Let's simplify things by saying it stays $O(\log n)$.
  - No one has proven that alternating between the predecessor and successor will fix this.
- Hibbard devised the algorithm in 1962. Still no algorithm for efficient deletion in Binary Search Trees!
- Overall, BSTs can have $O(n)$ worst-case for search, insert, and delete. We want to do better (see future lectures).

Lecture 18: Dictionaries and Binary Search Trees

- Dictionaries
- Binary Search Trees

Readings:

- Recommended Textbook: Chapters 3.2 (Pages 396–414)
- Website:
  - https://algs4.cs.princeton.edu/32bst/
- Visualization:
  - https://visualgo.net/en/bst

Worksheet

- Lecture 18 worksheet

Problem 1

- Draw the BST that results when you insert the keys 5, 1, 19, 25, 17, 5, 19, 20, 9, 15, 14 in that order.

Problem 2

- Inserting the keys in the order A X C S E R H into an initially empty BST gives a worst-case tree where every node has one null link (one child), except one at the bottom that has two null links (it's a leaf). Give five other orderings of these keys that produce worst-case trees.

Problem 3

- Give five orderings of the keys A X C S E R H that when inserted into an initially empty binary search tree, produce best-case trees.

ANSWER 1

- Draw the BST that results when you insert the keys 5, 1, 19, 25, 17, 5, 19, 20, 9, 15, 14 in that order.
- -2 indicates that this node has been updated to the second value associated with that key.

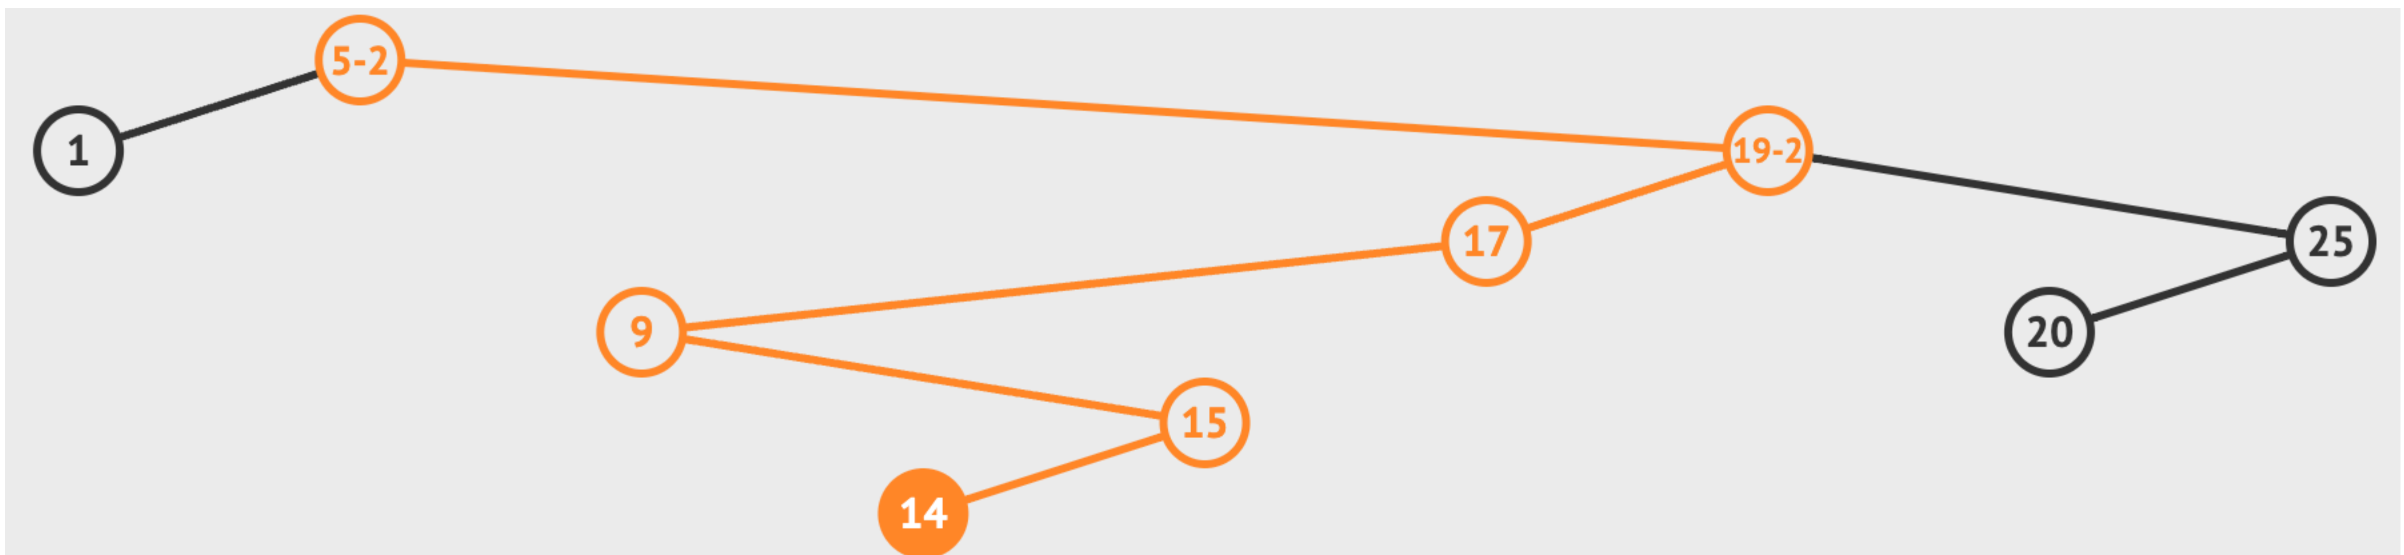

ANSWER 2

- Inserting the keys in the order A X C S E R H into an initially empty BST gives a worst-case tree where every node has one null link (one child), except one at the bottom that has two null links (it's a leaf). Give five other orderings of these keys that produce worst-case trees.
- A C E H R S X
- X S R H E C A
- X A S C R E H
- X A S C R H E
- A X C S E H R

ANSWER 3

- Inserting the keys in the order A X C S E R H into an initially empty BST gives a worst-case tree where every node has one null link (one child), except one at the bottom that has two null links (it's a leaf). Give five other orderings of these keys that produce worst-case trees.
- H C S A E R X
- H C A E S R X
- H C E A S R X
- H S R X C A E
- H S X R C A E