

Chapter 12: Binary Search Trees

A **binary search tree** is a binary tree with a special property called the **BST-property**, which is given as follows:

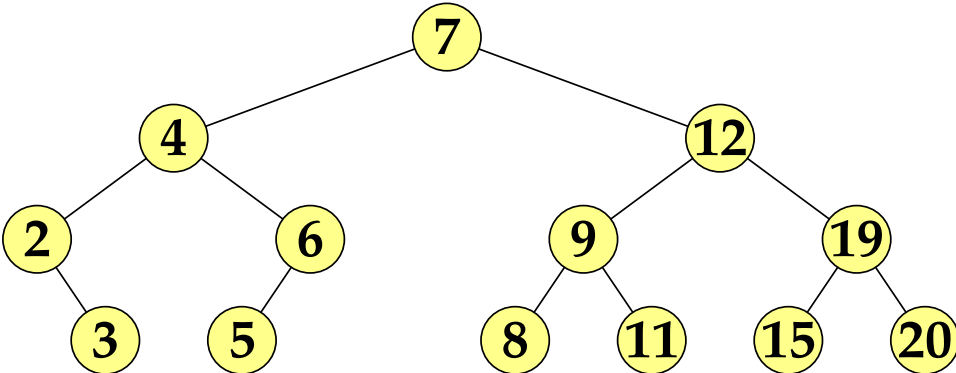
- ★ *For all nodes x and y , if y belongs to the left subtree of x , then the key at y is less than the key at x , and if y belongs to the right subtree of x , then the key at y is greater than the key at x .*

We will assume that the keys of a BST are pairwise distinct.

Each node has the following attributes:

- p , $left$, and $right$, which are pointers to the parent, the left child, and the right child, respectively, and
- key , which is key stored at the node.

An example



Traversal of the Nodes in a BST

By “traversal” we mean visiting all the nodes in a graph. Traversal strategies can be specified by the ordering of the three objects to visit: the current node, the left subtree, and the right subtree. We assume the the left subtree always comes before the right subtree. Then there are three strategies.

1. **Inorder**. The ordering is: the left subtree, the current node, the right subtree.
2. **Preorder**. The ordering is: the current node, the left subtree, the right subtree.
3. **Postorder**. The ordering is: the left subtree, the right subtree, the current node.

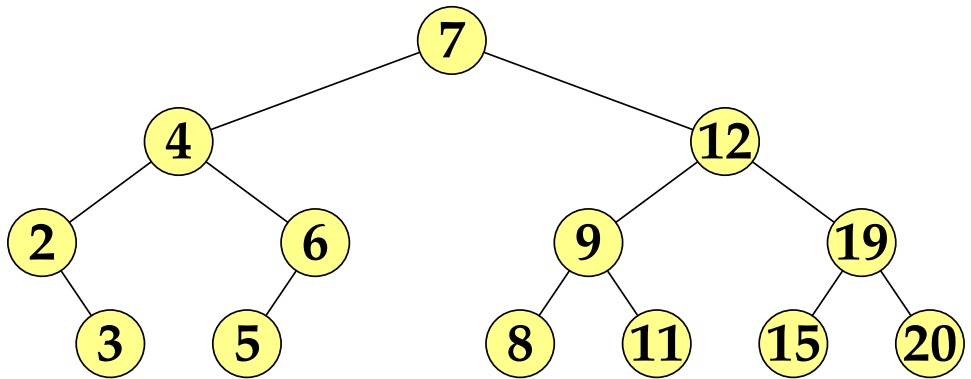
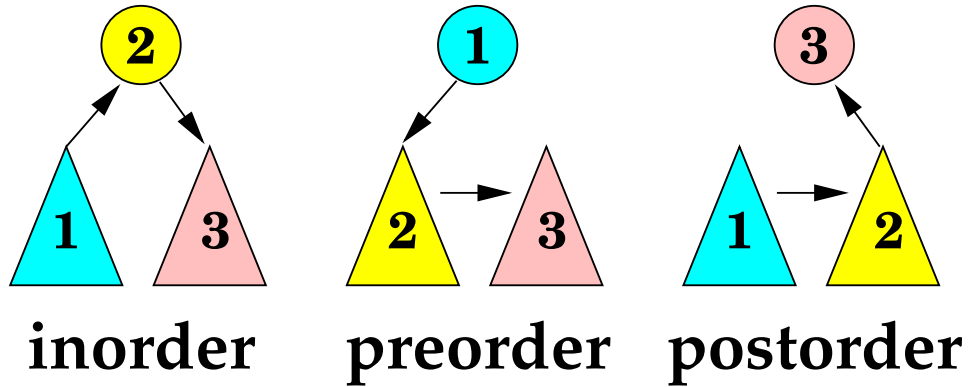
Inorder Traversal Pseudocode

This recursive algorithm takes as the input a pointer to a tree and executed inorder traversal on the tree. While doing traversal it prints out the key of each node that is visited.

Inorder-Walk(x)

- 1: **if** $x = \text{nil}$ **then return**
- 2: Inorder-Walk($\text{left}[x]$)
- 3: Print $\text{key}[x]$
- 4: Inorder-Walk($\text{right}[x]$)

We can write a similar pseudocode for preorder and postorder.



*What is the outcome of
inorder traversal on this BST?
How about postorder traversal
and preorder traversal?*

Inorder traversal gives: 2, 3,
4, 5, 6, 7, 8 , 9, 11, 12, 15,
19, 20.

Preorder traversal gives: 7, 4,
2, 3, 6, 5, 12, 9, 8, 11, 19,
15, 20.

Postorder traversal gives: 3,
2, 5, 6, 4, 8, 11, 9, 15, 20,
19, 12, 7.

So, inorder travel on a BST
finds the keys in
nondecreasing order!

Operations on BST

1. Searching for a key

We assume that a key and the subtree in which the key is searched for are given as an input. We'll take the full advantage of the BST-property.

Suppose we are at a node. If the node has the key that is being searched for, then the search is over. Otherwise, the key at the current node is either strictly smaller than the key that is searched for or strictly greater than the key that is searched for. If the former is the case, then by the BST property, all the keys in the left subtree are strictly less than the key that is searched for. That means that we do not need to search in the left subtree. Thus, we will examine only the right subtree. If the latter is the case, by symmetry we will examine only the right subtree.

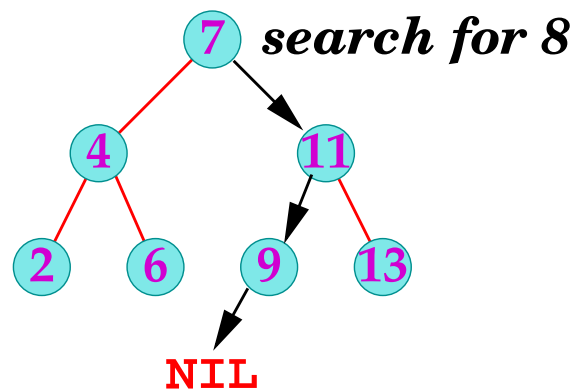
Algorithm

Here k is the key that is searched for and x is the start node.

BST-Search(x, k)

- 1: $y \leftarrow x$
- 2: **while** $y \neq \text{nil}$ **do**
- 3: **if** $\text{key}[y] = k$ **then return** y
- 4: **else if** $\text{key}[y] < k$ **then** $y \leftarrow \text{right}[y]$
- 5: **else** $y \leftarrow \text{left}[y]$
- 6: **return** ("NOT FOUND")

An Example



What is the running time of search?

2. The Maximum and the Minimum

To find the minimum identify the leftmost node, i.e. the farthest node you can reach by following only left branches.

To find the maximum identify the rightmost node, i.e. the farthest node you can reach by following only right branches.

BST-Minimum(x)

- 1: **if** $x = \text{nil}$ **then return** ("Empty Tree")
- 2: $y \leftarrow x$
- 3: **while** $\text{left}[y] \neq \text{nil}$ **do** $y \leftarrow \text{left}[y]$
- 4: **return** ($\text{key}[y]$)

BST-Maximum(x)

- 1: **if** $x = \text{nil}$ **then return** ("Empty Tree")
- 2: $y \leftarrow x$
- 3: **while** $\text{right}[y] \neq \text{nil}$ **do** $y \leftarrow \text{right}[y]$
- 4: **return** ($\text{key}[y]$)

3. Insertion

Suppose that we need to insert a node z such that $k = \text{key}[z]$. Using binary search we find a **nil** such that replacing it by z does not break the BST-property.

BST-Insert(x, z, k)

```
1: if  $x = \text{nil}$  then return "Error"  
2:  $y \leftarrow x$   
3: while true do {  
4:     if  $\text{key}[y] < k$   
5:     then  $z \leftarrow \text{left}[y]$   
6:     else  $z \leftarrow \text{right}[y]$   
7:     if  $z = \text{nil}$  break  
8: }  
9: if  $\text{key}[y] > k$  then  $\text{left}[y] \leftarrow z$   
10: else  $\text{right}[p[y]] \leftarrow z$ 
```

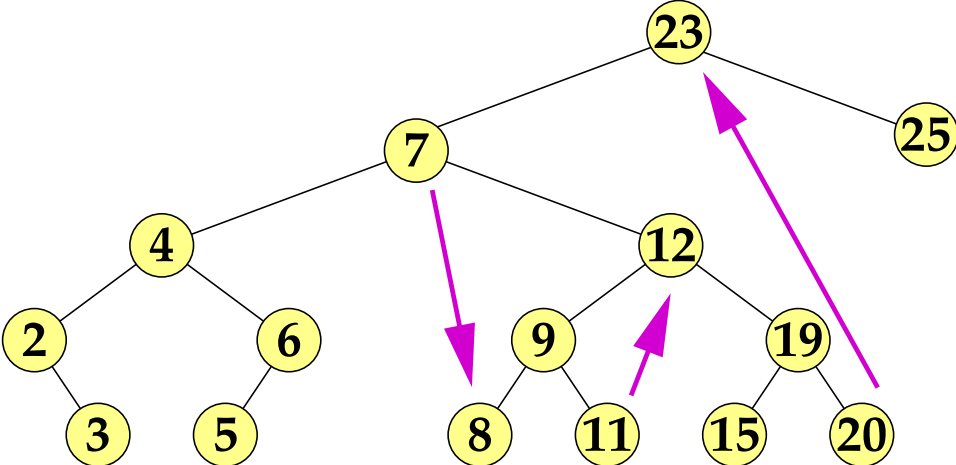
4. The Successor and The Predecessor

The successor (respectively, the predecessor) of a key k in a search tree is the smallest (respectively, the largest) key that belongs to the tree and that is strictly greater than (respectively, less than) k .

The idea for finding the successor of a given node x .

- If x has the right child, then the successor is the minimum in the right subtree of x .
- Otherwise, the successor is the parent of the farthest node that can be reached from x by following only right branches backward.

An Example



Algorithm

BST-Successor(x)

```
1: if  $right[x] \neq nil$  then
2:   {  $y \leftarrow right[x]$ 
3:     while  $left[y] \neq nil$  do  $y \leftarrow left[y]$ 
4:     return ( $y$ ) }
5: else
6:   {  $y \leftarrow x$ 
7:     while  $right[p[x]] = x$  do  $y \leftarrow p[x]$ 
8:     if  $p[x] \neq nil$  then return ( $p[x]$ )
9:     else return ("NO SUCCESSOR") }
```

The predecessor can be found similarly with the roles of left and right exchanged and with the roles of maximum and minimum exchanged.

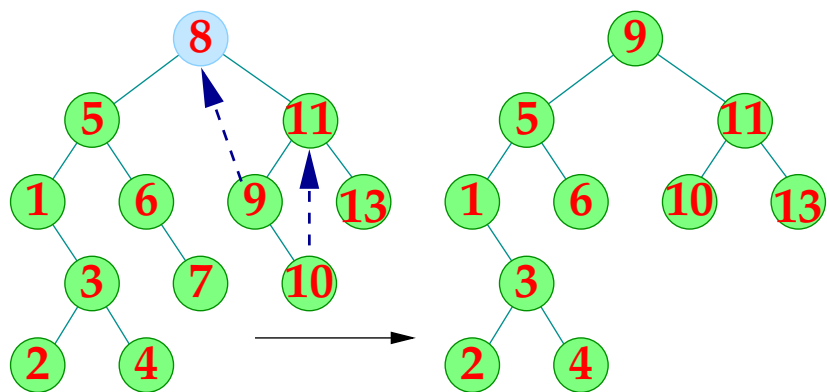
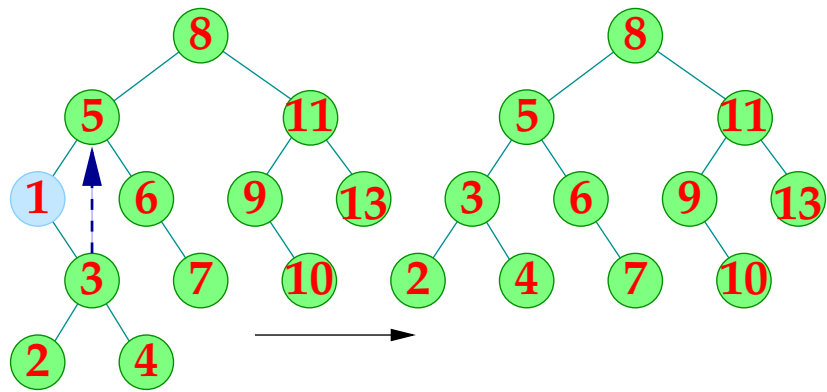
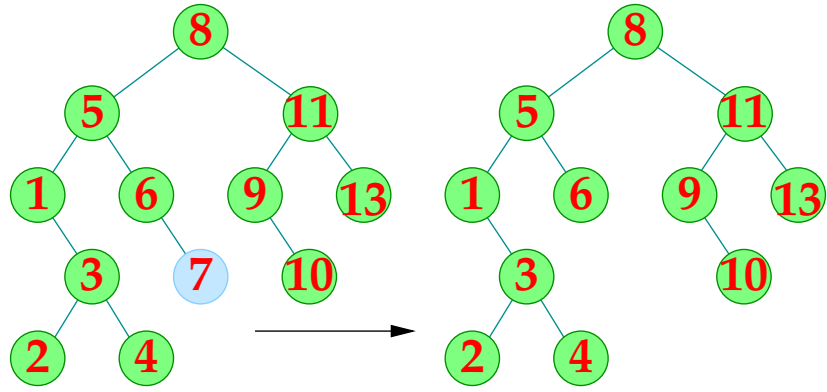
*For which node is the
successor undefined?*

*What is the running time of
the successor algorithm?*

5. Deletion

Suppose we want to delete a node z .

1. If z has no children, then we will just replace z by **nil**.
2. If z has only one child, then we will promote the unique child to z 's place.
3. If z has two children, then we will identify z 's successor. Call it y . The successor y either is a leaf or has only the right child. Promote y to z 's place. Treat the loss of y using one of the above two solutions.



Algorithm

This algorithm deletes z from BST T .

BST-Delete(T, z)

- 1: **if** $left[z] = \text{nil}$ **or** $right[z] = \text{nil}$
- 2: **then** $y \leftarrow z$
- 3: **else** $y \leftarrow \text{BST-Successor}(z)$
- 4: \triangleright y is the node that's actually removed.
- 5: \triangleright Here y does not have two children.
- 6: **if** $left[y] \neq \text{nil}$
- 7: **then** $x \leftarrow left[y]$
- 8: **else** $x \leftarrow right[y]$
- 9: \triangleright x is the node that's moving to y 's position.
- 10: **if** $x \neq \text{nil}$ **then** $p[x] \leftarrow p[y]$
- 11: \triangleright $p[x]$ is reset If x isn't NIL.
- 12: \triangleright Resetting is unnecessary if x is NIL.

Algorithm (cont'd)

```
13: if  $p[y] = \text{nil}$  then  $\text{root}[T] \leftarrow x$ 
14:  $\triangleright$  If  $y$  is the root, then  $x$  becomes the root.
15:  $\triangleright$  Otherwise, do the following.
16: else if  $y = \text{left}[p[y]]$ 
17:     then  $\text{left}[p[y]] \leftarrow x$ 
18:  $\triangleright$  If  $y$  is the left child of its parent, then
19:  $\triangleright$  Set the parent's left child to  $x$ .
20:     else  $\text{right}[p[y]] \leftarrow x$ 
21:  $\triangleright$  If  $y$  is the right child of its parent, then
22:  $\triangleright$  Set the parent's right child to  $x$ .
23: if  $y \neq z$  then
24:     {  $\text{key}[z] \leftarrow \text{key}[y]$ 
25:       Move other data from  $y$  to  $z$  }
27: return ( $y$ )
```

Summary of Efficiency Analysis

Theorem A On a binary search tree of height h , **Search**, **Minimum**, **Maximum**, **Successor**, **Predecessor**, **Insert**, and **Delete** can be made to run in $O(h)$ time.

Randomly built BST

Suppose that we insert n distinct keys into an initially empty tree. Assuming that the $n!$ permutations are equally likely to occur, what is the average height of the tree?

To study this question we consider the process of constructing a tree T by **inserting in order randomly selected n distinct keys** to an initially empty tree. Here the actual values of the keys do not matter. What matters is the position of the inserted key in the n keys.

The Process of Construction

So, we will view the process as follows:

A key x from the keys is selected uniformly at random and is inserted to the tree. Then all the other keys are inserted. Here all the keys greater than x go into the right subtree of x and all the keys smaller than x go into the left subtree. Thus, the height of the tree thus constructed is one plus the larger of the height of the left subtree and the height of the right subtree.

Random Variables

n = number of keys

X_n = height of the tree of n keys

$$Y_n = 2^{X_n}.$$

We want an upper bound on $E[Y_n]$.

For $n \geq 2$, we have

$$E[Y_n] = \frac{1}{n} \left(\sum_{i=1}^n 2E[\max\{Y_{i-1}, Y_{n-i}\}] \right).$$

$$\begin{aligned} E[\max\{Y_{i-1}, Y_{n-i}\}] &\leq E[Y_{i-1} + Y_{n-i}] \\ &\leq E[Y_{i-1}] + E[Y_{n-i}] \end{aligned}$$

Collecting terms:

$$E[Y_n] \leq \frac{4}{n} \sum_{i=1}^{n-1} E[Y_i].$$

Analysis

We claim that for all $n \geq 1$ $E[Y_n] \leq \frac{1}{4} \binom{n+3}{3}$.

We prove this by induction on n .

Base case: $E[Y_1] = 2^0 = 1$.

Induction step: We have

$$E[Y_n] \leq \frac{4}{n} \sum_{i=1}^{n-1} E[Y_i]$$

Using the fact that

$$\sum_{i=0}^{n-1} \binom{i+3}{3} = \binom{n+3}{4}$$

$$E[Y_n] \leq \frac{4}{n} \cdot \frac{1}{4} \cdot \binom{n+3}{4}$$

$$E[Y_n] \leq \frac{1}{4} \cdot \binom{n+3}{3}$$

Jensen's inequality

A function f is **convex** if for all x and y , $x < y$, and for all λ , $0 \leq \lambda \leq 1$,

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$$

Jensen's inequality states that for all random variables X and for all convex function f

$$f(E[X]) \leq E[f(X)].$$

Let this X be X_n and $f(x) = 2^x$. Then $E[f(X)] = E[Y_n]$. So, we have

$$2^{E[X_n]} \leq \frac{1}{4} \binom{n+3}{3}.$$

The right-hand side is at most $(n+3)^3$. By taking the log of both sides, we have

$$E[X_n] = O(\log n).$$

Thus the average height of a randomly build BST is $O(\log n)$.