

Search trees are data structures that support many dynamic-set operations, including SEARCH, MINIMUM, MAXIMUM, PREDECESSOR, SUCCESSOR, INSERT, and DELETE. Thus, a search tree can be used both as a dictionary and as a priority queue.

Basic operations on a binary search tree take time proportional to the height of the tree. For a complete binary tree with  $n$  nodes, such operations run in  $\Theta(\lg n)$  worst-case time. If the tree is a linear chain of  $n$  nodes, however, the same operations take  $\Theta(n)$  worst-case time. We shall see in Section 12.4 that the expected height of a randomly built binary search tree is  $O(\lg n)$ , so that basic dynamic-set operations on such a tree take  $\Theta(\lg n)$  time on average.


In practice, we can't always guarantee that binary search trees are built randomly, but there are variations of binary search trees whose worst-case performance on basic operations can be guaranteed to be good. Chapter 13 presents one such variation, red-black trees, which have height  $O(\lg n)$ . Chapter 18 introduces B-trees, which are particularly good for maintaining databases on random-access, secondary (disk) storage.

After presenting the basic properties of binary search trees, the following sections show how to walk a binary search tree to print its values in sorted order, how to search for a value in a binary search tree, how to find the minimum or maximum element, how to find the predecessor or successor of an element, and how to insert into or delete from a binary search tree. The basic mathematical properties of trees appear in Appendix B.

---

### 12.1 What is a binary search tree?

A binary search tree is organized, as the name suggests, in a binary tree, as shown in Figure 12.1. Such a tree can be represented by a linked data structure in which each node is an object. In addition to a *key* field and satellite data, each node



**Figure 12.1** Binary search trees. For any node  $x$ , the keys in the left subtree of  $x$  are at most  $key[x]$ , and the keys in the right subtree of  $x$  are at least  $key[x]$ . Different binary search trees can represent the same set of values. The worst-case running time for most search-tree operations is proportional to the height of the tree. (a) A binary search tree on 6 nodes with height 2. (b) A less efficient binary search tree with height 4 that contains the same keys.

contains fields *left*, *right*, and *p* that point to the nodes corresponding to its left child, its right child, and its parent, respectively. If a child or the parent is missing, the appropriate field contains the value *NIL*. The root node is the only node in the tree whose parent field is *NIL*.

The keys in a binary search tree are always stored in such a way as to satisfy the **binary-search-tree property**:

Let  $x$  be a node in a binary search tree. If  $y$  is a node in the left subtree of  $x$ , then  $key[y] \leq key[x]$ . If  $y$  is a node in the right subtree of  $x$ , then  $key[x] \leq key[y]$ .

Thus, in Figure 12.1(a), the key of the root is 5, the keys 2, 3, and 5 in its left subtree are no larger than 5, and the keys 7 and 8 in its right subtree are no smaller than 5. The same property holds for every node in the tree. For example, the key 3 in Figure 12.1(a) is no smaller than the key 2 in its left subtree and no larger than the key 5 in its right subtree.

The binary-search-tree property allows us to print out all the keys in a binary search tree in sorted order by a simple recursive algorithm, called an **inorder tree walk**. This algorithm is so named because the key of the root of a subtree is printed between the values in its left subtree and those in its right subtree. (Similarly, a **preorder tree walk** prints the root before the values in either subtree, and a **postorder tree walk** prints the root after the values in its subtrees.) To use the following procedure to print all the elements in a binary search tree  $T$ , we call `INORDER-TREE-WALK(root[ $T$ ])`.

INORDER-TREE-WALK( $x$ )

```

1  if  $x \neq \text{NIL}$ 
2    then INORDER-TREE-WALK( $\text{left}[x]$ )
3         print  $\text{key}[x]$ 
4         INORDER-TREE-WALK( $\text{right}[x]$ )

```

As an example, the inorder tree walk prints the keys in each of the two binary search trees from Figure 12.1 in the order 2, 3, 5, 5, 7, 8. The correctness of the algorithm follows by induction directly from the binary-search-tree property.

It takes  $\Theta(n)$  time to walk an  $n$ -node binary search tree, since after the initial call, the procedure is called recursively exactly twice for each node in the tree—once for its left child and once for its right child. The following theorem gives a more formal proof that it takes linear time to perform an inorder tree walk.

**Theorem 12.1**

If  $x$  is the root of an  $n$ -node subtree, then the call INORDER-TREE-WALK( $x$ ) takes  $\Theta(n)$  time.

**Proof** Let  $T(n)$  denote the time taken by INORDER-TREE-WALK when it is called on the root of an  $n$ -node subtree. INORDER-TREE-WALK takes a small, constant amount of time on an empty subtree (for the test  $x \neq \text{NIL}$ ), and so  $T(0) = c$  for some positive constant  $c$ .

For  $n > 0$ , suppose that INORDER-TREE-WALK is called on a node  $x$  whose left subtree has  $k$  nodes and whose right subtree has  $n - k - 1$  nodes. The time to perform INORDER-TREE-WALK( $x$ ) is  $T(n) = T(k) + T(n - k - 1) + d$  for some positive constant  $d$  that reflects the time to execute INORDER-TREE-WALK( $x$ ), exclusive of the time spent in recursive calls.

We use the substitution method to show that  $T(n) = \Theta(n)$  by proving that  $T(n) = (c + d)n + c$ . For  $n = 0$ , we have  $(c + d) \cdot 0 + c = c = T(0)$ . For  $n > 0$ , we have

$$\begin{aligned}
 T(n) &= T(k) + T(n - k - 1) + d \\
 &= ((c + d)k + c) + ((c + d)(n - k - 1) + c) + d \\
 &= (c + d)n + c - (c + d) + c + d \\
 &= (c + d)n + c,
 \end{aligned}$$

which completes the proof. ■

## Exercises

### 12.1-1

For the set of keys  $\{1, 4, 5, 10, 16, 17, 21\}$ , draw binary search trees of height 2, 3, 4, 5, and 6.

### 12.1-2

What is the difference between the binary-search-tree property and the min-heap property (see page 129)? Can the min-heap property be used to print out the keys of an  $n$ -node tree in sorted order in  $O(n)$  time? Explain how or why not.

### 12.1-3

Give a nonrecursive algorithm that performs an inorder tree walk. (*Hint:* There is an easy solution that uses a stack as an auxiliary data structure and a more complicated but elegant solution that uses no stack but assumes that two pointers can be tested for equality.)

### 12.1-4

Give recursive algorithms that perform preorder and postorder tree walks in  $\Theta(n)$  time on a tree of  $n$  nodes.

### 12.1-5

Argue that since sorting  $n$  elements takes  $\Omega(n \lg n)$  time in the worst case in the comparison model, any comparison-based algorithm for constructing a binary search tree from an arbitrary list of  $n$  elements takes  $\Omega(n \lg n)$  time in the worst case.


---

## 12.2 Querying a binary search tree

A common operation performed on a binary search tree is searching for a key stored in the tree. Besides the SEARCH operation, binary search trees can support such queries as MINIMUM, MAXIMUM, SUCCESSOR, and PREDECESSOR. In this section, we shall examine these operations and show that each can be supported in time  $O(h)$  on a binary search tree of height  $h$ .

### Searching

We use the following procedure to search for a node with a given key in a binary search tree. Given a pointer to the root of the tree and a key  $k$ , TREE-SEARCH returns a pointer to a node with key  $k$  if one exists; otherwise, it returns NIL.



**Figure 12.2** Queries on a binary search tree. To search for the key 13 in the tree, we follow the path  $15 \rightarrow 6 \rightarrow 7 \rightarrow 13$  from the root. The minimum key in the tree is 2, which can be found by following *left* pointers from the root. The maximum key 20 is found by following *right* pointers from the root. The successor of the node with key 15 is the node with key 17, since it is the minimum key in the right subtree of 15. The node with key 13 has no right subtree, and thus its successor is its lowest ancestor whose left child is also an ancestor. In this case, the node with key 15 is its successor.

TREE-SEARCH( $x, k$ )

```

1  if  $x = \text{NIL}$  or  $k = \text{key}[x]$ 
2    then return  $x$ 
3  if  $k < \text{key}[x]$ 
4    then return TREE-SEARCH( $\text{left}[x], k$ )
5  else return TREE-SEARCH( $\text{right}[x], k$ )

```

The procedure begins its search at the root and traces a path downward in the tree, as shown in Figure 12.2. For each node  $x$  it encounters, it compares the key  $k$  with  $\text{key}[x]$ . If the two keys are equal, the search terminates. If  $k$  is smaller than  $\text{key}[x]$ , the search continues in the left subtree of  $x$ , since the binary-search-tree property implies that  $k$  could not be stored in the right subtree. Symmetrically, if  $k$  is larger than  $\text{key}[x]$ , the search continues in the right subtree. The nodes encountered during the recursion form a path downward from the root of the tree, and thus the running time of TREE-SEARCH is  $O(h)$ , where  $h$  is the height of the tree.

The same procedure can be written iteratively by “unrolling” the recursion into a **while** loop. On most computers, this version is more efficient.

ITERATIVE-TREE-SEARCH( $x, k$ )

```

1  while  $x \neq \text{NIL}$  and  $k \neq \text{key}[x]$ 
2    do if  $k < \text{key}[x]$ 
3        then  $x \leftarrow \text{left}[x]$ 
4        else  $x \leftarrow \text{right}[x]$ 
5  return  $x$ 

```

### Minimum and maximum

An element in a binary search tree whose key is a minimum can always be found by following *left* child pointers from the root until a NIL is encountered, as shown in Figure 12.2. The following procedure returns a pointer to the minimum element in the subtree rooted at a given node  $x$ .

TREE-MINIMUM( $x$ )

```
1  while  $left[x] \neq \text{NIL}$ 
2      do  $x \leftarrow left[x]$ 
3  return  $x$ 
```

The binary-search-tree property guarantees that TREE-MINIMUM is correct. If a node  $x$  has no left subtree, then since every key in the right subtree of  $x$  is at least as large as  $key[x]$ , the minimum key in the subtree rooted at  $x$  is  $key[x]$ . If node  $x$  has a left subtree, then since no key in the right subtree is smaller than  $key[x]$  and every key in the left subtree is not larger than  $key[x]$ , the minimum key in the subtree rooted at  $x$  can be found in the subtree rooted at  $left[x]$ .

The pseudocode for TREE-MAXIMUM is symmetric.

TREE-MAXIMUM( $x$ )

```
1  while  $right[x] \neq \text{NIL}$ 
2      do  $x \leftarrow right[x]$ 
3  return  $x$ 
```

Both of these procedures run in  $O(h)$  time on a tree of height  $h$  since, as in TREE-SEARCH, the sequence of nodes encountered forms a path downward from the root.

### Successor and predecessor

Given a node in a binary search tree, it is sometimes important to be able to find its successor in the sorted order determined by an inorder tree walk. If all keys are distinct, the successor of a node  $x$  is the node with the smallest key greater than  $key[x]$ . The structure of a binary search tree allows us to determine the successor of a node without ever comparing keys. The following procedure returns the successor of a node  $x$  in a binary search tree if it exists, and NIL if  $x$  has the largest key in the tree.

TREE-SUCCESSOR( $x$ )

```

1  if  $right[x] \neq \text{NIL}$ 
2      then return TREE-MINIMUM( $right[x]$ )
3   $y \leftarrow p[x]$ 
4  while  $y \neq \text{NIL}$  and  $x = right[y]$ 
5      do  $x \leftarrow y$ 
6       $y \leftarrow p[y]$ 
7  return  $y$ 

```

The code for TREE-SUCCESSOR is broken into two cases. If the right subtree of node  $x$  is nonempty, then the successor of  $x$  is just the leftmost node in the right subtree, which is found in line 2 by calling TREE-MINIMUM( $right[x]$ ). For example, the successor of the node with key 15 in Figure 12.2 is the node with key 17.

On the other hand, as Exercise 12.2-6 asks you to show, if the right subtree of node  $x$  is empty and  $x$  has a successor  $y$ , then  $y$  is the lowest ancestor of  $x$  whose left child is also an ancestor of  $x$ . In Figure 12.2, the successor of the node with key 13 is the node with key 15. To find  $y$ , we simply go up the tree from  $x$  until we encounter a node that is the left child of its parent; this is accomplished by lines 3–7 of TREE-SUCCESSOR.

The running time of TREE-SUCCESSOR on a tree of height  $h$  is  $O(h)$ , since we either follow a path up the tree or follow a path down the tree. The procedure TREE-PREDECESSOR, which is symmetric to TREE-SUCCESSOR, also runs in time  $O(h)$ .

Even if keys are not distinct, we define the successor and predecessor of any node  $x$  as the node returned by calls made to TREE-SUCCESSOR( $x$ ) and TREE-PREDECESSOR( $x$ ), respectively.

In summary, we have proved the following theorem.

### Theorem 12.2

The dynamic-set operations SEARCH, MINIMUM, MAXIMUM, SUCCESSOR, and PREDECESSOR can be made to run in  $O(h)$  time on a binary search tree of height  $h$ . ■

### Exercises

#### 12.2-1

Suppose that we have numbers between 1 and 1000 in a binary search tree and want to search for the number 363. Which of the following sequences could *not* be the sequence of nodes examined?

- a. 2, 252, 401, 398, 330, 344, 397, 363.
- b. 924, 220, 911, 244, 898, 258, 362, 363.

- c. 925, 202, 911, 240, 912, 245, 363.
- d. 2, 399, 387, 219, 266, 382, 381, 278, 363.
- e. 935, 278, 347, 621, 299, 392, 358, 363.

**12.2-2**

Write recursive versions of the TREE-MINIMUM and TREE-MAXIMUM procedures.

**12.2-3**

Write the TREE-PREDECESSOR procedure.

**12.2-4**

Professor Bunyan thinks he has discovered a remarkable property of binary search trees. Suppose that the search for key  $k$  in a binary search tree ends up in a leaf. Consider three sets:  $A$ , the keys to the left of the search path;  $B$ , the keys on the search path; and  $C$ , the keys to the right of the search path. Professor Bunyan claims that any three keys  $a \in A$ ,  $b \in B$ , and  $c \in C$  must satisfy  $a \leq b \leq c$ . Give a smallest possible counterexample to the professor's claim.

**12.2-5**

Show that if a node in a binary search tree has two children, then its successor has no left child and its predecessor has no right child.

**12.2-6**

Consider a binary search tree  $T$  whose keys are distinct. Show that if the right subtree of a node  $x$  in  $T$  is empty and  $x$  has a successor  $y$ , then  $y$  is the lowest ancestor of  $x$  whose left child is also an ancestor of  $x$ . (Recall that every node is its own ancestor.)

**12.2-7**

An inorder tree walk of an  $n$ -node binary search tree can be implemented by finding the minimum element in the tree with TREE-MINIMUM and then making  $n-1$  calls to TREE-SUCCESSOR. Prove that this algorithm runs in  $\Theta(n)$  time.

**12.2-8**

Prove that no matter what node we start at in a height- $h$  binary search tree,  $k$  successive calls to TREE-SUCCESSOR take  $O(k + h)$  time.

**12.2-9**

Let  $T$  be a binary search tree whose keys are distinct, let  $x$  be a leaf node, and let  $y$  be its parent. Show that  $\text{key}[y]$  is either the smallest key in  $T$  larger than  $\text{key}[x]$  or the largest key in  $T$  smaller than  $\text{key}[x]$ .