

Binary Trees (Ch. 12)

Binary trees are an efficient way of storing data so that searches can be $O(\lg n)$.

Used when we need a data structure that supports dynamic set operations, e.g., for tree S and key x (searching is primary purpose of BST)

- o INSERT(S, x)
- o Search(S, k)
- o MINIMUM(S), MAXIMUM(S)
- o SUCCESSOR(S, x), PREDESSOR(S, x)
- o Preorder, Inorder and Postorder traversal help to evaluate expressions

Binary Search Trees

Requirements for Binary Search Tree (BST):

1. Must be a binary tree.
2. All keys must be unique (key values are like individual ID numbers).
3. Each node in tree is the root of a BST such that:
 - All nodes to left of root will have keys $<$ root.key, and
 - all nodes to right will have keys $>$ root.key.

Main advantage of BST is **rapid search** and low memory use; memory use is dependent only on size of data set.

Time of search = $O(\text{depth of BST})$ = maximum number of nodes on root to leaf path.

Binary Search Trees

Every binary tree node (internal or leaf) contains a key. These keys are unique.

Binary Search Tree Property: For every node x in tree,

- o $y.\text{key} < x.\text{key}$ for every y in $x.\text{left}$ (left subtree of x)
- o $y.\text{key} > x.\text{key}$ for every y in $x.\text{right}$ (right subtree of x)

BST has stronger requirements than heaps do.

These dynamic set operations are supported on BST S and key x

- o INSERT(S, x), DELETE(S, x)
- o SEARCH(S, x), MINIMUM(S), MAXIMUM(S)
- o SUCCESSOR(S, x), PREDESSOR(S, x)
- o InorderTraversal(S, x) (to list or print sorted set)

Binary Search Trees

Binary search trees, like heaps, can do most operations quickly because the operations depend on the height (depth) of the tree.

Unlike heaps, binary search trees keep all the data in semi-sorted order such that the cost to print all the data in sorted order is linear in the number of items in the tree. For heapsort, this cost is $O(n \lg n)$.

We would like all the dynamic set functions mentioned on the last slide to be $O(\lg n)$. However, this good running time depends on the bst implementation.

Binary Search Trees

A BST can be implemented as either a hierarchical list or as a sequential array.

The hierarchical list representation is better in terms of storage required, because only the amount of storage needed is used. Most algorithms are written for a hierarchical tree with fields for key, right, and left subtrees (and a parent pointer if needed).


The sequential array implementation has better performance when the BST is complete...otherwise there are holes in the array = wasted memory.

left subtree of $A[i] = A[2i]$ right subtree $A[i] = A[2i + 1]$

parent of $A[i]$ is at $A[\text{floor of } i/2]$


Structure of bst nodes

Each bst node contains fields *left*, *right*, and *key*. Each bst node is the root of a binary search tree. Assume all keys are unique and all leaves are NIL (singly-linked bst)



Structure of bst nodes

A doubly-linked bst node contains fields *left*, *right*, *parent* and *key*. Each bst node is the root of a binary search tree.



Binary Search Trees

The BST's total ordering does the heap's partial ordering one better; not only is there a relationship between a BST node and its children, but there is also a relationship between the children, i.e. the value of a node's left child is always less than the value of its right child.

BST Insert

```

Insert(T, z)
1.  y = NIL
2.  x = T.root
3.  while x ≠ NIL
4.    y = x
5.    if z.key < x.key
6.      x = x.left
7.    else
8.      x = x.right
9.  z.parent == y // x = NIL & parent of z is set to y
10. if y == NIL // z is first node in T, which was previously empty
11.   T.root = z
12. else if z.key < y.key
13.   y.left = z // z is added to tree as a leaf, the left child of y
14. else
15.   y.right = z // z is added to tree as a leaf, the right child of y

```

Input is a BST T and a node z such that z.left = z.right = z.parent = NIL

All leaves in T are NIL

Every node is inserted as a leaf

BST Search

```

Iterative-Tree-Search(x, k)
1.  while (x != NIL) and (k != x.key)
2.    if k < x.key
3.      x = x.left
4.    else
5.      x = x.right
6.  return x

```

The iterative version is more efficient, in terms of space used, on most computers.

Both have running times of $O(h)$, where h is the height of the tree.

```

Recursive-Tree-Search(x, k)
1.  if (x == NIL) or (k == x.key) // base cases
2.    return x
3.  if (k < x.key) // recursive case 1: search left
4.    return Recursive-Tree-Search (x.left, k)
5.  else // recursive case 2: search right
6.    return Recursive-Tree-Search (x.right, k)

```

BST Min & Max

The minimum element in a BST can always be found by following left child pointers to a leaf (until a NIL left child pointer is encountered). Likewise, the maximum element can be found by following right child pointers to a leaf.

```

Tree-Minimum(x)
while x.left ≠ NIL
  x = x.left
return x

```

```

Tree-Maximum(x)
while x.right ≠ NIL
  x = x.right
return x

```

Both have running times of $O(h)$, where h is the height of the tree.

BST Inorder Successor

In a BST, the Inorder Successor of x can also be defined as the node with the smallest key greater than the key x . This algorithm is used when deleting a node from a BST.


1. If x has a right child, then x .successor is the smallest node in the subtree rooted at x .right.
2. If x has no right child, then x .successor is the nearest ancestor of x whose left child is either an ancestor of x or x itself.

```

Tree-Successor(x)
1.  if x.right != NIL
2.    return Tree-Minimum(x.right)
3.  temp = x.parent
4.  while temp != NIL and x == temp.right
5.    x = temp
6.  temp = temp.parent
7.  return temp

```

BST Inorder Successor




Case where x.right is not equal to NIL

```
Tree-Successor(x)
1. if x.right != NIL
2.   return Tree-Minimum(x.right)
```

Tree-Successor(x)

```
1. if x.right != NIL
2.   return Tree-Minimum(x.right)
3. temp = x.parent
4. while temp != NIL and x == temp.right
5.   x = temp
6.   temp = temp.parent
7. return temp
```


Cases where x.right = NIL



BST Inorder Predecessor

```
Tree-Predecessor(x)
1. if x.left != NIL then
2.   then return Tree-Maximum(x.left)
3. temp = x.parent
4. while temp != NIL and x == temp.left
5.   x = temp
6.   temp = temp.parent
7. return temp
```


- o If x has a leftchild, then x.predecessor is the largest node in the subtree rooted at x.left.
- o If x has no leftchild, then x.predecessor is the nearest ancestor of x whose right child is either an ancestor of x, or x itself



Cases where x.left = NIL

Case where x.left not equal to NIL

```
Tree-Predecessor(x)
1. if x.left != NIL then
2.   then return Tree-Maximum(x.left)
3. temp = x.parent
4. while temp != NIL and x == temp.left
5.   x = temp
6.   temp = temp.parent
7. return temp
```



Delete(T, z)

```
1. if z.left == NIL or z.right == NIL
2.   y = z
3. else
4.   y = Tree-Successor(z)
5. if y.left != NIL
6.   x = y.left
7. else
8.   x = y.right
9. if x != NIL
10.  x.parent = y.parent
11. if y.parent == NIL
12.  T.root = x
13. else if y == y.parent.left
14.  y.parent.left = x
15. else
16.  y.parent.right = x
17. if y != z
18.  swap key[z] and key[y]
19. return y
```

BST Delete

Input: BST T and node z to be deleted. Three cases:

- 1) z has no children. Just remove it.
- 2) z has only one child. Splice out z, by letting z's child replace z.

Delete(T, z)

```
1. if z.left == NIL or z.right == NIL
2.   y = z
3. else
4.   y = Tree-Successor(z)
5. if y.left != NIL
6.   x = y.left
7. else
8.   x = y.right
9. if x != NIL
10.  x.parent = y.parent
11. if y.parent == NIL
12.  T.root = x
13. else if y == y.parent.left
14.  y.parent.left = x
15. else
16.  y.parent.right = x
17. if y != z
18.  swap key[z] and key[y]
19. return y
```

BST Delete

- 3) z has two children. Find z's successor y, which has at most one child.

Since y is z's successor, then y can have no left child, but it may have a right child.

If y is z's right child, then replace z by y, leaving y's right child as is.

If y is in z's right subtree, but is not z's right child, first replace y by its own right child and replace z by y.

```

Transplant(T, u, v)
1. if u.parent == NIL
2.   t.root = v
3. else if u == u.parent.left
4.   u.parent.left = v
5. else
6.   u.parent.right = v
7. if v != NIL
8.   v.p = u.p

```

BST Transplant

Replaces one subtree (rooted at u) with another subtree (rooted at v).

When Transplant replaces the subtree rooted at node u with the subtree rooted at node v, node u's parent becomes node v's parent and node u's parent sets node v as its appropriate child.

This procedure simplifies the code for deleting a node from a BST.

BST Tree-Delete with Transplant

Tree-Delete(T, z)

```

1. if z.left == NIL           ;; Case 1 or 2
2.   Transplant(T, z, z.right)
3. else if z.right == NIL    ;; Case 2
4.   Transplant(T, z, z.left)
5. else                       ;; Case 3
6.   y = Tree-Minimum(z.right)
7.   if y.p != z
8.     Transplant(T, y, y.right)
9.     y.right = z.right
10.    y.right.p = y
11.   Transplant(T, z, y)
12.   y.left = z.left
13.   y.left.p = y

```

Case 1: z is a leaf
Case 2: z has one child
Case 3: z has two children

```

Transplant(T, u, v)
1. if u.parent == NIL
2.   t.root = v
3. else if u == u.parent.left
4.   u.parent.left = v
5. else
6.   u.parent.right = v
7. if v != NIL
8.   v.p = u.p

```

Postorder Traversal


Postorder traversal is a recursive algorithm used to produce a postfix expression from an expression tree. All leaf nodes have NIL left and right children.

Postorder-Tree-Walk(x) /** start at root x **/

1. if x != NIL
2. Postorder-Tree-Walk(x.left)
3. Postorder-Tree-Walk(x.right)
4. visit the root

Running time = $\theta(n)$ (each node must be visited at least once)

Postorder traversal example



Iterative postorder traversal: (S1 and S2 are stacks)
1. Push root on S1.
2. Loop while S1 is not empty
2.1 U = Pop node from S1 and push it on S2
2.2 Push non-NIL L and R children of U on S1
3. Pop and print contents of S2 until it is empty

Post-traversal using 2 stacks:

1. push 1 on S1
2. pop 1 from S1 and push it on S2
3. push 2 and 3 on S1
4. pop 3 from S1 and push it on S2
5. push 6 and 7 on S1
6. pop 7 from S1 and push it on S2

7. pop 6 from S1 and push it on S2
 8. pop 2 from S1 and push it on S2
 9. push 4 and 5 on S1
 10. pop 5 from S1 and push it on S2
 11. pop 4 from S1 and push it on S2
- Algorithm done because S1 is empty
Pop then Print each number of S2.

Preorder Traversal

Preorder traversal is a recursive algorithm that is used to get a prefix expression from an expression tree. All leaf nodes have NIL left and right children.

Preorder-Tree-Walk(x) /** start at root x **/

1. if x != NIL
2. visit the root
3. Preorder-Tree-Walk(x.left)
4. Preorder-Tree-Walk(x.right)

Running time = $\theta(n)$ (each node must be visited at least once)

Inorder Traversal

Inorder-Tree-Walk(x) /** start at root x **/


1. if x != NIL
2. Inorder-Tree-Walk(x.left)
3. visit the root
4. Inorder-Tree-Walk(x.right)

In a BST, inorder-Tree-Walk(root) visits the keys in ascending order.

Running time = $\theta(n)$ (each node must be visited at least once)

Expression Trees

A binary expression tree is a specific kind of a binary tree used to represent arithmetic expressions. An inorder traversal yields an infix arithmetic expression.



Minimizing Running Time

Problem: worst case for binary search tree height is $\Theta(n)$ - no better than a linked list.

Solution: Guarantee tree has small height by making sure it is balanced so that $h = O(\lg n)$.

Method: restructure the tree if necessary. No extra work for searching, but requires extra work when inserting or deleting.

Red-black, AVL and 2-3 trees: special cases of binary trees that avoid the worst-case behavior by ensuring that the tree is nearly balanced at all times.