Tirgul 7

Binary Search Trees

DAST 2005

Motivation

- We would like to have a dynamic ADT that efficiently supports the following common operations:
 - Insert & Delete
 - · Search for an element
 - Minimum & Maximum
 - Predecessor & Successor
- Use a binary tree! All operations take Θ(h)
- When the tree is balanced, *h*=log(*n*)

DAST 2005

Binary search tree

- A binary search tree has a root, internal nodes with at most two children each, and leaf nodes
- Each node x has *left*(x), *right*(x), *parent*(x), and *key*(x) fields (and possibly other fields as well).

Binary-search-tree properties:

- Let *x* be the root of a sub-tree, and *y* a node below it. *y* resides in the:
 - left sub-tree if key(y) < key(x)
 - right sub-tree if key(y) > key(x)







Tree traversal

<u>Inorder-Tree-Walk</u>(x)

```
if x ≠ null
then Inorder-Tree-Walk(left[x])
print key[x]
Inorder-Tree-Walk(right[x])
```

- Recurrence equation:
 - $T(0) = \Theta(1)$
 - $T(n)=T(k) + T(n-k-1) + \Theta(1)$
- **Complexity:** $\Theta(n)$

DAST 2005





- The successor of *x* is the <u>smallest element</u> *y* with a key greater than that of *x*
- The successor of *x* can be found without comparing the keys. It is either:
 - 1. *null* if *x* is the *maximum node*.
 - 2. the *minimum* of the right child of t when t has a right child.
 - 3. or else, the lowest ancestor of x whose <u>left child</u> is also an ancestor of x.



Tree-Successor routine

Tree-Successor(x)
if right[x] ≠ null // Case 2
then return Tree-Minimum(right[x])
y ← parent[x] // Case 3
while y ≠ null and x = right[y] do
x ← y
y ← parent[y]
return y

DAST 2005

Insert

- · Insert is very similar to search:
- We search for the value, if we do not find it, we continue searching along either the left or right branch.
- Eventually we will reach a null leaf, and simply add the value at that position.
- The complexity is proportional to the height of the tree





Delete

- Delete is more complicated than insert. There are three cases to delete node *z*:
 - 1. z has no children
 - 2. z has one child
 - 3. z has two children
- Case 1: delete z and update the child's parent child to null.
- Case 2: delete z and connect its parent to its child.
- Case 3: more complex; we can't just take the node out and reconnect its parent with its children, because the tree will no longer be a binary tree!







- For case 3, the solution is to replace the node by its successor (or predecessor), and "pull" the successor, which necessarily has one child at most.
- Claim: if a node has two children, its successor has at most one child.
- <u>Proof</u>: This is because if the node has two children, its successor is the minimum of its right sub-tree. This minimum cannot have a left child because then the child would be the minimum...
- <u>Invariant</u>: in all cases the binary search tree property is preserved after the deletion.











Complexity analysis

- <u>Delete</u>: The two first cases take *O*(1) operations: they involve switching the pointers of the parent and the child (if it exists) of the node that is deleted.
- The third case requires a call to Tree-Successor, and thus can take *O*(*h*) time.
- In conclusion: all dynamic operations on a binary search tree take *O*(*h*), where *h* is the height of the tree.
- In the worst case, the height of the tree can be O(*n*)