

Trees



Searching

Suppose we want to search for things in a list

One possibility is to keep the items in a 'randomly' ordered list, so insertion is $O(1)$, but then a search takes $O(n)$ time

Or, we could keep them in a sorted list, in which case we can use a binary search which takes $O(\log(n))$ time, but then new items would have to be added in the middle, which takes $O(n)$ time

When there is a mixture of search and insert operations, and both operations need to be well below $O(n)$, then the items can usefully be stored in an ordered binary tree

A tree is created out of cells, with each cell having two pointers

We will create ordered binary trees, without worrying about how well balanced the tree is

Balancing techniques include:

- Reorder the input data (assuming few updates)
- Occasionally re-construct the tree
- Randomise the data (treap)
- Use a self-balancing tree (red-black, AVL, 2-3, ...)

Tree structure

Here's a struct for holding one node in a tree of ints:

```
struct node {  
    struct node *left;  
    int key;  
    struct node *right;  
};  
typedef struct node node;
```

This is essentially the same as

$\text{Tree } a = \text{Tip} \mid \text{Node } (\text{Tree } a) \ a \ (\text{Tree } a)$
in Haskell (using `NULL` for `Tip`)

New node

Here's a function to create a new node (a one-element tree):

```
node *new_node(int n) {  
    node *p = malloc(sizeof(node));  
    *p = (node) { NULL, n, NULL };  
    return p;  
}
```

Recursive Insertion

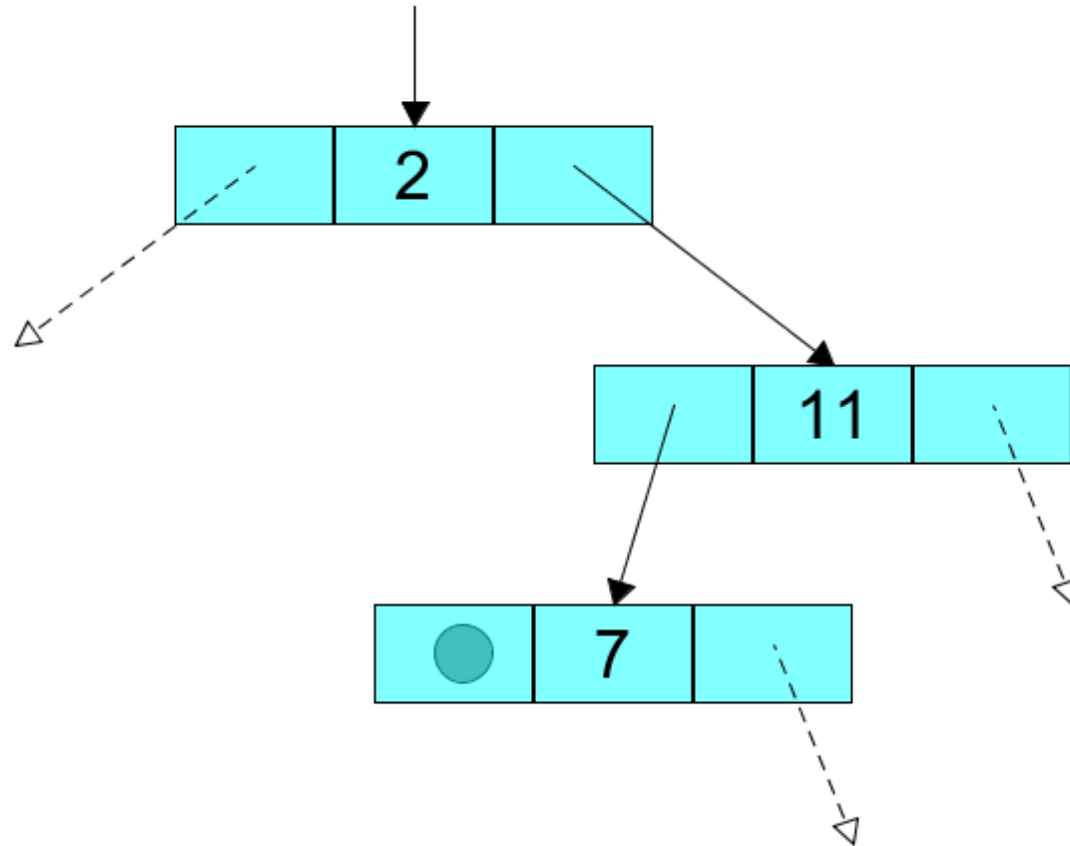
Here's a recursive insertion function:

```
node *insert_node(node *p, int n) {
    if (p == NULL) p = new_node(n);
    else if (n < p->key)
        p->left = insert_node(p->left, n);
    else if (n > p->key)
        p->right = insert_node(p->right, n);
    return p;
}
```

It uses `p` as a current-node variable

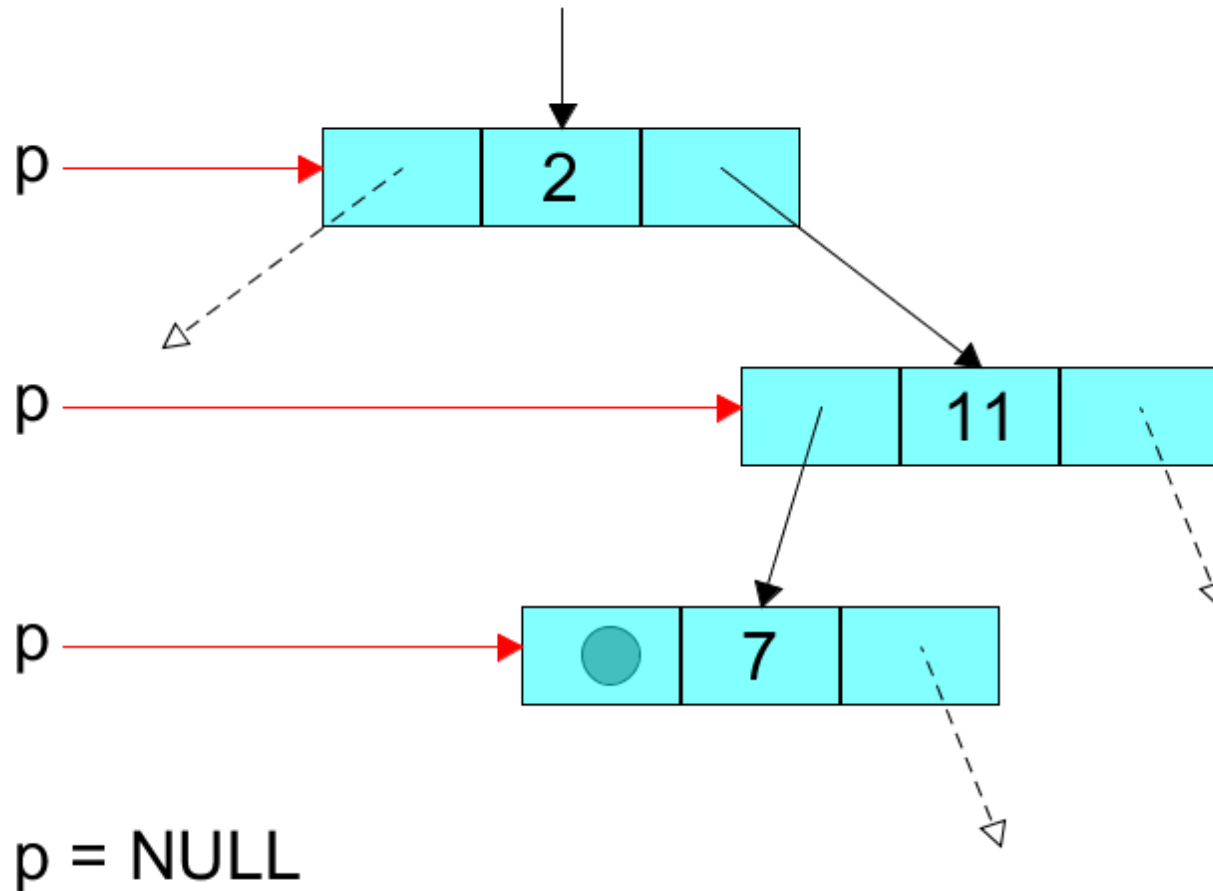
When you call it, it returns a possibly updated node, which you have to put back where you got it

Visualise tree



Pointers are shown pointing to the 'middle' of nodes, but that's only for symmetry

Visualise insert



Inserting 5, pointer p points to nodes, moves down the nodes, ends as **NULL**

Alternative Recursive Insertion

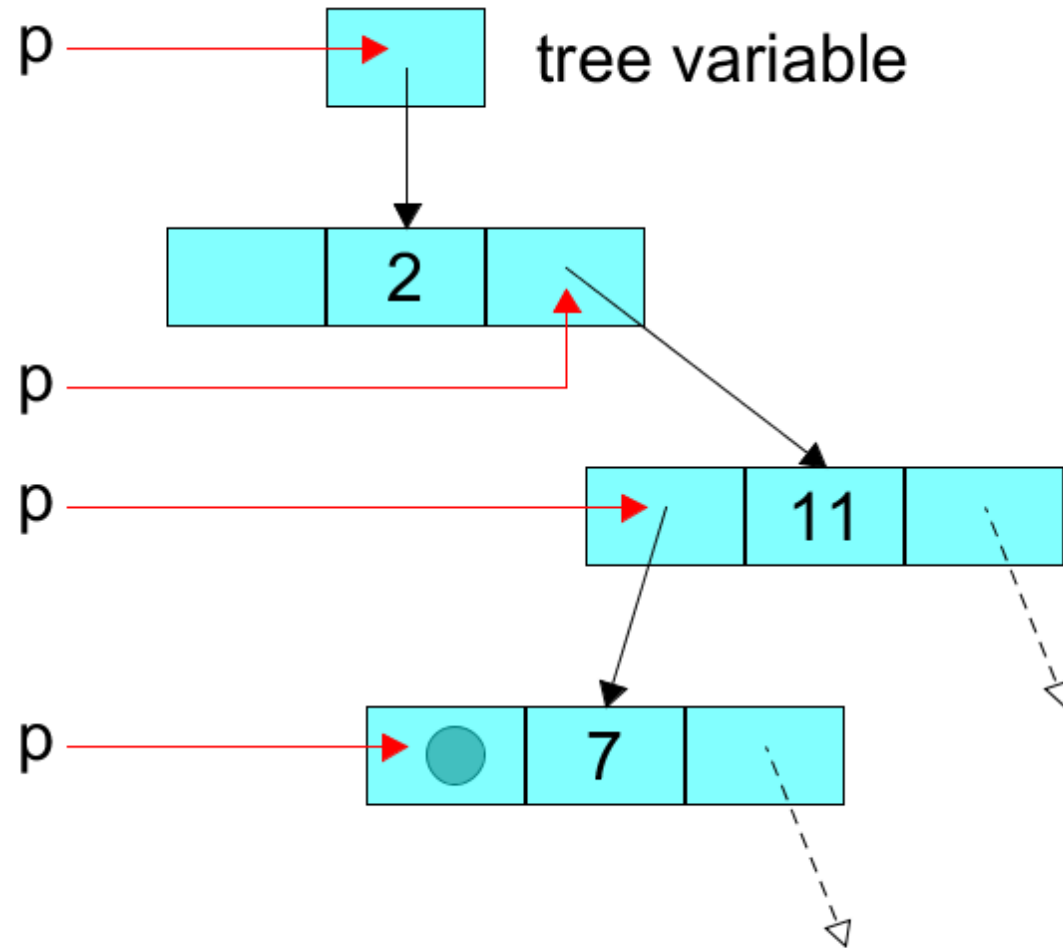
Here's a version which doesn't return anything, but uses a pointer to a pointer:

```
void insert_node(node **p, int n) {  
    if (*p == NULL) *p = new_node(n);  
    else if (n < (*p)->key)  
        insert_node(&(*p)->left, n);  
    else if (n > (*p)->key)  
        insert_node(&(*p)->right, n);  
}
```

It updates in place, and only does it once

When we reach the right place, we have a *pointer* to NULL, so we can replace it

Visualise



Now **p** points to a pointer, ends pointing to a box containing **NULL**, and the box can be updated

Iterative Insertion

Here's a complicated iterative version:

```
node *insert_node(node *p, int n) {
    bool done = false;
    if (p == NULL) { p = new_node(n); done = true; }
    while (! done) {
        if (n == p->key) done = true;
        else if (n < p->key) {
            if (p->left != NULL) p = p->left;
            else { p->left = new_node(n); done = true; }
        }
        else {
            if (p->right != NULL) p = p->right;
            else { p->right = new_node(n); done = true; }
        }
    }
    return p;
}
```

Alternative Iterative Insertion

The structure is simpler using a pointer to a pointer:

```
void insert_node(node **p, int n) {
    bool done = false;
    while (!done) {
        if (*p == NULL) {
            *p = new_node(n); done = true;
        }
        else if (n == (*p)->key) done = true;
        else if (n < (*p)->key) p = &(*p)->left;
        else p = &(*p)->right;
    }
}
```

Should you use recursive or iterative insertion, and should you use pointers-to-pointers or not?

You can disregard what anybody says about efficiency - what matters is complexity (and balance)

Use whichever you like - but when you write functions which use *both* left and right subtrees instead of just one, recursion stays simple while iteration gets nastier, and pointers-to-pointers are complex

So most programmers use recursion, and not pointers-to-pointers

A wrapper

Functions on trees are inconvenient if we force callers to use the nodes directly, either they have to catch the output (or pass a pointer-to-a-pointer)

So we need a wrapping structure for a tree:

```
struct tree {  
    struct node *root;  
};  
typedef struct tree tree;
```

This can also be a useful place to put global information about the tree

New tree

Here's a reasonable function to create a new tree:

```
tree *new_tree() {  
    tree *t = malloc(sizeof(tree));  
    t->root = NULL;  
    return t;  
}
```

A wrapped insertion function is:

```
void insert(tree *t, int n) {  
    t->root = insert_node(t->root, n);  
}
```

The `insert_node` function is the pointer-to-node recursive version, but the user can't tell which version we are using

Recursive searching

Searching is a bit simpler, and can also be done recursively or iteratively - here's a recursive version:

```
node *find_node(node *p, int n) {  
    if (p == NULL) { }  
    else if (n < p->key)  
        p = find_node(p->left, n);  
    else if (n > p->key)  
        p = find_node(p->right, n);  
    return p;  
}
```

Iterative searching

Here's an iterative version:

```
node *find_node(node *p, int n) {
    bool done = false;
    while (!done) {
        if (p == NULL) done = true;
        else if (n == p->key) done = true;
        else if (n < p->key) p = p->left;
        else p = p->right;
    }
    return p;
}
```

Wrapping

Again, we want a wrapper function

It shouldn't export any nodes to the user, it should just return (e.g.) a boolean to say whether the number is in the tree or not

```
bool contains(tree *t, int n) {  
    return find_node(t->root, n) != NULL;  
}
```

A *map* is a structure which maps keys to values

For example, when counting words, you might want to map word strings as keys, to integer counts as values:

```
struct node {  
    struct node *left;  
    char word[20];  
    int count;  
    struct node *right;  
};
```

The tree would be structured according to the words, and functions would retrieve or update the counts


Self balancing trees

There are many types of self-balancing tree, with red-black trees being the most popular in libraries because you only need one extra bit per node

The different types (AVL trees, 2-3 trees, ...) all use the same *mechanism* but have different *policies*

Rotation

The mechanism used for balancing is rotation:



In both cases, $A < P < B < Q < C$

A rotate function

Here's a function to rotate right:

```
node *rotate_right(node *q) {  
    node *p = q->left;  
    q->left = p->right;  
    p->right = q;  
    return p;  
}
```

A *policy* is an algorithm which decides what rotations to do and when, according to some extra info in each node, and which guarantees $O(\log(n))$ depth