

Today's reading: Goodrich & Tamassia, Section 10.3.

SPLAY TREES
=====

A splay tree is a type of balanced binary search tree. Structurally, it is identical to an ordinary binary search tree; the only difference is in the algorithms for finding, inserting, and deleting entries.

All splay tree operations run in $O(\log n)$ time on average, where n is the number of entries in the tree. Any single operation can take $\Theta(n)$ time in the worst case. But any sequence of k splay tree operations, with the tree initially empty and never exceeding n items, takes $O(k \log n)$ worst-case time.

Although 2-3-4 trees make a stronger guarantee (every operation on a 2-3-4 tree takes $O(\log n)$ time), splay trees have several advantages. Splay trees are simpler and easier to program. Because of their simplicity, splay tree insertions and deletions are typically faster in practice (sometimes by a constant factor, sometimes asymptotically). Find operations can be faster or slower, depending on circumstances.

Splay trees are designed to give especially fast access to entries that have been accessed recently, so they really excel in applications where a small fraction of the entries are the targets of most of the find operations.

Splay trees have become the most widely used basic data structure invented in the last 30 years, because they're the fastest type of balanced search tree for many applications.

Tree Rotations

Like many types of balanced search trees, splay trees are kept balanced with the help of structural changes called rotations. There are two types--a left rotation and a right rotation--and each is the other's reverse. Suppose that X and Y are binary tree nodes, and A , B , and C are subtrees. A rotation transforms either of the configurations illustrated above to the other. Observe that the binary search tree invariant is preserved: keys in A are less than or equal to X ; keys in C are greater than or equal to Y ; and keys in B are $\geq X$ and $\leq Y$.

Rotations are also used in AVL trees and red-black trees, which are discussed by Goodrich and Tamassia, but are not covered in this course.

Unlike 2-3-4 trees, splay trees are not kept perfectly balanced, but they tend to stay reasonably well-balanced most of the time, thereby averaging $O(\log n)$ time per operation in the worst case (and sometimes achieving $O(1)$ average running time in special cases).

Splay Tree Operations

```
[1] Entry find(Object k);
```


The `find()` operation in a splay tree begins just like the `find()` operation in an ordinary binary search tree: we walk down the tree until we find the entry with key k , or reach a dead end (a node from which the next logical step leads to a null pointer).

However, a splay tree isn't finished its job. Let X be the node where the search ended, whether it contains the key k or not. We splay X up the tree through a sequence of rotations, so that X becomes the root of the tree. Why? One reason is so that recently accessed entries are near the root of the tree, and if we access the same few entries repeatedly, accesses will be very fast. Another reason is because if X lies deeply down an unbalanced branch of the tree, the splay operation will improve the balance along that branch.

When we splay a node to the root of the tree, there are three cases that determine the rotations we use.

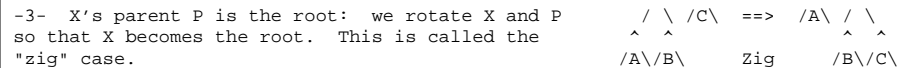


The mirror image of this case--where X is a left child and P is a right child--uses the same rotations in mirror image: rotate X and P right, then X and G left. Both the case illustrated above and its mirror image are called the "zig-zag" case.



The mirror image of this case--where X and P are both right children--uses the same rotations in mirror image: rotate G and P left, then P and X left. Both the case illustrated above and its mirror image are called the "zig-zig" case.

We repeatedly apply zig-zag and zig-zig rotations to X ; each pair of rotations raises X two levels higher in the tree. Eventually, either X will reach the root (and we're done), or X will become the child of the root. One more case handles the latter circumstance.



Here's an example of "find(7)". Note how the tree's balance improves.

