

I thought the following four [rules] would be enough, provided that I made a firm and constant resolution not to fail even once in the observance of them. The first was never to accept anything as true if I had not evident knowledge of its being so. . . . The second, to divide each problem I examined into as many parts as was feasible, and as was requisite for its better solution. The third, to direct my thoughts in an orderly way. . . . establishing an order in thought even when the objects had no natural priority one to another. And the last, to make throughout such complete enumerations and such general surveys that I might be sure of leaving nothing out.

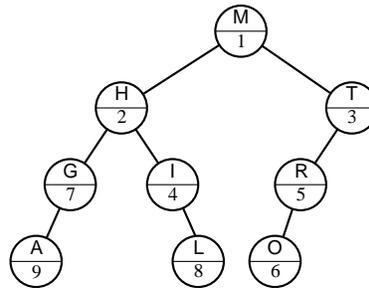
— René Descartes, *Discours de la Méthode* (1637)

9 Treaps and Skip Lists (October 13 and 18)

In this lecture, we consider two randomized alternatives to balanced binary search structures such as AVL trees, red-black trees, B-trees, or splay trees.

9.1 Treaps

A *treap* is a binary tree in which every node has both a *search key* and a *priority*, where the inorder sequence of search keys is sorted and each node's priority is smaller than the priorities of its children.¹ In other words, a treap is simultaneously a binary search tree for the search keys and a (min-)heap for the priorities. In our examples, we will use letters for the search keys and numbers for the priorities.



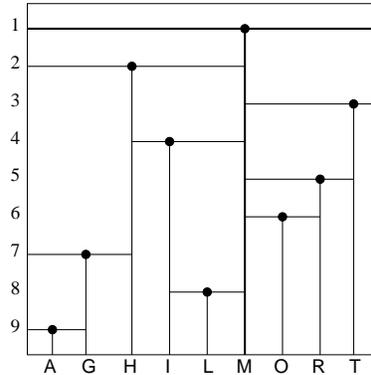
A treap. The top half of each node shows its search key and the bottom half shows its priority.

I'll assume from now on that all the keys and priorities are distinct. Under this assumption, we can easily prove by induction that the structure of a treap is completely determined by the search keys and priorities of its nodes. Since it's a heap, the node v with highest priority must be the root. Since it's also a binary search tree, any node u with $key(u) < key(v)$ must be in the left subtree, and any node w with $key(w) > key(v)$ must be in the right subtree. Finally, since the subtrees are treaps, by induction, their structures are completely determined. The base case is the trivial empty treap.

Another way to describe the structure is that a treap is exactly the binary tree that results by inserting the nodes one at a time into an initially empty tree, in order of increasing priority, using the usual insertion algorithm. This is also easy to prove by induction.

A third way interprets the keys and priorities as the coordinates of a set of points in the plane. The root corresponds to a T whose joint lies on the topmost point. The T splits the plane into three parts. The top part is (by definition) empty; the left and right parts are split recursively. This interpretation has some interesting applications in computational geometry, which (unfortunately) we probably won't have time to talk about.

¹Sometimes I hate English. Normally, 'higher priority' means 'more important', but 'first priority' is also more important than 'second priority'. Maybe 'posteriority' would be better; one student suggested 'unimportance'.



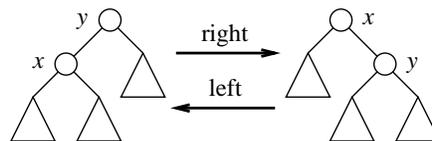
A geometric interpretation of the same treap.

Treaps were first discovered by Jean Vuillemin in 1980, but he called them *Cartesian trees*.² The word ‘treap’ was first used by Edward McCreight around 1980 to describe a slightly different data structure, but he later switched to the more prosaic name *priority search trees*.³ Treaps were rediscovered and used to build randomized search trees by Cecilia Aragon and Raimund Seidel in 1989.⁴ A different kind of randomized binary search tree, which uses random rebalancing instead of random priorities, was later discovered and analyzed by Conrado Martínez and Salvador Roura in 1996.⁵

9.1.1 Binary Search Tree Operations

The search algorithm is the usual one for binary search trees. The time for a successful search is proportional to the depth of the node. The time for an unsuccessful search is proportional to the depth of either its successor or its predecessor.

To insert a new node z , we start by using the standard binary search tree insertion algorithm to insert it at the bottom of the tree. At the point, the search keys still form a search tree, but the priorities may no longer form a heap. To fix the heap property, as long as z has smaller priority than its parent, perform a *rotation* at z , a local operation that decreases the depth of z by one and increases its parent’s depth by one, while maintaining the search tree property. Rotations can be performed in constant time, since they only involve simple pointer manipulation.

A right rotation at x and a left rotation at y are inverses.

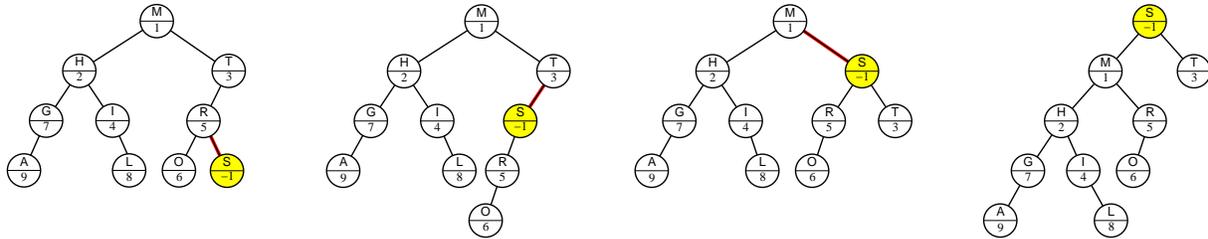
The overall time to insert z is proportional to the depth of z before the rotations—we have to walk down the treap to insert z , and then walk back up the treap doing rotations. Another way to say this is that the time to insert z is roughly twice the time to perform an unsuccessful search for $key(z)$.

²J. Vuillemin, A unifying look at data structures. *Commun. ACM* 23:229–239, 1980.

³E. M. McCreight. Priority search trees. *SIAM J. Comput.* 14(2):257–276, 1985.

⁴R. Seidel and C. R. Aragon. Randomized search trees. *Algorithmica* 16:464–497, 1996.

⁵C. Martínez and S. Roura. Randomized binary search trees. *J. ACM* 45(2):288–323, 1998. The results in this paper are virtually identical (including the constant factors!) to the corresponding results for treaps, although the analysis techniques are quite different.



Left to right: After inserting $(S, 10)$, rotate it up to fix the heap property.
 Right to left: Before deleting $(S, 10)$, rotate it down to make it a leaf.

Deleting a node is *exactly* like inserting a node, but in reverse order. Suppose we want to delete node z . As long as z is not a leaf, perform a rotation at the child of z with smaller priority. This moves z down a level and its smaller-priority child up a level. The choice of which child to rotate preserves the heap property everywhere except at z . When z becomes a leaf, chop it off.

We sometimes also want to *split* a treap T into two treaps $T_{<}$ and $T_{>}$ along some pivot key π , so that all the nodes in $T_{<}$ have keys less than π and all the nodes in $T_{>}$ have keys bigger than π . A simple way to do this is to insert a new node z with $key(z) = \pi$ and $priority(z) = -\infty$. After the insertion, the new node is the root of the treap. If we delete the root, the left and right sub-treaps are exactly the trees we want. The time to split at π is roughly twice the time to (unsuccessfully) search for π .

Similarly, we may want to *merge* two treaps $T_{<}$ and $T_{>}$, where every node in $T_{<}$ has a smaller search key than any node in $T_{>}$, into one super-treap. Merging is just splitting in reverse—create a dummy root whose left sub-treap is $T_{<}$ and whose right sub-treap is $T_{>}$, rotate the dummy node down to a leaf, and then cut it off.

9.1.2 Analysis

The cost of each of these operations is proportional to the depth $d(v)$ of some node v in the treap.

- **Search:** A successful search for key k takes $O(d(v))$ time, where v is the node with $key(v) = k$. For an unsuccessful search, let v^- be the inorder *predecessor* of k (the node whose key is just barely smaller than k), and let v^+ be the inorder *successor* of k (the node whose key is just barely larger than k). Since the last node examined by the binary search is either v^- or v^+ , the time for an unsuccessful search is either $O(d(v^+))$ or $O(d(v^-))$.
- **Insert/Delete:** Inserting a new node with key k takes either $O(d(v^+))$ time or $O(d(v^-))$ time, where v^+ and v^- are the predecessor and successor of the new node. Deletion is just insertion in reverse.
- **Split/Merge:** Splitting a treap at pivot value k takes either $O(d(v^+))$ time or $O(d(v^-))$ time, since it costs the same as inserting a new dummy root with search key k and priority $-\infty$. Merging is just splitting in reverse.

Since the depth of a node in a treap is $\Theta(n)$ in the worst case, each of these operations has a worst-case running time of $\Theta(n)$.

9.1.3 Random Priorities

A *randomized binary search tree* is a treap in which the priorities are *independently and uniformly distributed continuous random variables*. That means that whenever we insert a new search key into the treap, we generate a random real number between (say) 0 and 1 and use that number as

the priority of the new node. The only reason we're using real numbers is so that the probability of two nodes having the same priority is zero, since equal priorities make the analysis messy. In practice, we could just choose random integers from a large range, like 0 to $2^{31} - 1$, or random floating point numbers. Also, since the priorities are independent, each node is equally likely to have the smallest priority.

The cost of all the operations we discussed—search, insert, delete, split, join—is proportional to the depth of some node in the tree. Here we'll see that the *expected* depth of *any* node is $O(\log n)$, which implies that the expected running time for any of those operations is also $O(\log n)$.

Let x_k denote the node with the k th smallest search key. To analyze the expected depth, we define an indicator variable

$$A_k^i = [x_i \text{ is a proper ancestor of } x_k].$$

(The superscript doesn't mean power in this case; it just a reminder of which node is supposed to be further up in the tree.) Since the depth $d(v)$ of v is just the number of proper ancestors of v , we have the following identity:

$$d(x_k) = \sum_{i=1}^n A_k^i.$$

Now we can express the *expected* depth of a node in terms of these indicator variables as follows.

$$E[d(x_k)] = \sum_{i=1}^n \Pr[A_k^i = 1]$$

(Just as in our analysis of matching nuts and bolts in Lecture 3, we're using linearity of expectation and the fact that $E[X] = \Pr[X = 1]$ for any indicator variable X .) So to compute the expected depth of a node, we just have to compute the probability that some node is a proper ancestor of some other node.

Fortunately, we can do this easily once we prove a simple structural lemma. Let $X(i, k)$ denote either the subset of treap nodes $\{x_i, x_{i+1}, \dots, x_k\}$ or the subset $\{x_k, x_{k+1}, \dots, x_i\}$, depending on whether $i < k$ or $i > k$. $X(i, k)$ and $X(k, i)$ always denote precisely the same subset, and this subset contains $|k - i| + 1$ nodes. $X(1, n) = X(n, 1)$ contains all n nodes in the treap.

Lemma 1. *For all $i \neq k$, x_i is a proper ancestor of x_k if and only if x_i has the smallest priority among all nodes in $X(i, k)$.*

Proof: If x_i is the root, then it is an ancestor of x_k , and by definition, it has the smallest priority of *any* node in the treap, so it must have the smallest priority in $X(i, k)$.

On the other hand, if x_k is the root, then x_i is not an ancestor of x_k , and indeed x_i does not have the smallest priority in $X(i, k)$ — x_k does.

On the gripping hand⁶, suppose some other node x_j is the root. If x_i and x_k are in different subtrees, then either $i < j < k$ or $i > j > k$, so $x_j \in X(i, k)$. In this case, x_i is not an ancestor of x_k , and indeed x_i does not have the smallest priority in $X(i, k)$ — x_j does.

Finally, if x_i and x_k are in the same subtree, the lemma follows inductively (or, if you prefer, recursively), since the subtree is a smaller treap. The empty treap is the trivial base case. \square

Since each node in $X(i, k)$ is equally likely to have smallest priority, we immediately have the probability we wanted:

⁶See Larry Niven and Jerry Pournelle, *The Gripping Hand*, Pocket Books, 1994.

$$\Pr[A_k^i = 1] = \frac{[i \neq k]}{|k - i| + 1} = \begin{cases} \frac{1}{k - i + 1} & \text{if } i < k \\ 0 & \text{if } i = k \\ \frac{1}{i - k + 1} & \text{if } i > k \end{cases}$$

To compute the expected depth of a node x_k , we just plug this probability into our formula and grind through the algebra.

$$\begin{aligned} \mathbb{E}[d(x_k)] &= \sum_{i=1}^n \Pr[A_k^i = 1] = \sum_{i=1}^{k-1} \frac{1}{k - i + 1} + \sum_{i=k+1}^n \frac{1}{i - k + 1} \\ &= \sum_{j=2}^k \frac{1}{j} + \sum_{i=2}^{n-k} \frac{1}{i} \\ &= H_k - 1 + H_{n-k} - 1 \\ &< \ln k + \ln(n - k) - 2 \\ &< 2 \ln n - 2. \end{aligned}$$

In conclusion, every search, insertion, deletion, split, and merge operation in an n -node randomized binary search tree takes $O(\log n)$ expected time.

Since a treap is exactly the binary tree that results when you insert the keys in order of increasing priority, a randomized treap is the result of inserting the keys in *random* order. So our analysis also automatically gives us the expected depth of any node in a binary tree built by random insertions (without using priorities).

9.1.4 Randomized Quicksort (Again?!)

We've already seen two completely different ways of describing randomized quicksort. The first is the familiar recursive one: choose a random pivot, partition, and recurse. The second is a less familiar iterative version: repeatedly choose a new random pivot, partition whatever subset contains it, and continue. But there's a third way to describe randomized quicksort, this time in terms of binary search trees.

RANDOMIZEDQUICKSORT:
 $T \leftarrow$ an empty binary search tree
 insert the keys into T in *random order*
 output the inorder sequence of keys in T

Our treap analysis tells us that this algorithm will run in $O(n \log n)$ expected time, since each key is inserted in $O(\log n)$ expected time.

Why is this quicksort? Just like last time, all we've done is rearrange the order of the comparisons. Intuitively, the binary tree is just the recursion tree created by the normal version of quicksort. In the recursive formulation, we compare the initial pivot against everything else and then recurse. In the binary tree formulation, the first "pivot" becomes the root of the tree without any comparisons, but then later as each other key is inserted into the tree, it is compared against the root. Either way, the first pivot chosen is compared with everything else. The partition splits the remaining items into a left subarray and a right subarray; in the binary tree version, these are exactly the items that go into the left subtree and the right subtree. Since both algorithms define the same two subproblems, by induction, both algorithms perform the same comparisons.

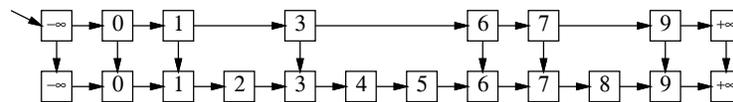
We even saw the probability $\frac{1}{|k-i|+1}$ before, when we were talking about sorting nuts and bolts with a variant of randomized quicksort. In the more familiar setting of sorting an array of numbers, the probability that randomized quicksort compares the i th largest and k th largest elements is exactly $\frac{2}{|k-i|+1}$. The binary tree version compares x_i and x_k if and only if x_i is an ancestor of x_k or vice versa, so the probabilities are exactly the same.

9.2 Skip Lists

Skip lists, which were first discovered by Bill Pugh in the late 1980's,⁷ have many of the usual desirable properties of balanced binary search trees, but their structure is completely different.

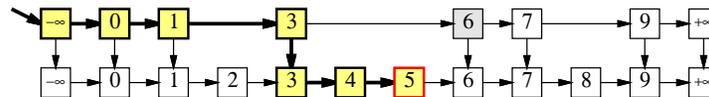
9.2.1 Shortcuts

At a high level, a skip list is just a sorted, singly linked list with some shortcuts. To do a search in a normal singly-linked list of length n , we obviously need to look at n items in the worst case. To speed up this process, we can make a second-level list that contains roughly half the items from the original list. Specifically, for each item in the original list, we duplicate it with probability $1/2$. We then string together all the duplicates into a second sorted linked list, and add a pointer from each duplicate back to its original. Just to be safe, we also add sentinel nodes at the beginning and end of both lists.



A linked list with some randomly-chosen shortcuts.

Now we can find a value x in this augmented structure using a two-stage algorithm. First, we scan for x in the shortcut list, starting at the $-\infty$ sentinel node. If we find x , we're done. Otherwise, we reach some value bigger than x and we know that x is not in the shortcut list. Let w be the largest item less than x in the shortcut list. In the second phase, we scan for x in the original list, starting from w . Again, if we reach a value bigger than x , we know that x is not in the data structure.



Searching for 5 in a list with shortcuts.

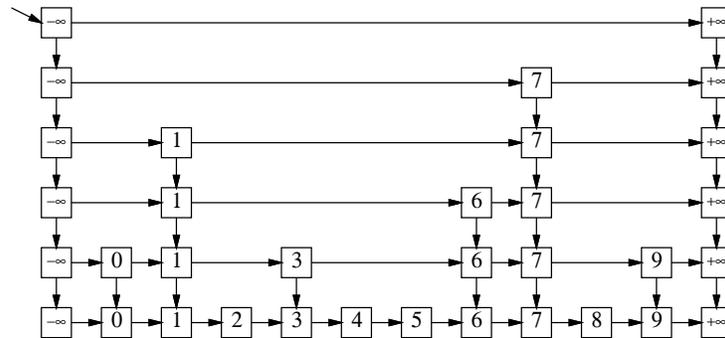
Since each node appears in the shortcut list with probability $1/2$, the expected number of nodes examined in the first phase is at most $n/2$. Only one of the nodes examined in the second phase has a duplicate. The probability that any node is followed by k nodes without duplicates is 2^{-k} , so the expected number of nodes examined in the second phase is at most $1 + \sum_{k \geq 0} 2^{-k} = 2$. Thus, by adding these random shortcuts, we've reduced the cost of a search from n to $n/2 + 2$, roughly a factor of two in savings.

9.2.2 Skip lists

Now there's an obvious improvement—add shortcuts to the shortcuts, and repeat recursively. That's exactly how skip lists are constructed. For each node in the original list, we flip a coin

⁷William Pugh. Skip lists: A probabilistic alternative to balanced trees. *Commun. ACM* 33(6):668–676, 1990.

over and over until we get tails. For each heads, we make a duplicate of the node. The duplicates are stacked up in levels, and the nodes on each level are strung together into sorted linked lists. Each node v stores a search key ($\text{key}(v)$), a pointer to its next lower copy ($\text{down}(v)$), and a pointer to the next node in its level ($\text{right}(v)$).

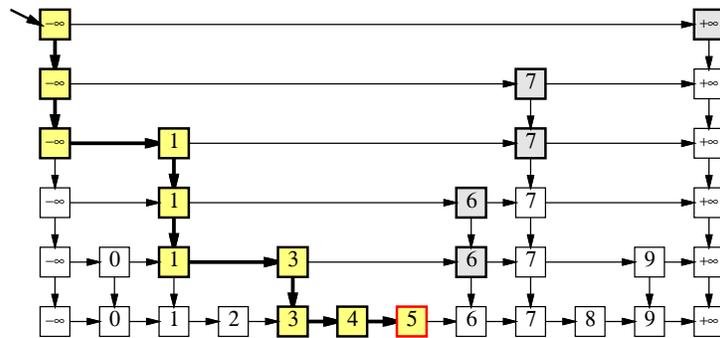


A skip list is a linked list with recursive random shortcuts.

The search algorithm for skip lists is very simple. Starting at the leftmost node L in the highest level, we scan through each level as far as we can without passing the target value x , and then proceed down to the next level. The search ends when we either reach a node with search key x or fail to find x on the lowest level.

```

SKIPLISTFIND( $x, L$ ):
   $v \leftarrow L$ 
  while ( $v \neq \text{NULL}$  and  $\text{key}(v) \neq x$ )
    if  $\text{key}(\text{right}(v)) > x$ 
       $v \leftarrow \text{down}(v)$ 
    else
       $v \leftarrow \text{right}(v)$ 
  return  $v$ 
    
```



Searching for 5 in a skip list.

Intuitively, Since each level of the skip lists has about half the number of nodes as the previous level, the total number of levels should be about $O(\log n)$. Similarly, each time we add another level of random shortcuts to the skip list, we cut the search time in half except for a constant overhead. So after $O(\log n)$ levels, we should have a search time of $O(\log n)$. Let's formalize each of these two intuitive observations.

9.2.3 Number of Levels

The actual values of the search keys don't affect the skip list analysis, so let's assume the keys are the integers 1 through n . Let $L(x)$ be the number of levels of the skip list that contain some search key x , not counting the bottom level. Each new copy of x is created with probability $1/2$ from the previous level, essentially by flipping a coin. We can compute the expected value of $L(x)$ recursively—with probability $1/2$, we flip tails and $L(x) = 0$; and with probability $1/2$, we flip heads, increase $L(x)$ by one, and recurse:

$$E[L(x)] = \frac{1}{2} \cdot 0 + \frac{1}{2}(1 + E[L(x)])$$

Solving this equation gives us $E[L(x)] = 1$.

In order to analyze the expected cost of a search, however, we need a bound on the number of levels $L = \max_x L(x)$. Unfortunately, we can't compute the average of a maximum the way we would compute the average of a sum. Instead, we will derive a stronger result, showing that the depth is $O(\log n)$ with high probability. 'High probability' is a technical term that means the probability is at least $1 - 1/n^c$ for some constant $c \geq 1$.

In order for a search key x to appear on the k th level, we must have flipped k heads in a row, so $\Pr[L(x) \geq k] = 2^{-k}$. In particular,

$$\Pr[L(x) \geq 2 \lg n] = \frac{1}{n^2}.$$

(There's nothing special about the number 2 here.) The skip list has at least $2 \lg n$ levels if and only if $L(x) \geq 2 \lg n$ for at least one of the n search keys.

$$\Pr[L \geq 2 \lg n] = \Pr[(L(1) \geq 2 \lg n) \vee (L(2) \geq 2 \lg n) \vee \dots \vee (L(n) \geq 2 \lg n)]$$

Since $\Pr[A \vee B] \leq \Pr[A] + \Pr[B]$ for any random events A and B , we can simplify this as follows:

$$\Pr[L \geq 2 \lg n] \leq \sum_{x=1}^n \Pr[L(x) \geq 2 \lg n] = \sum_{x=1}^n \frac{1}{n^2} = \frac{1}{n}.$$

So with high probability, a skip list has $O(\log n)$ levels.

9.2.4 Logarithmic Search Time

It's a little easier to analyze the cost of a search if we imagine running the algorithm backwards. `SKIPLISTFIND` takes the output from `SKIPLISTFIND` as input and traces back through the data structure to the upper left corner. Skip lists don't really have up and left pointers, but we'll pretend that they do so we don't have to write '`(v)up`' or '`(v)left`'.⁸

| |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $\begin{array}{l} \text{SKIPLISTFIND}(v): \\ \text{while } (v \neq L) \\ \quad \text{if up}(v) \text{ exists} \\ \quad \quad v \leftarrow \text{up}(v) \\ \quad \text{else} \\ \quad \quad v \leftarrow \text{left}(v) \end{array}$ |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

⁸ Leonardo da Vinci used to write everything this way, but not because he wanted to keep his discoveries secret. He just had really bad arthritis in his right hand!

Now for *every* node v in the skip list, $\text{up}(v)$ exists with probability $1/2$. So for purposes of analysis, FLIPWALK is equivalent to the following algorithm:

| |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>FLIPWALK(v): while ($v \neq L$) if COINFLIP = HEADS $v \leftarrow \text{up}(v)$ else $v \leftarrow \text{left}(v)$</pre> |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Obviously, the expected number of heads is exactly the same as the expected number of tails. Thus, the expected running time of this algorithm is twice the expected number of upward jumps. Since we already know that the number of upward jumps is $O(\log n)$ with high probability, we can conclude that the expected search time is $O(\log n)$.