








1 2-3 Trees and LLRB's

- (a) Draw what the following 2-3 tree would look like after inserting 18, 38, 12, 13, and 20.



- (b) Now, convert the resulting 2-3 tree to a left-leaning red-black tree.



- (c) If a 2-3 tree has depth H (that is, the leaves are at distance H from the root), what is the maximum number of comparisons done in the corresponding red-black tree to find whether a certain key is present in the tree?

$2H + 2$ comparisons.

The maximum number of comparisons occur from a root to leaf path with the most nodes. Because the height of the tree is H , we know that there is a path down the leaf-leaning red-black tree that consists of at most H black links, for black links in the left-leaning red-black tree are the links that add to the height of the corresponding 2-3 tree. This means that there are $H + 1$ nodes on the path from the root to the leaf, since there is one less link than nodes,

In the worst case, in the 2-3 tree representation, this path can consist entirely of nodes with two items, meaning in the left-leaning red-black tree representation, each blank link is followed by a red link. This doubles the amount of nodes on this path from the root to the leaf.

This example will represent our longest path, which is $2H + 2$ nodes long, meaning we make at most $2H + 2$ comparisons in the left-leaning red-black tree.

2 Hashing

- (a) Here are three potential implementations of the `Integer`'s `hashCode()` function. Categorize each as either a valid or an invalid hash function. If it is invalid, explain why. If it is valid, point out a flaw or disadvantage. For the 2nd implementation, note that `intValue()` will return that `Integer`'s number value as an `int`

```
public int hashCode() {
    return -1;
}
```

Valid. As required, this hash function returns the same `hashCode` for `Integers` that are `equals()` to each other. However, this is a terrible hash code because collisions are extremely frequent (collisions occur 100% of the time).

```
public int hashCode() {
    return intValue() * intValue();
}
```

Valid. Similar to (a), this hash function returns the same `hashCode` for integers that are `equals()`. However, integers that share the same absolute values will collide (for example, $x = 5$ and $x = -5$ will have the same hash code). A better hash function would be to just return the `intValue()` itself.

```
public int hashCode() {
    return super.hashCode();
}
```

Invalid. When we call `super.hashCode()` here, we will ultimately end up returning `Object.hashCode()`. This hash function returns some number corresponding to the `Integer` object's location in memory.

However, note what happens below:

```
Integer x = new Integer(5);
Integer y = new Integer(5);
```

Here, `x` and `y` are both instantiated as separate objects, meaning they'll be at two distinct memory addresses! Then `x` and `y` would receive different hashcodes according to our current `hashCode()` implementation. However, this is not desired/valid behavior, because `x` and `y` are equal according to the `.equals()` method, so their hashcodes should be the same.

- (b) For each of the following questions, answer **Always**, **Sometimes**, or **Never**.
1. If you were able to modify a key that has been inserted into a `HashMap` would you be able to retrieve that entry again later? For example, let us suppose you were mapping ID numbers to student names, and you did a `put(303, "Anjali")` operation. Now, let us suppose we somehow went to that item in our `HashMap` and manually changed the key to be 304. If we later do `get(304)`, will we be able to find and return "Anjali"? Explain.

Sometimes. If the `hashCode` for the key happens to change as a result of the modification, then we won't be able to retrieve the entry in our hashtable (unless we were to recompute which bucket the new key would belong to). Changing the key can potentially (and likely) change which bucket a key would now be indexed to.

In our example, let us suppose the `hashCode()` for a key was just its integer value, and the bucket was calculated as that number modulo the number of buckets. Let's say we had 10 buckets. In this scenario, 303 would be mapped to bucket 3, and 304 would be mapped to bucket 4. So when we `put(303, "Anjali")`, this item goes to bucket 3. Then we change the key to be 304, but don't change anything else, so this item remains in bucket 3. Now, let us suppose we later do `get(304)`. We will compute which bucket the key 304 would correspond to, which would be bucket 4. We look in bucket 4 and don't see the item there, so we cannot successfully return the item.

2. When you modify a value that has been inserted into a `HashMap` will you be able to retrieve that entry again? For example, in the above scenario, suppose we first inserted `put(303, "Anjali")` and then changed that item's value from "Anjali" to "Ajay". If we later do `get(303)`, will we be able to find and return "Ajay"? Explain.

Always. The bucket index for an entry in a `HashMap` is decided by the key, not the value. Mutating the value does not affect the lookup procedure.

In our example this would work, because when we do `get(303)` we will still go to the correct bucket.

3 A Side of Hash Browns

We want to map food items to their yumminess. We want to be able to find this information in constant time, so we've decided to use java's built-in `HashMap` class! Here, the key is an `String` representing the food item and the value is an `int` yumminess rating.

For simplicity, let's say that here a `String`'s hashcode is the first letter's position in the alphabet (`A = 0, B = 1... Z = 25`). For example, the `String` "hash brown" starts with "h", and "h" is 7th letter in the alphabet (0 indexed), so the hashcode would be 7. Note that in reality, a `String` has a much more complicated `hashCode()` implementation.

Our `HashMap` will compute the index as the key's hashcode value modulo the number of buckets in our `HashMap`. Assume the initial size is 4 buckets, and we double the size our `HashMap` as soon as the load factor reaches 3/4.


(a) Draw what our `HashMap` would look like after the following operations.

```

1  HashMap<Integer, String> hm = new HashMap<>();
2  hm.put("Hashbrowns", 7);
3  hm.put("Dim sum", 10);
4  hm.put("Escargot", 5);
5  hm.put("Brown bananas", 1);
6  hm.put("Burritos", 10);
7  hm.put("Buffalo wings", 8);
8  hm.put("Banh mi", 9);

```

Note that we resize from 4 to 8 when adding escargot (because we have 3 items and 4 buckets) and then from 8 to 16 when adding buffalo wings (because we have 6 items and 8 buckets).



(b) Do you see a potential problem here with the behavior of our `HashMap`? How could we solve this?

Here, adding a bunch of food items that start with the letter "B" result in one bucket with a lot of items. No matter how many times we resize, our current `hashCode()` will result in this problem! Imagine if we added in 100 more items that started with the letter b. Though we would resize and keep our load factor low, it wouldn't change the fact that our operations will now be slow (hint: linear time) because now we essentially have to iterate over a linked list with pretty much all the items in the `HashMap` to do things like `get("Burrito")`, for example.

A solution to this would be to have a better `hashCode()` implementation. A better implementation would distribute the Strings more randomly and evenly. While knowing how to write such a `hashCode()` is difficult and out of scope for this class, you can look at the real `hashCode()` implementation for Java Strings if you're curious!