Binary Index Trees

A cumulative frequency array allows us to calculate the sum of the range of values in O(1), ***as long as there are no changes to the data once the queries start.***

But consider situations where we might change the value at an index in an array, then query for the sum of a range of values in the array, followed by some more changes and more queries. In this sort of situation, a cumulative frequency array would still give us O(1) query times, but it would take O(n) time to update after each change!!! (Basically, if we change one index in a cumulative frequency array, all other indexes above it would have to have this value added to it as well.) Here is a quick illustration:

Current cumulative frequency array:

Index	0	1	2	3	4	5	6	7	8
Value	2	4	4	4	6	8	11	13	17

Now, consider adding the value 2 to the data, recalling that index i stores the number of values less than or equal to i. The adjusted array is:

Index	0	1	2	3	4	5	6	7	8
Value	2	4	5	5	7	9	12	14	18

We had to edit each index 2 or greater, which would take O(n) for a cumulative frequency array of size n.

Thus, we want a new arrangement where both the query AND the update are relatively fast. The creative insight here is that perhaps if we store data in a different way, perhaps we can reduce the update time by quite a bit while incurring only a modest increase in query time. A binary index tree (also called a Fenwick tree), achieves just this. It achieves a O(lg n) time for both the update and the query.

What to Store in each Index

In a cumulative frequency array, array[i] stores $\sum_{k=0}^{i} x[k]$, where x[k] represents the original items. In a binary index tree (which is really just stored in an array), each index will store a sum of values from the original array as well, but the set of values summed at each index will be more complicated than in a cumulative frequency array.

Number of Values summed at index i

The value stored in index i of a binary index tree is based upon the binary representation of i, hence the name, binary index tree.

For this description, let x[i] represent the original values.

Each index will always store a sum of a set of values of x[i] that has a size that is a perfect power of 2. In particular, the number of values added to get the value stored in index i is 2^L, where L

represents the place of the lowest 1 bit. For example if the binary representation of i is 10000100, then there are $2^2 = 4$ values summed at index i, since the least significant bit equal to one is in the 2^2 place, third from the right. If i in binary is 1010100000, then the corresponding value is the sum of 2^5 values in the original array.

Which values are summed at index i

Now that we know HOW many values of the original array are summed at index i, we must specify WHICH values comprised that sum.

The values that comprise the sum at index i are the *last* 2^L values of x[i]. Thus, we finally can write this mathematically:

BIT[i] = $\sum_{k=i-2^L+1}^{i} x[k]$, where L is the lowest one bit location as previously defined.

Here is a table of what is stored in the first 16 indexes of a binary index tree:

Index	Values Summed From Original Array
1	x[1]
2	x[1]+x[2]
3	x[3]
4	x[1]+x[2]+x[3]+x[4]
5	x[5]
6	x[5]+x[6]
7	x[7]
8	x[1]+x[2]+x[3]+x[4]+x[5]+x[6]+x[7]+x[8]
9	x[9]
10	x[9]+x[10]
11	x[11]
12	x[9]+x[10]+x[11]+x[12]
13	x[13]
14	x[13]+x[14]
15	x[15]
16	x[1]+x[2]+x[3]+…x[15]+x[16]

Indexes to Change for an Update

Consider an update to an arbitrary index x[i] of the original array. If you look at the table above, we can see that the total number of indexes that refer to x[i] is logarithmic in the size of the whole table. Consider x[5] - it's part of BIT[5], BIT[6], BIT[8] and BIT[16]. To get from 5 to 6, we add 1, the value of the lowest one bit in 5. To get from 6 to 8, we add 2, the value of the lowest one bit in 6. To get from 8 to 16, we add 8, the value of the lowest one bit of 8, and so forth.

Thus, the algorithm for adding D to the value of x[i] is as follows:

1. Let curIndex = i.
2. while curIndex < sizeof(BIT)
    2a. BIT[curIndex] += D
    2b. curIndex += valueOfLowestOneBit(curIndex)

Indexes to Add for a Query

Consider just supporting queries in the range 1 to i. (Any range query can be represented as the difference of two range queries of this time.) Let's consider summing the values x[1] + x[2] +… + x[13]. First we'd add BIT[13]. This just adds x[13]. Then we would add BIT[12], adding x[9] + x[10] + x[11] + x[12]. Finally we would add BIT[8], which would add x[1]+x[2]+x[3]+x[4]+x[5]+x[6]+x[7]+x[8].

In short, we add our current index we are at in the BIT array, and then subtract the value of the lowest one bit from the current index. Thus, the algorithm for the query sum of x[1] to x[i] is as follows:

1. Let curIndex = i.
2. Let sum = 0.
3. while curIndex > 0
    2a. sum += BIT[curIndex]
    2b. curIndex -= valueOfLowestOneBit(curIndex)

Java Implementation

The nice thing is that Java's Integer class has a method, lowestOneBit that returns the numeric value of the lowest one's bit of an integer. For example, Integer.lowestOneBit(24) returns 8, since the binary representation of 24 is 11000 and the value of the right most one is 8. So, the add method, which adds value to the given index of the original data looks like this:

```
public void add(int index, long value) {
    while (index < cumfreq.length) {
        cumfreq[index] += value;
        index += Integer.lowestOneBit(index);
    }
}
```

Here is the sum method, which returns the sum of all items of the original array upto index:

```
public long sum(int index) {
    long ans = 0;
    while (index > 0) {
        ans += cumfreq[index];
        index -= (Integer.lowestOneBit(index));
    }
    return ans;
}
```

Alternate Code for Lowest One Bit

In other languages, there might not be a method which returns the lowest one bit, but there's a nice expression using bitwise operators that equals the value that this method returns. Thus, in the code shown on the previous page, we can replace

```
Integer.lowestOneBit(index)
```

with

```
(index&(-index))
```

It's obviously the case that this expression can NOT turn on a bit that's smaller than the lowest one bit of index, since index has all 0s in its binary representation after its lowest one bit. Due to how two's complement works to store negative numbers, it's guaranteed that any bit that is on in index that is more significant than the lowest one bit of index will be off in -index. Thus, we can write the add method as follows:

```
public void add(int index, long value) {
    while (index < cumfreq.length) {
        cumfreq[index] += value;
        index += (index&(-index));
    }
}
```

Sample Problem: Counting Inversions

An inversion in an array all pairs of indexes i and j such that i < j and array[i] > array[j]. (Essentially, the pair is out of order.) Consider the problem of counting the total number of inversions in an array. Here is a sample array:

3, 9, 2, 8, 7, 1, 4, 5, 6

For each value j > 0, we want to know how many values in the array before it are greater than array[j].

For 9, there are no values before it greater than it.
For 2, there are 2 values before it (3 and 9) greater than it.
For 8, there is 1 value before it (9) greater than it.
For 7, there are 2 values before it greater than it.
For 1, there are 5 values before it greater than it.
For 4, there are 3 values before it greater than it.
For 5, there are 3 values before it greater than it.
For 6, there are 3 values before it greater than it.
This is a total of 2 + 1 + 2 + 5 + 3 + 3 + 3 = 19 inversions.

Normally, using a straight-forward algorithm, this would take $O(n^2)$ time for an array with n elements. We would like to speed up the inner for loop. Instead of running through each item array[0] through array[j-1] and seeing how many are greater than array[j], we'd like to answer this question without looking at each value.

A binary index tree helps us as follows:

1. Start with an empty binary index tree.
2. As we loop through each element in the array, query the BIT to see the sum of the BIT from array[i]+1 to MAX. This is the number of elements greater than array[i] that come before it.
3. Now, after adding that to our total number of inversions, add 1 to array[i] in the BIT, indicating that for the future, there is a number array[i] that has already been processed.

Here is a trace through of the new algorithm with the same array, included for convenience:
3, 9, 2, 8, 7, 1, 4, 5, 6

So, we start with 0 inversions.
When we process 3, there are no sum greater than 3 in our BIT.
Then we add 1 to index 3 of our BIT.
When we process 9, there is no sum greater than 9 in our BIT.
Then we add 1 to index 9 of our BIT.
When we process 2, there is a sum of 2 for values greater than 2 in our BIT.
Then we add 1 to index 2 of our BIT.
When we process 8, there is a sum of 1 for values greater than 8 in our BIT.
Then we add 1 to index 8 of our BIT.
When we process 7, there is a sum of 2 for values greater than 7 in our BIT.

Then we add 1 to index 7 of our BIT.
When we process 1, there is a sum of 5 for values greater than 1 in our BIT.
Then we add 1 to index 1 of our BIT.
When we process 4, there is a sum of 3 for values greater than 4 in our BIT.
Then we add 1 to index 4 of our BIT.
When we process 5, there is a sum of 3 for values greater than 5 in our BIT.
Then we add 1 to index 5 of our BIT.
When we process 6, there is a sum of 3 for values greater than 6 in our BIT.
Then we add 1 to index 6 of our BIT.

A Kattis problem that can be solved by counting inversions in an array quickly is Bread Sorting:

https://open.kattis.com/problems/bread

See if you can solve it!

Sample Problem: Rotating Cards

In this problem, we have cards in a stack labeled 1 through n and must remove the cards with the labels 1, then 2, then 3, etc. To remove a card successfully, we must bring it to the top of the stack. The only modifications we can make to the stack are taking the card on the top and putting it on the bottom or taking the card on the bottom and putting it on the top. The cost of moving a card is the number on its label. The goal is to determine the minimum cost of removing all of the cards. Consider the following initial stack:

3
1
4
2
5

For this stack, it's clear that we want to first move 3 from top to bottom (cost 3), then 1 is at the top so we can remove it. More generally, for every move, we have two choices: remove cards from the top until our next card reaches the top, OR move cards from the bottom to the top until our card goes from bottom to top. The better choice is simply whichever direction has a sum of cards that is less. This is where the binary index tree comes in!

We can place these values in a BIT (3 at index 1, 1 at index 2, etc.). We can use the BIT to quickly give us both sums. (Alternatively, we can use the BIT to give us one sum and use subtraction from the whole to figure out the other sum.)

We'll choose the direction that has a smaller sum, add this to our running minimum result, and then 0 out the BIT at that location, also updating our current position to that location:

3
0 (* store in variable as current location)
4
2
5

We can see that 2 is stored in index 4 (store this in a reverse list) and the sum of moving cards from top to bottom would only be 4, but the sum of moving cards from bottom to top would be 10 (3 + 5 + 2). So, after our next move we have:

3
0
4
0 (* store in variable as current location)
5

We continue in this fashion until all of the cards are removed.

In the posted implementation, the binary index tree is twice the necessary size and the numbers are all stored in the even indexes of the BIT. The odd indexes are used for the queries so that there are no issues with off by 1 errors. This is a classic technique to reduce bugs with a constant overhead in cost (twice as much memory, just a bit more time, pun intended!)

See if you can rewrite the posted code to work for a BIT indexed from 1 to n!

Sample Problem: ZigZag2

A zigzag sequence is one that goes up and down. For example, 3, 19, 17, 14 is a zigzag sequence as is 17, 13, 16, 1, 2, 1, 5, 4. Given an arbitrary sequence of length up to 10^5, determine the number of zigzag subsequences contained in the sequence of length 3 or greater. All of the values in the sequence are in between -10^9 and 10^9.

Consider just solving the problem for all zigzag sequences, regardless of their length. (Afterwards, we can figure out how to subtract out all the sequences we over-counted of lengths 1 and 2.)

First, note that we can compress the input data by noting how many unique items (call this m) there are and reassigning each one a number from 0 to m-1. For example, if the original sequence was 18, 3, -5, 14, 9, 3, 14, then the mapping would be f(-5) = 0, f(3) = 1, f(9) = 2 f(14) = 3 and f(18) = 4. Our mapped sequence would then be 4, 1, 0, 3, 2, 1, 3. The number of zigzag subsequences in this sequence is equal to the number of zigzag subsequences in the original, since we haven't changed any relative values.

Let this mapped input sequence be a[0], a[1], … a[n-1], with all values in the range from 0 to m-1. (Note that $m \leq n$.) Let numUpSeq[i] store the number of zigzag subsequences that end at the

value i that are currently going up and let numDownSeq[i] store the number of zigzag sequences that end at the value i, currently going down.

Consider we arrive at some number, x after processing some part of the input. We want to know how many sequences we can tack x onto. Well, there are two options:

1) Tacking x onto a sequence that is currently heading up that ends at x+1 or higher.
2) Tacking x onto a sequence that is currently heading down that ends at x-1 or lower.

The new sequences in the first set are built by adding numSeqUp[x+1] + numSeqUp[x+2]+…numSeqUp[m-1], a query perfect for a binary index tree. These new sequences are going down and should be added to numSeqDown[x].

The new sequences in the second set are built by adding numSeqDown[0]+numSeqDown[1]+…+numSeqDown[x-1]. These new sequences are going up and should be added numSeqUp[x]. At the end of running the Binary Index Trees for numSeqUp and numSeqDown, the sum of ALL the values in both trees represents all the unique zigzag subsequences we saw.

To solve the given problem though, we have to subtract all subsequences of lengths 1 or 2. Sequences of length 1 are easy: there is one of these for each number for both binary index trees. For sequences of length 2, we need keep a third binary index tree running just of each item, one at a time. When I get to an item x, I just want to know how many items prior to x were not equal to x (so the sum of the items above x and below x). All of these form valid zigzag subsequences of length 2. Here is the critical code to solve the problem, after the input value have been remapped (stored in nums):

```
long sub = 0;
bit up = new bit(index, MOD);
bit down = new bit(index, MOD);
bit seqBit = new bit(index, MOD);

for (int i=0; i<n; i++) {
    long prevUp = up.above(nums[i]);
    long prevDown = down.below(nums[i]);
    down.add(nums[i], prevUp+1);
    up.add(nums[i], prevDown+1);
    sub = (sub+2+ seqBit.above(nums[i]) + seqBit.below(nums[i]))%MOD;
    seqBit.add(nums[i], 1);
}

long res = (up.all() + down.all() - sub + MOD)%MOD;
```

Note: the methods above and below return the sum of items above and below the input value to the methods, respectively.