

Recursion Problem Solutions: Group B

1. Given a number n , generate all distinct ways to write n as the sum of positive integers. For example, with $n = 4$, the options are 4, 3 + 1, 2 + 2, 2 + 1 + 1, and 1 + 1 + 1 + 1.

The main challenge in solving this problem is figuring out how to avoid getting duplicate solutions like 3 + 1 and 1 + 3. The approach I've outlined in this solution works by always generating the numbers in nonincreasing order, so we will never generate solutions like 1 + 3.

Here's a solution that works by strengthening the recursion. We will actually answer the harder question "list all distinct ways to write n as the sum of positive integers, where all the integers in the sum are less than or equal to k ." Given this, we can make the following insight:

- There is just one way to write 0 as the sum of integers up to and including k : namely, it's the empty sum of no numbers.
- Otherwise, the sum must have at least one integer in it. For each possible first integer i (which must be in the range of 1 to the minimum of n and k), try writing the number as the sum of i , plus all possible ways of summing to $n - i$ using numbers no greater than i .

This gives the following recursive solution:

```
function allPartitionsOf(n) {
    return restrictedAllPartitions(n, n);
}
function restrictedAllPartitions(n, k) {
    if n is zero, the only solution is the empty sum.
    for i from 1 to the minimum of n and k, inclusive:
        record all solutions of the form i + restrictedAllPartitions(n - i, i)
    return those solutions
}
```

This solution will recompute many subproblems, so you could consider using memoization or DP to try to avoid the recomputations. This unfortunately increases the memory usage, so it's up to you to decide whether it's a good idea or not.

2. In a binary tree, a *common value subtree* is a complete subtree where every node has the same value. (A complete subtree is a subtree consisting of a node and all its children). Determine the largest common value subtree in a nonempty binary tree.

For notational simplicity, let's call a common value subtree a CVS. The key observation necessary here is the following:

- A tree with just one node is a CVS whose value is the value in the root.
- A tree with root r and subtrees t_1, t_2, \dots, t_n is a CVS if each of t_1, t_2, \dots, t_n is a CVS and has the same value as the root's value.

This means that we can do a bottom-up pass over the tree and find all of the CVS's as we go. We can then determine which CVS found this way is the largest. The pseudocode below splits this into several passes over the tree just for simplicity, and can be significantly space-optimized by combining everything together into one pass.

```
function findLargestCVS(t) {
    annotateTreeSizes(t)
    annotateTreeCVS(t)
    return largestCVSIn(t)
}
function annotateTreeSizes(t) {
    if t is a leaf, set t.size = 1
    else:
        call annotateTreeSizes(c) for each child c of t.
        set t.size = 1 + sum(c.size) for each child c of t
}
function annotateCVS(t) {
    if t is a leaf, set t.isCVS to true.
    else:
        call annotateCVS(c) for each child c of t.
        set t.isCVS to whether c.value = t.value and c.isCVS for each child c of t.
}
function largestCVSIn(t) {
    if t.isCVS return t.size
    else
        let  $d_k = \text{largestCVSIn}(c_k)$  for each child  $c_k$  of t.
        return  $\text{argmax}\{d_k.\text{size}\}$  for all  $d_k$ .
}
```

This code runs in time $O(n)$, where n is the number of nodes in the trees. To see this, note that every function call takes time $O(1 + \text{num children})$, so summing up across all nodes in the tree we can charge $O(1)$ total work to each node in the tree. Summing up across all n nodes gives a runtime of $O(n)$.

3. Suppose you have a multiway tree where each node has an associated integer value. Find a set of nodes with the maximum possible sum, subject to the constraint that you cannot choose a node and any of its children at the same time.

The insight necessary to solve this problem is to split the problem into two cases – first, solving this problem when the root node *is* included in the solution, and second when the root node is *not* included in the solution. If we include the root node, then all of its children must not be included, and we'd like optimal solutions for each of its subtrees subject to the restriction that their root nodes aren't included. If we exclude the root node, then for each child, we should take the best solution possible, whether or not we choose to include the root node.

If we code this up using a naïve recursion, we get this solution:

```
function maxSum(t) {
    return maxSumUnrestricted(t)
}

// Gives the largest possible sum that can be made with the tree rooted at t when
// there are no restrictions on whether t must be included.
function maxSumUnrestricted(t) {
    return the max of maxSumRestricted(t, true) and maxSumRestricted(t, false);
}

// Gives the largest possible sum that can be made with the tree rooted at t
// subject to the restriction that the root of t either must be or must not be
// included
function maxSumRestricted(t, include) {
    if t is null, return 0.

    if mustInclude is true:
        return t.value + sum of maxSumRestricted(c, false) for all children c of t.

    if mustInclude is false:
        return the sum of maxSumUnrestricted(c) for all children c of t
}
```

This recursion is correct but highly inefficient – we'll end up recomputing the same subproblems on many subtrees, leading to a very slow runtime. However, there are only $2n$ possible unique calls in a tree of n nodes – for each node, we can either include it or exclude it – and so we can memoize the results or use dynamic programming to compute the results bottom-up. If we do this, the net runtime is only $O(n)$ because we do $O(1)$ work per node $O(n)$ total times, though it now requires $O(n)$ storage space.

4. Suppose that you have a group of people that you need to assign into different houses. For each house, you know the number of people that the house can hold. Additionally, you know that some people *insist* that they not be put into the same house as some other people. Given the list of pairs of people that can't be put into houses and the house capacities, determine how to distribute the people into the houses, or report that it's impossible.

This problem is actually equivalent to the graph coloring problem – you can imagine that you have a node for each person, an edge between two people if they insist on not being in the same house, and one color for each house. Since this problem is, in general, NP-hard, it's probably best to approach this through some kind of backtracking approach. One natural one is the following:

- If all people are assigned, you're done. Report success.
- Otherwise, choose a person.
- Determine what possible houses they can be in by checking where their neighbors have been assigned.
- If there are no options, backtrack.
- Otherwise, for each choice, try that choice. If any of them work, report success; otherwise report failure.