

CS 6363.005.19S Lecture 7–February 5, 2019

Main topics are `#dynamic_programming` with `#example/fibonacci_numbers` and `#example/rod_cutting`.

Prelude

- Homework 1 is due today! Please submit it ASAP to eLearning.
- Homework 2 will be released shortly to be due Tuesday, February 19th.
- Midterm 1 will be held in class on Tuesday, February 26th. If you have a conflict, please let me know ASAP. I'm doing this as a closed book, no notes exam or calculators exam, because the QE is also closed book. It will contain a combination of questions designed to check your understanding of the past lectures as well as one or two algorithm design questions.

Fibonacci Numbers

- By now, I assume you're all familiar with Fibonacci numbers.
- These were described by Leonardo Pisano around the 12th century, but the Fibonacci recurrence was originally discovered by Indian scholar Virahanka 500 years earlier during his study of classical Sanskrit poetry.
- In this class, Fibonacci numbers are defined using the following recurrence.
 - $F_0 = 0$
 - $F_1 = 1$
 - $F_n = F_{n-1} + F_{n-2}$

Recursion is Sometimes Slow

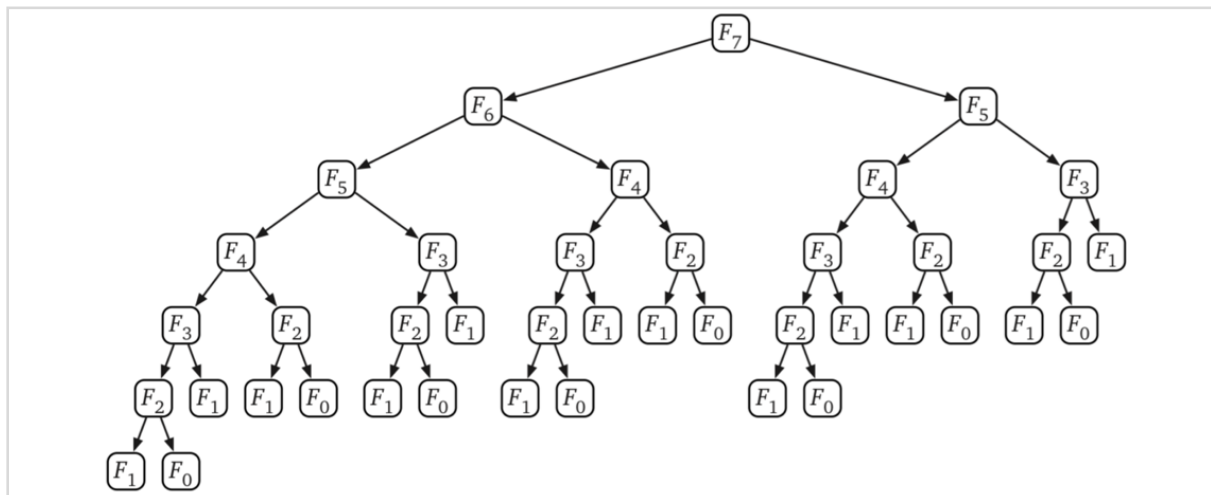
- As we've seen in the homework, Fibonacci numbers describe many different phenomena, and it would be good if there was some way to calculate them.
- Fibonacci numbers are defined recursively, so we can design a recursive algorithm pretty easily:

```
REC FIBO(n):  
  if n = 0  
    return 0  
  else if n = 1  
    return n  
  else  
    return REC FIBO(n - 1) + REC FIBO(n - 2)
```

- Unfortunately, this algorithm is horribly slow.
- Let $T(n)$ be the number of subproblems for computing F_n assuming arithmetic operations takes constant time. Outside the recursive calls, we do only a constant number of steps, so

we can write $T(n) = T(n - 1) + T(n - 2) + 1$ with $T(0) = 1$ and $T(1) = 1$.

- It's almost the same recurrence once again! And if we wrote out the first few terms we could correctly guess $T(n) = 2F_{n+1} - 1$.
- But that means computing F_n takes twice as long as counting to it.
- Using techniques that some of you may have seen in discrete math, we can figure out $F_n = \Theta(\phi^n)$ where $\phi = (\sqrt{5} + 1) / 2 \sim 1.61803$, the *golden ratio*. This algorithm is exponential in n !
- But is there an intuitive explanation for why is it so slow? Let's look at a recursion tree. I'll mark which number we're trying to compute.



- All these numbers come from adding up values computed in the recursive calls, so there must be F_n leaves with value 1. These are the calls to `RecFibo(1)`. There's at most that many calls to `RecFibo(0)`, so $\Theta(F_n)$ leaves total. It's a full binary tree, so there must be $\Theta(F_n)$ nodes. That's a lot of recursive calls!

Memoization

- Of course, even in this small example, you can tell we're wasting a lot of time recomputing the same numbers over and over again. You can even prove via induction that `RecFibo(n - k)` is called $F_{k - 1}$ times.
- Now one thing we could do is try to remember which values we've already computed by storing them in a global array. This is called *memoization*.

```

MEMFIBO(n):
  if n = 0
    return 0
  else if n = 1
    return 1
  else
    if F[n] is undefined
      F[n] ← MEMFIBO(n - 1) + MEMFIBO(n - 2)
    return F[n]

```

- If we look through the recursive calls of `MemFibo`, we see that the array `F[]` is filled from

the bottom up. First $F[2]$, then $F[3]$, and so on. In a sense, we've significantly trimmed the recursion tree!

- Outside recursive calls, it only takes $O(1)$ time to evaluate each F_i , and we evaluate them once, so the total number of additions is only $O(n)$. That's a *lot* faster.

Filling Deliberately and Dynamic Programming

- But now that we see how the table is filled, why don't we just do so deliberately to make the algorithm even simpler. We'll just use a simple iterative algorithm that fills the array entries one by one.

```
ITERFIBO(n):  
  F[0] ← 0  
  F[1] ← 1  
  for i ← 2 to n  
    F[i] ← F[i - 1] + F[i - 2]  
  return F[n]
```

- Not only does this algorithm avoid recursion and probably perform better in practice, but it's a *lot* easier to analyze. There's a single for loop over n entries, so it performs only $O(n)$ additions. We can easily tell that it stores $O(n)$ integers as well.
- This is an example of a *dynamic programming* algorithm, formalized and popularized by Richard Bellman in the mid-1950s, although others including Virahanka and Fibonacci already applied it to this example centuries earlier.
- Bellman claimed the term dynamic programming was used to hide the fact he was doing mathematics research from his industry-minded military bosses, although there are reasons to doubt that story. The word programming here doesn't refer to writing code. It's being used in the same way that television and radio stations plan or schedule their "programs", typically by filling a table. They were originally trying to optimize for certain time-varying processes, so he used the term dynamic.
- Dynamic programming is now a standard tool for multistage planning in a variety of areas, and it's one of the most useful tools we have for designing algorithms.

Saving Space

- As I'll explain later, our job with dynamic programming was (ultimately) to recognize how to fill a table with Fibonacci numbers, but keeping a whole table around when all you need is a single answer is somewhat wasteful.
- Since we only refer back to the last two entries in each iteration of the for loop, we can easily save some space in this instance.

ITERFIBO2(n):

```
prev ← 1
curr ← 0
for i ← 1 to n
    next ← curr + prev
    prev ← curr
    curr ← next
return curr
```

- Now we're only storing a constant number of integers.
- I *usually* won't ask you about space usage in this class, but I may for the next couple weeks since it's easy enough to figure out and you can sometimes make easy improvements.

Rod Cutting

- This dynamic programming technique of filling out a table of recursive subproblem solutions or recurrence relation evaluations is typically applied to optimization problems.
- An *optimization problem* is one where there may be many valid or feasible solutions, and each solution has a numerical value. We wish to find the solution with optimal (maximum or minimum) value.
- Let's look at one more example before I discuss the general strategy behind dynamic programming.
- We'll look at what CLRS calls the *rod-cutting problem*. Imagine we're working for a company that buys large lengths of steel rod and cuts them up to resale the smaller pieces. We want to maximize the revenue we make.
- For our problem input, we are given an *integer* n representing the length of our original rod. We are also given an array $P[1 .. n]$ where $P[i]$ is the revenue we obtain selling a rod piece of length i . A solution to our problem is a list of lengths i_1, i_2, \dots, i_k such that $\sum_{j=1}^k i_j = n$. We want to maximize the value of our solution $\sum_{j=1}^k P[i_j]$.
- For example, if $n = 4$, and $P[1 .. n] = \langle 1, 5, 8, 9 \rangle$, we could sell the whole rod for 9 dollars. Or we could instead cut the rod into four pieces of length 1 for $4 * 1 = 4$ dollars. But the best option is to cut the rod into two pieces of length 2 for $5 + 5 = 10$ dollars.
- Now, when initially designing algorithms for optimization problems, it's easiest if we focus not on computing the solution itself but instead computing the optimal *value* of the solution. So that's what we'll do to start. We can come back to outputting the actual solutions later.
- The first and hardest step in designing a dynamic programming algorithm is to formulate and solve your problem recursively.
- I like to think of these kinds of optimization problems as us having to make a sequence of decisions. Once we've made a decision, we're left trying to deal with the consequences, and that hopefully means solving a smaller version of the same problem.

- So, we need to find a collection of rod lengths that sum up to n . **what is a first decision we might make for rod cutting?**
- Let's go with length of the first piece.
- Say we chop off a first piece of length j to sell it. We're left with a rod of length $n - j$. The best thing to do at this point is to maximize revenue on a rod of length $n - j$!
- We can most easily write down this relationship between the max revenue for length n and the max revenue for length $n - j$ using a recurrence relation.
- Let $\text{CutRod}(i)$ denote the maximum revenue obtainable cutting a rod of length n given the array $P[1 \dots n]$.
- $\text{CutRod}^*(i) =$
 - 0 if $i = 0$
 - $P[j] + \text{CutRod}(i - j)$ otherwise
- But wait. What is j ? We don't know in advance which length j is the best. Therefore, we'll have to try them all! We need to use the j that maximizes the total revenue from the first cut and the remaining cuts.
- $\text{CutRod}(i) =$
 - 0 if $i = 0$
 - $\max_{\{1 \leq j \leq i\}} P[j] + \text{CutRod}(i - j)$ otherwise
- This recurrence immediately implies an algorithm for computing the optimal revenue. Simply evaluate $\text{CutRod}(n)$ using the recurrence definition.
- This idea of trying each option for exactly one decision (e.g. length of the first piece) and then recursively solving problem instances consistent with each option is called *backtracking*.
- The specific idea that an optimal solution to this optimization problem incorporates optimal solutions to related subproblems is called the *optimal substructure* property by CLRS.

- OK, so we have *an* algorithm, but there's no telling how long it will take to run. (spoiler: it will take exponential time)
- We already saw the second step in dynamic programming: we need to solve all the subproblems by writing their solutions into a data structure.
- An array $\text{CutRod}[0 \dots n]$ will do. It has one entry per parameter value to the recurrence.
- We need to figure out what order we can fill the array. I like to do this by drawing a picture.
- We draw the array. We write down the solution to any base cases. Then we pick a generic element and draw arrays *into* it depicting which subproblems are used to get its solution.
- $[0 \quad \rightarrow \rightarrow \rightarrow \rightarrow \quad]$
- So we fill in the table from left to right (low index to high index).
- Now we're ready to write the iterative algorithm.
- For code:

- RodCutting(n):
 - $R[0] \leftarrow 0$
 - for $i \leftarrow 1$ to n
 - $q = 0$
 - for $j \leftarrow 1$ to i
 - $q = \max(q, P[j] + R[i - j])$
 - $R[j] = q$
 - return $R[n]$
- Now it's easy to figure out the running time. We have two for loops where i goes from 1 to n and j goes from 1 to at most n . The algorithm takes $O(n^2)$ time.

Dynamic Programming

- Dynamic programming is *recursion without repetition*. You store the solutions to the intermediate subproblems, often but not always in an array or table.
- Many students focus on the table, because tables are easy and familiar. But you need to focus on the much more important (and difficult) problem of finding a correct recurrence. If you memoize the correct recurrence, you may not even need an explicit table, but if the recursion is incorrect, nothing works. In particular, if you jump straight to making a table on your homework or exam solutions, you're going to have a much harder time convincing me your algorithm is correct, and my default assumption will be that it's actually wrong.

**Dynamic programming is *not* about filling in tables.
It's about smart recursion!**

- That said, there is a framework that you can and should follow to make things easier.
 1. Formulate the problem recursively. Write down a recursive formula or algorithm for the whole problem in terms of smaller subproblems. This is the hard part. (I greatly prefer recursive formulas.)
 1. Specification. Describe the problem you want to solve recursively in coherent and precise English. Not *how* to solve the problem, but *what* the problem is. Without this, we cannot tell if your solution is correct (the solution to what?).
 2. Solution. Give a clear recursive formula or algorithm for the whole problem in terms of answers to smaller instances of *exactly* the same problem.
 2. Build solutions to the recurrence from the bottom up. Write an algorithm that starts with base cases for your recurrence and works its way up to the final solution by considering intermediate subproblems in the correct order. This is the easy(er) part, and can be broken down into smaller relatively mechanical steps (an algorithm for designing algorithms!)
 1. Identify the subproblems. What are the different ways your recursive algorithm

can call itself? RecFibo took integers between 0 and n.

2. Choose a memoization data structure. Find a structure to store the solution to every subproblem from part (a). This is usually *but not always* a multidimensional array.
 3. Identify dependencies. Except for base cases, every subproblem depends on other subproblems. Which ones? Draw a picture of your data structure, pick a generic element, and draw arrows from each element it depends upon.
 4. Find a good evaluation order. Order the subproblems so that each one comes *after* the subproblems it depends on. Base cases first, then subproblems that depends only on base cases, and so on.
 5. Analyze space and running time. The number of distinct subproblems determine the space complexity of your algorithm. For total running time, add up the running time of all possible subproblems, assuming deeper recursive calls have already been memoized.
 - Usually, running time is [number of subproblems] * [time per subproblem].
 6. Write down the algorithm. You know the order to consider subproblems and how to solve them, so do that! If the data structure is an array, you'll usually write some nested for-loops around the recurrence.
- Don't forget to prove these steps are correct. If you write the wrong recurrence or fill the table in the wrong order, the algorithm won't work!

Greed is Bad, Actually

- If you're very lucky, you might be able to run a greedy algorithm.
- This means committing to a good looking first decision directly, without looking at any recursive subproblems, and then letting the Recursion Fairy do everything else. So backtracking without every looking tracking back.
- It seems natural, but very few problems can be solved correctly in this way.
- For example, for rod cutting, you might just take make a first cut of highest revenue and then commit to recursively cutting the rest of the rod.
- But even a simple example $n = 3, P = \langle 2, 2, 3 \rangle$ show that doesn't work.
- So remember, GREEDY ALGORITHMS NEVER WORK (except very rarely).
- If you ever get an inkling that a greedy approach might work, you really want to use dynamic programming instead. Really.
- In a couple weeks, we'll go over what's involved in designing a greedy algorithm and proving it correct. Until then, don't even try to use one.