Potes enim videre in hac margine, qualiter hoc operati fuimus, scilicet quod iunximus primum numerum cum secundo, videlicet 1 cum 2; et secundum cum tercio; et tercium cum quarto; et quartum cum quinto, et sic deinceps....

[You can see in the margin here how we have worked this; clearly, we combined the first number with the second, namely 1 with 2, and the second with the third, and the third with the fourth, and the fourth with the fifth, and so forth....]

— Leonardo Pisano, *Liber Abaci* (1202)

Those who cannot remember the past are condemned to repeat it.

— Jorge Agustín Nicolás Ruiz de Santayana y Borrás, *The Life of Reason, Book I: Introduction and Reason in Common Sense* (1905)

You know what a learning experience is?
A learning experience is one of those things that says,
"You know that thing you just did? Don't do that."

— Douglas Adams, *The Salmon of Doubt* (2002)

3

Dynamic Programming

3.1 Mātrāvṛtta

One of the earliest examples of recursion arose in India more than 2000 years ago, in the study of poetic meter, or prosody. Classical Sanskrit poetry distinguishes between two types of syllables (*akṣara*): *light* (*laghu*) and *heavy* (*guru*). In one class of meters, variously called *mātrāvṛtta* or *mātrāchandas*, each line of poetry consists of a fixed number of "beats" (*mātrā*), where each light syllable lasts one beat and each heavy syllable lasts two beats. The formal study of *mātrā-vṛtta* dates back to the *Chandaḥśāstra*, written by the scholar Piṅgala between 600BCE and 200BCE. Piṅgala observed that there are exactly five 4-beat meters: ——, —••, •—•, ••—, and ••••. (Here each "—" represents a long syllable and each "•" represents a short syllable.)[1]

[1] In Morse code, a "dah" lasts three times as long as a "dit", but each "dit" or "dah" is followed by a pause with the same duration as a "dit". Thus, each "dit-pause" is a *laghu akṣara*, each

Although Piṅgala's text *hints* at a systematic rule for counting meters with a given number of beats,[2] it took about a millennium for that rule to be stated explicitly. In the 7th century CE, another Indian scholar named Virahāṇka wrote a commentary on Piṅgala's work, in which he observed that the number of meters with n beats is the sum of the number of meters with $(n-2)$ beats and the number of meters with $(n-1)$ beats. In more modern notation, Virahāṇka's observation implies a recurrence for the total number $M(n)$ of n-beat meters:

$$M(n) = M(n-2) + M(n-1)$$

It is not hard to see that $M(0) = 1$ (there is only one empty meter) and $M(1) = 1$ (the only one-beat meter consists of a single short syllable).

The same recurrence reappeared in Europe about 500 years after Virahāṇka, in Leonardo of Pisa's 1202 treatise *Liber Abaci*, one of the most influential early European works on "algorism". In full compliance with Stigler's Law of Eponymy,[3] the modern *Fibonacci numbers* are defined using Virahāṇka's recurrence, but with different base cases:

$$F_n = \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ F_{n-1} + F_{n-2} & \text{otherwise} \end{cases}$$

In particular, we have $M(n) = F_{n+1}$ for all n.

Backtracking Can Be Slow

The recursive definition of Fibonacci numbers immediately gives us a recursive algorithm for computing them. Here is the same algorithm written in pseudocode:

"dah-pause" is a *guru akṣara*, and there are exactly five letters (M, D, R, U, and H) whose codes last four *mātrā*.

[2]The *Chandaḥśāstra* contains *two* systematic rules for listing all meters with a given number of *syllables*, which correspond roughly to writing numbers in binary from left to right (like Greeks) or from right to left (like Egyptians). The same text includes a recursive algorithm to compute 2^n (the number of meters with n syllables) by repeated squaring, and (arguably) a recursive algorithm to compute binomial coefficients (the number of meters with k short syllables and n syllables overall).

[3]*"No scientific discovery is named after its original discoverer."* In his 1980 paper that gives the law its name, the statistician Stephen Stigler jokingly claimed that this law was first proposed by sociologist Robert K. Merton. However, similar statements were previously made by Vladimir Arnol'd in the 1970's ("Discoveries are rarely attributed to the correct person."), Carl Boyer in 1968 ("Clio, the muse of history, often is fickle in attaching names to theorems!"), Alfred North Whitehead in 1917 ("Everything of importance has been said before by someone who did not discover it."), and even Stephen's father George Stigler in 1966 ("If we should ever encounter a case where a theory is named for the correct man, it will be noted."). We will see *many* other examples of Stigler's law in this book.

```
RecFibo(n):
  if n = 0
      return 0
  else if n = 1
      return 1
  else
      return RecFibo(n − 1) + RecFibo(n − 2)
```

Unfortunately, this naive recursive algorithm is horribly slow. Except for the recursive calls, the entire algorithm requires only a constant number of steps: one comparison and possibly one addition. Let $T(n)$ denote the number of recursive calls to RecFibo; this function satisfies the recurrence

$$T(0) = 1, \quad T(1) = 1, \quad T(n) = T(n-1) + T(n-2) + 1,$$

which looks an awful lot like the recurrence for Fibonacci numbers themselves! Writing out the first several values of $T(n)$ suggests the closed-form solution $\boldsymbol{T(n) = 2F_{n+1} - 1}$, which we can verify by induction (hint, hint). So computing F_n using this algorithm takes about twice as long as just counting to F_n. Methods beyond the scope of this book[4] imply that $F_n = \Theta(\phi^n)$, where $\phi = (\sqrt{5}+1)/2 \approx 1.61803$ is the so-called *golden ratio*. In short, the running time of this recursive algorithm is exponential in n.

We can actually see this exponential growth directly as follows. Think of the recursion tree for RecFibo as a binary tree of additions, with only 0s and 1s at the leaves. Since the eventual output is F_n, exactly F_n of the leaves must have value 1; these leaves represent the calls to RecRibo(1). An easy inductive argument (hint, hint) implies that RecFibo(0) is called exactly F_{n-1} times. (If we just want an asymptotic bound, it's enough to observe that the number of calls to RecFibo(0) is at most the number of calls to RecFibo(1).) Thus, the recursion tree has exactly $F_n + F_{n-1} = F_{n+1} = O(F_n)$ leaves, and therefore, because it's a full binary tree, $2F_{n+1} - 1 = O(F_n)$ nodes altogether.

Memo(r)ization: Remember Everything

The obvious reason for the recursive algorithm's lack of speed is that it computes the same Fibonacci numbers over and over and over. A single call to RecFibo(n) results in one recursive call to RecFibo($n-1$), two recursive calls to RecFibo($n-2$), three recursive calls to RecFibo($n-3$), five recursive calls to RecFibo($n-4$), and in general F_{k-1} recursive calls to RecFibo($n-k$) for any integer $0 \le k < n$. Each call is recomputing some Fibonacci number from scratch.

We can speed up our recursive algorithm considerably by writing down the results of our recursive calls and looking them up again if we need them later.

[4] See http://algorithms.wtf for notes on solving backtracking recurrences.

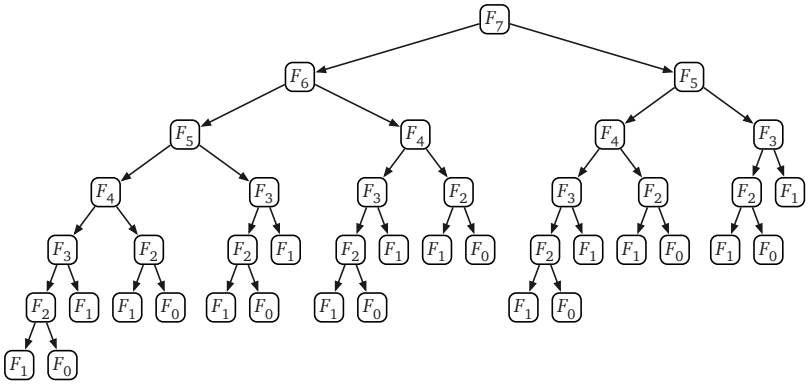

Figure 3.1. The recursion tree for computing F_7; arrows represent recursive calls.

This optimization technique, now known as *memoization* (yes, without an R), is usually credited to Donald Michie in 1967, but essentially the same technique was proposed in 1959 by Arthur Samuel.[5]

```
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]
```

Memoization clearly decreases the running time of the algorithm, but by how much? If we actually trace through the recursive calls made by MEMFIBO, we find that the array $F[\,]$ is filled from the bottom up: first $F[2]$, then $F[3]$, and so on, up to $F[n]$. This pattern can be verified by induction: Each entry $F[i]$ is filled only after its predecessor $F[i-1]$. If we ignore the time spent in recursive calls, it requires only constant time to evaluate the recurrence for each Fibonacci number F_i. But by design, the recurrence for F_i is evaluated only once for each index i. We conclude that MEMFIBO performs only $O(n)$ additions, an *exponential* improvement over the naïve recursive algorithm!

[5]Michie proposed that programming languages should support an abstraction he called a "memo function", consisting of both a standard function ("rule") and a dictionary ("rote"), instead of separately supporting arrays and functions. Whenever a memo function computes a function value for the first time, it "memorises" (yes, with an R) that value into its dictionary. Michie was inspired by Samuel's use of "rote learning" to speed up the recursive evaluation of checkers game trees; Michie describes his more general proposal as "enabling the programmer to 'Samuelize' any functions he pleases." (As far as I can tell, Michie never used the term "memoisation" himself.) Memoization was used even earlier by Claude Shannon's maze-solving robot "Theseus", which he designed and constructed in 1950.

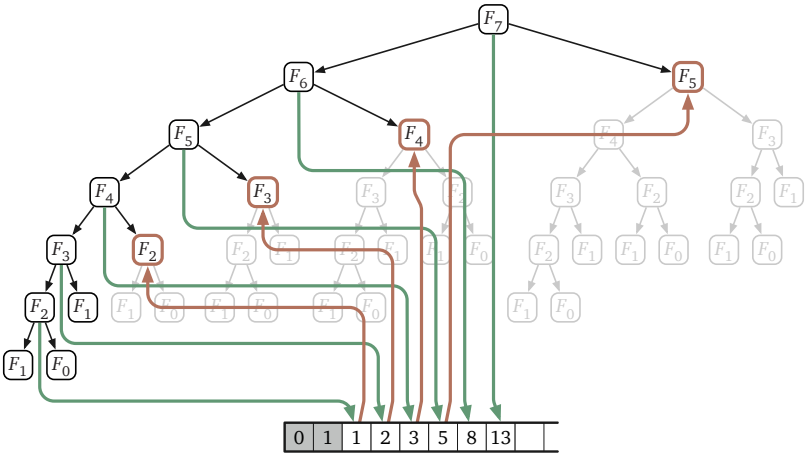

Figure 3.2. The recursion tree for F_7 trimmed by memoization. Downward green arrows indicate writing into the memoization array; upward red arrows indicate reading from the memoization array.

Dynamic Programming: Fill Deliberately

Once we see how the array $F[\,]$ is filled, we can replace the memoized recurrence with a simple for-loop that *intentionally* fills the array in that order, instead of relying on a more complicated recursive algorithm to do it for us accidentally.

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

Now the time analysis is immediate: IterFibo clearly uses $O(n)$ ***additions*** and stores $O(n)$ *integers*.

This is our first explicit ***dynamic programming*** algorithm. The dynamic programming paradigm was formalized and popularized by Richard Bellman in the mid-1950s, while working at the RAND Corporation, although he was far from the first to use the technique. In particular, this iterative algorithm for Fibonacci numbers was already proposed by Virahāṇka and later Sanskrit prosodists in the 12th century, and again by Fibonacci at the turn of the 13th century![6]

[6]More general dynamic programming techniques were independently deployed several times in the late 1930s and early 1940s. For example, Pierre Massé used dynamic programming algorithms to optimize the operation of hydroelectric dams in France during the Vichy regime. John von Neumann and Oskar Morgenstern developed dynamic programming algorithms to determine the winner of any two-player game with perfect information (for example, checkers). Alan Turing and his cohorts used similar methods as part of their code-breaking efforts at Bletchley

Many years after the fact, Bellman claimed that he deliberately chose the name "dynamic programming" to hide the mathematical character of his work from his military bosses, who were actively hostile toward anything resembling mathematical research.[7] The word "programming" does not refer to writing code, but rather to the older sense of *planning* or *scheduling*, typically by filling in a table. For example, sports programs and theater programs are schedules of important events (with ads); television programming involves filling each available time slot with a show (and ads); degree programs are schedules of classes to be taken (with ads). The Air Force funded Bellman and others to develop methods for constructing training and logistics schedules, or as they called them, "programs". The word "dynamic" was not only a reference to the multistage, time-varying processes that Bellman and his colleagues were attempting to optimize, but also a marketing buzzword that would resonate with the Futuristic Can-Do Zeitgeist™ of post-WW II America.[8] Thanks in part to Bellman's proselytizing, dynamic programming is now a standard tool for multistage planning in economics, robotics, control theory, and several other disciplines.

Don't Remember Everything After All

In many dynamic programming algorithms, it is not necessary to retain *all* intermediate results through the entire computation. For example, we can significantly reduce the space requirements of our algorithm IterFibo by maintaining only the two newest elements of the array:

Park. Both Massé's work and von Neumann and Mergenstern's work were first published in 1944, six years before Bellman coined the phrase "dynamic programming". The details of Turing's "Banburismus" were kept secret until the mid-1980s.

[7]Charles Erwin Wilson became Secretary of Defense in January 1953, after a dozen years as the president of General Motors. "Engine Charlie" reorganized the Department of Defense and significantly decreased its budget in his first year in office, with the explicit goal of running the Department much more like an industrial corporation. Bellman described Wilson in his 1984 autobiography as follows:

> We had a very interesting gentleman in Washington named Wilson. He was secretary of Defense, and he actually had a pathological fear and hatred of the word "*research*". I'm not using the term lightly; I'm using it precisely. His face would suffuse, he would turn red, and he would get violent if people used the term "*research*" in his presence. You can imagine how he felt, then, about the term "*mathematical*". . . . I felt I had to do something to shield Wilson and the Air Force from the fact that I was really doing mathematics inside the RAND Corporation. What title, what name, could I choose?

However, Bellman's first published use of the term "dynamic programming" already appeared in 1952, several months before Wilson took office, so this story is at least *slightly* embellished.

[8]. . . and just possibly a riff on the iconic brand name "Dynamic-Tension" for Charles Atlas's famous series of exercises, which Charles Roman coined in 1928. Hero of the Beach!

```
ITERFIBO2(n):
  prev ← 1
  curr ← 0
  for i ← 1 to n
      next ← curr + prev
      prev ← curr
      curr ← next
  return curr
```

(This algorithm uses the non-standard but consistent base case $F_{-1} = 1$ so that ITERFIBO2(0) returns the correct value 0.) Although saving space can be absolutely crucial in practice, we won't focus on space issues in this book.

♥3.2 Aside: Even Faster Fibonacci Numbers

Although the previous algorithm is simple and attractive, it is *not* the fastest algorithm to compute Fibonacci numbers. We can derive a faster algorithm by exploiting the following matrix reformulation of the Fibonacci recurrence:

$$\begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} y \\ x+y \end{bmatrix}$$

In other words, multiplying a two-dimensional vector by the matrix $\left[\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right]$ has exactly the same effect as one iteration of the inner loop of ITERFIBO2. It follows that multiplying by the matrix n times is the same as iterating the loop n times:

$$\begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}^n \begin{bmatrix} 1 \\ 0 \end{bmatrix} = \begin{bmatrix} F_{n-1} \\ F_n \end{bmatrix}.$$

So if we want the nth Fibonacci number, we only need to compute the nth power of the matrix $\left[\begin{smallmatrix} 0 & 1 \\ 1 & 1 \end{smallmatrix}\right]$. If we use repeated squaring,[9] computing the nth power of something requires only $O(\log n)$ multiplications. Here, because "something" is a 2×2 matrix, that means $O(\log n)$ 2×2 matrix multiplications, each of which reduces to a constant number of integer multiplications and additions. Thus, we can compute F_n in only ***O(log n) integer arithmetic operations***.

We can achieve the same speedup using the identity $F_n = F_m F_{n-m-1} + F_{m+1} F_{n-m}$, which holds (by induction!) for all integers m and n. In particular, this identity implies the following mutual recurrence for pairs of adjacent Fibonacci numbers, first proposed by Édouard Lucas in 1898:

$$\begin{aligned} F_{2n-1} &= F_{n-1}^2 + F_n^2 \\ F_{2n} &= F_n(F_{n-1} + F_{n+1}) = F_n(2F_{n-1} + F_n) \end{aligned}$$

[9] as suggested by Piṅgala for powers of 2 elsewhere in *Chandaḥśāstra*

(We can also derive this mutual recurrence directly from the matrix-squaring algorithm.) These recurrences translate directly into the following algorithm:

```
⟨⟨Compute the pair F_{n-1}, F_n⟩⟩
FastRecFibo(n):
  if n = 1
      return 0, 1
  m ← ⌊n/2⌋
  hprv, hcur ← FastRecFibo(m)     ⟨⟨F_{m-1}, F_m⟩⟩
  prev ← hprv² + hcur²            ⟨⟨F_{2m-1}⟩⟩
  curr ← hcur · (2 · hprv + hcur) ⟨⟨F_{2m}⟩⟩
  next ← prev + curr              ⟨⟨F_{2m+1}⟩⟩
  if n is even
      return prev, curr
  else
      return curr, next
```

Our standard recursion tree technique implies that this algorithm performs only $O(\log n)$ integer arithmetic operations.

This is an exponential speedup over the standard iterative algorithm, which was already an exponential speedup over our original recursive algorithm. Right?

Whoa! Not so fast!

Well, not exactly. Fibonacci numbers grow exponentially fast. The nth Fibonacci number is approximately $n \log_{10} \phi \approx n/5$ decimal digits long, or $n \log_2 \phi \approx 2n/3$ bits. So we can't possibly compute F_n in logarithmic *time* — we need $\Omega(n)$ time just to write down the answer!

The way out of this apparent paradox is to observe that ***we can't perform arbitrary-precision arithmetic in constant time***. Let $M(n)$ denote the time required to multiply two n-digit numbers. The running time of FastRecFibo satisfies the recurrence $T(n) = T(\lfloor n/2 \rfloor) + M(n)$, which solves to $T(n) = \boldsymbol{O(M(n))}$ via recursion trees. The fastest integer multiplication algorithm known (as of 2019) runs in $O(n \log n)$ time, so that is also the running time of the fastest algorithm known (as of 2019) to compute Fibonacci numbers.

Is this algorithm slower than our "linear-time" iterative algorithms? Actually, no—addition isn't free, either! Adding two n-digit numbers requires $O(n)$ time, so the iterative algorithms IterFibo and IterFibo2 actually run in $\boldsymbol{O(n^2)}$ ***time***. (Do you see why?) So FastRecFibo is significantly faster than the iterative algorithms, just not *exponentially* faster.

In the original recursive algorithm, the extra cost of arbitrary-precision arithmetic is overwhelmed by the huge number of recursive calls. The correct

recurrence is $T(n) = T(n-1) + T(n-2) + O(n)$, which still has the solution $T(n) = O(\phi^n)$.

3.3 *Interpunctio Verborum Redux*

For our next dynamic programming algorithm, let's consider the text segmentation problem from the previous chapter. We are given a string $A[1..n]$ and a subroutine IsWord that determines whether a given string is a word (whatever that means), and we want to know whether A can be partitioned into a sequence of words.

We solved this problem by defining a function $Splittable(i)$ that returns True if and only if the suffix $A[i..n]$ can be partitioned into a sequence of words. We need to compute $Splittable(1)$. This function satisfies the recurrence

$$Splittable(i) = \begin{cases} \text{True} & \text{if } i > n \\ \bigvee_{j=i}^{n} \big(IsWord(i,j) \wedge Splittable(j+1) \big) & \text{otherwise} \end{cases}$$

where $IsWord(i,j)$ is shorthand for $IsWord(A[i..j])$. This recurrence translates directly into a recursive backtracking algorithm that calls the IsWord subroutine $O(2^n)$ times in the worst case.

But for any fixed string $A[1..n]$, there are only n different ways to call the recursive function $Splittable(i)$—one for each value of i between 1 and $n+1$—and only $O(n^2)$ different ways to call IsWord(i,j)—one for each pair (i,j) such that $1 \le i \le j \le n$. Why are we spending exponential time computing only a polynomial amount of stuff?

Each recursive subproblem is specified by an integer between 1 and $n+1$, so we can memoize the function $Splittable$ into an array $SplitTable[1..n+1]$. Each subproblem $Splittable(i)$ depends only on results of subproblems $Splittable(j)$ where $j > i$, so the memoized recursive algorithm fills the array in *decreasing* index order. If we fill the array in this order deliberately, we obtain the dynamic programming algorithm shown in Figure 3.3. The algorithm makes $O(n^2)$ calls to IsWord, an exponential improvement over our earlier backtracking algorithm.

3.4 The Pattern: Smart Recursion

In a nutshell, dynamic programming is *recursion without repetition*. Dynamic programming algorithms store the solutions of intermediate subproblems, often *but not always* in some kind of array or table. Many algorithms students

```
FastSplittable(A[1..n]):
  SplitTable[n+1] ← True
  for i ← n down to 1
      SplitTable[i] ← False
      for j ← i to n
          if IsWord(i, j) and SplitTable[j+1]
              SplitTable[i] ← True
  return SplitTable[1]
```

Figure 3.3. Interpunctio verborum velox

(and instructors, and textbooks) make the mistake of focusing on the table—because tables are easy and familiar—instead of the *much* more important (and difficult) task of finding a correct recurrence. As long as we memoize the correct recurrence, an explicit table isn't really necessary, but if the recurrence is incorrect, we are well and truly hosed.

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

Dynamic programming algorithms are best developed in two distinct stages, each of which can be broken into several simpler steps. (This list of steps may look intimidating at first, but most of the steps are straightforward.)

1. **Formulate the problem recursively.** Write down a recursive formula or algorithm for the whole problem in terms of the answers to smaller subproblems. This is the hard part. A complete recursive formulation has two parts:

   (a) **Imagine solving the problem yourself,** by hand, not by writing code. What sequence of decisions do you need to make? For each decision, what are your options?

   (b) **Describe the "shape" of your recursive subproblems.** After you've made a few decisions, what portion(s) of the input data have you not yet considered? A prefix, a suffix, one or more intervals, or something else? What other information do you need to remember about your past decisions to make consistent future decisions?

   (c) **Specify the "shape" of each subproblem using indices**. For example, you can specify a prefix by its length (not the subarray), an interval by its starting and ending index (not the subarray), and an input value by its position in the input array (not the actual numerical value).

   (d) **Specification.** Describe the problem that you want to solve recursively, in coherent and precise English—not *how* to solve the problem, but *what*

problem you are trying to solve. Equivalently, after you've made some decisions, what precise question are you asking about the remaining input data? ***This is the most important step!!*** Without this specification, it is impossible, even in principle, to determine whether your algorithm is correct.[10]

(e) **Reduce the original problem** to the recursive function you specified in the previous step. This often requires only a single invocation of the recursive function, but sometimes requires multiple function calls.

(f) **Recurrence.** Finally, develop a recurrence that answers the precise question you specified in step 1(d), formalizing your earlier intuition. Your recursive algorithm should explore all possible options for exactly one decision, and each option should be evaluated by recursively solving smaller instances of *exactly* the same problem.

2. **Build solutions to your recurrence from the bottom up.** Write an algorithm that starts with the base cases of your recurrence and works its way up to the final solution, by considering intermediate subproblems in the correct order. This stage can be broken down into several smaller, relatively mechanical steps:

(a) **Identify the subproblems.** What are all the different ways your recursive algorithm can call itself, starting with some initial input? For example, the argument to RecFibo is always an integer between 0 and n.

(b) **Choose a memoization data structure.** Find a data structure that can store the solution to *every* subproblem you identified in step (a). This is usually *but not always* a 1-, 2-, or 3-dimensional array. If your first thought is "hash table", you can do better.

(c) **Identify dependencies.** Except for the 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 of the other elements it depends on. Then formalize your picture.

(d) **Find a good evaluation order.** Order the subproblems so that each one comes *after* the subproblems it depends on. Your order should consider the base cases first, eventually building up to the original top-level problem. The dependencies you identified in the previous step define a partial order over the subproblems; you need to find a linear extension of that partial order. ***Be careful!***

(e) **Analyze space and running time.** The number of distinct subproblems determines the space complexity of your memoized algorithm. To

[10] And for the love of all that is sane and holy, give your function a name that actually reflects its *meaning*. Calling your function *OPT* or *dp* angers the Great Old Ones. Iä! Shub-Niggurath!

compute the total running time, add up the running times of all possible subproblems, *assuming deeper recursive calls are already memoized*. You can actually do this immediately after step (a).

(f) **Write down the algorithm.** You know what order to consider the subproblems, and you know how to solve each subproblem. So do that! If your data structure is an array, this usually means writing a few nested for-loops around your original recurrence, and replacing the recursive calls with array look-ups.

Of course, you have to prove that each of these steps is correct. If your recurrence is wrong, or if you try to build up answers in the wrong order, your algorithm won't work!

3.5 Warning: Greed is Stupid

If we're incredibly lucky, we can bypass all the recurrences and tables and so forth, and solve the problem using a *greedy* algorithm. Like a backtracking algorithm, a greedy algorithm constructs a solution through a series of decisions, but it makes those decisions directly, *without* solving at any recursive subproblems. While this approach seems very natural, it almost never works; optimization problems that can be solved correctly by a greedy algorithm are quite rare. Nevertheless, for many problems that should be solved by backtracking or dynamic programming, many students' first intuition is to apply a greedy strategy.

For example, a greedy algorithm for the text segmentation problem might find the shortest (or, if you prefer, longest) prefix of the input string that is a word, accept that prefix as the first word in the segmentation, and then recursively segment the remaining suffix of the input string. Similarly, a greedy algorithm for the longest increasing subsequence problem might look for the smallest element of the input array, accept that element as the start of the target subsequence, and then recursively look for the longest increasing subsequence to the right of that element. If these sound like stupid hacks to you, pat yourself on the back; these aren't even *close* to correct solutions.

Everyone should tattoo the following sentence on the back of their hands, right under all the rules about logarithms and big-Oh notation:

Greedy algorithms never work!
Use dynamic programming instead!

What, never?
No, never!

What, *never*?

Well. . . hardly ever.[11]

Because the greedy approach is so incredibly tempting, but so rarely correct, I strongly advocate the following policy in any algorithms course, even (or perhaps *especially*) for courses that do not normally ask for proofs of correctness.[12]

You will not receive *any* credit for *any* greedy algorithm, on *any* homework or exam, even if the algorithm is correct, without a *formal* proof of correctness.

Moreover, the vast majority of problems for which students are tempted to submit a greedy algorithm are actually best solved using dynamic programming. So I always offer the following advice to my algorithms students.

Whenever you write—or even *think*—the word "greeDY", your subconscious is telling you to use DYnamic programming.

Even for problems that *can* be correctly solved by greedy algorithms, it's usually more productive to develop a backtracking or dynamic programming algorithm first. First make it work, then make it fast. We will see techniques for proving greedy algorithms correct in the next chapter.

3.6 Longest Increasing Subsequence

Another problem we considered in the previous chapter was computing the length of the longest increasing subsequence of a given array $A[1..n]$ of numbers. We developed two different recursive backtracking algorithms for this problem. Both algorithms run in $O(2^n)$ time in the worst case; both algorithms can be sped up significantly via dynamic programming.

First Recurrence: Is This Next?

Our first backtracking algorithm evaluated the function $LISbigger(i, j)$, which we defined as the length of the longest increasing subsequence of $A[j..n]$ in which every element is larger than $A[i]$. We derived the following recurrence

[11] They hardly ever ever work! Then give three cheers, and one cheer more, for the rig'rous Captain of the *Pinafore*! Then give three cheers, and one cheer more, for the Captain of the *Pinafore*!

[12] Introducing this policy in my own algorithms courses significantly improved students' grades, because it significantly reduced the frequency of incorrect greedy algorithms.

for this function:

$$LISbigger(i,j) = \begin{cases} 0 & \text{if } j > n \\ LISbigger(i, j+1) & \text{if } A[i] \geq A[j] \\ \max\left\{ \begin{array}{l} LISbigger(i, j+1) \\ 1 + LISbigger(j, j+1) \end{array} \right\} & \text{otherwise} \end{cases}$$

To solve the original problem, we can add a sentinel value $A[0] = -\infty$ to the array and compute $LISbigger(0,1)$.

Each recursive subproblem is identified by two indices i and j, so there are only $O(n^2)$ distinct recursive subproblems to consider. We can memoize the results of these subproblems into a two-dimensional array $\boldsymbol{LISbigger[0..n, 1..n]}$.[13] Moreover, each subproblem can be solved in $O(1)$ time, not counting recursive calls, so we should expect the final dynamic programming algorithm to run in $\boldsymbol{O(n^2)}$ ***time***.

The order in which the memoized recursive algorithm fills this array is not immediately clear; all we can tell from the recurrence is that each entry $LISbigger[i,j]$ is filled in *after* the entries $LISbigger[i,j+1]$ and $LISbigger[j,j+1]$ in the next column, as indicated on the left in Figure 3.4.

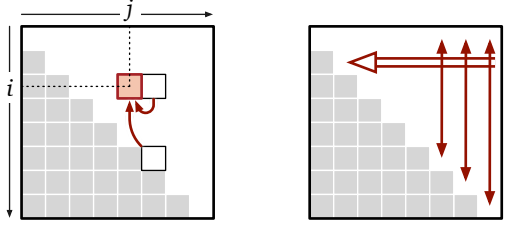

Figure 3.4. Subproblem dependencies for longest increasing subsequence, and a valid evaluation order

Fortunately, this partial information is enough to give us a *valid* evaluation order. If we fill the table one column at a time, from right to left, then whenever we reach an entry in the table, the entries it depends on are already available. This may not be the order that the recursive algorithm would use, but it works, so we'll go with it. The right figure in Figure 3.4 illustrates this evaluation order, with a double arrow indicating the outer loop and single arrows indicating the inner loop. In this case, the single arrows are bidirectional, because the order that we use to fill each column doesn't matter.

And we're done! Pseudocode for our dynamic programming algorithm is shown below; as expected, our algorithm clearly runs in $O(n^2)$ time. If necessary,

[13] In fact, we only need half of this array, because we always have $i < j$. But even if we cared about constant factors in this book (we don't), this would be the wrong time to worry about them. First make it work; then make it better.

we can reduce the space bound from $O(n^2)$ to $O(n)$ by maintaining only the two most recent columns of the table, $LISbigger[\cdot, j]$ and $LISbigger[\cdot, j+1]$.[14]

```
FastLIS(A[1..n]):
  A[0] ← −∞                          ⟨⟨Add a sentinel⟩⟩
  for i ← 0 to n                     ⟨⟨Base cases⟩⟩
      LISbigger[i, n+1] ← 0
  for j ← n down to 1
      for i ← 0 to j − 1             ⟨⟨...or whatever⟩⟩
          keep ← 1 + LISbigger[j, j+1]
          skip ← LISbigger[i, j+1]
          if A[i] ≥ A[j]
              LISbigger[i, j] ← skip
          else
              LISbigger[i, j] ← max{keep, skip}
  return LISbigger[0, 1]
```

Second Recurrence: What's Next?

Our second backtracking algorithm evaluated the function $LISfirst(i)$, which we defined as the length of the longest increasing subsequence of $A[i..n]$ that begins with $A[i]$. We derived the following recurrence for this function:

$$LISfirst(i) = 1 + \max\left\{LISfirst(j) \;\middle|\; j > i \text{ and } A[j] > A[i]\right\}$$

Here, we assume that $\max \varnothing = 0$, so that the base cases like $LISfirst(n) = 1$ fall out of the recurrence automatically. To solve the original problem, we can add a sentinel value $A[0] = -\infty$ to the array and compute $LISfirst(0) - 1$.

In this case, recursive subproblems are indicated by a single index i, so we can memoize the recurrence into a one-dimensional array $LISfirst[1..n]$. Each entry $LISfirst[i]$ depends only on entries $LISfirst[j]$ with $j > i$, so we can fill the array in decreasing index order. To compute each $LISfirst[i]$, we need to consider $LISfirst[j]$ for *all* indices $j > i$, but we don't need to consider those indices j in any particular order. The resulting dynamic programming algorithm runs in $\boldsymbol{O(n^2)}$ ***time*** and uses $O(n)$ space.

```
FastLIS2(A[1..n]):
  A[0] = −∞                          ⟨⟨Add a sentinel⟩⟩
  for i ← n down to 0
      LISfirst[i] ← 1
      for j ← i + 1 to n             ⟨⟨...or whatever⟩⟩
          if A[j] > A[i] and 1 + LISfirst[j] > LISfirst[i]
              LISfirst[i] ← 1 + LISfirst[j]
  return LISfirst[0] − 1             ⟨⟨Don't count the sentinel⟩⟩
```

[14] See, I told you not to worry about constant factors yet!

3.7 Edit Distance

The ***edit distance*** between two strings is the minimum number of letter insertions, letter deletions, and letter substitutions required to transform one string into the other. For example, the edit distance between FOOD and MONEY is at most four:

$$\underline{F}OOD \rightarrow MO\underline{O}D \rightarrow MON\underline{\wedge}D \rightarrow MONE\underline{D} \rightarrow MONEY$$

This distance function was independently proposed by Vladimir Levenshtein in 1965 (working on coding theory), Taras Vintsyuk in 1968 (working on speech recognition), and Stanislaw Ulam in 1972 (working with biological sequences). For this reason, edit distance is sometimes called ***Levenshtein distance*** or ***Ulam distance*** (but strangely, never "Vintsyuk distance").

We can visualize this editing process by aligning the strings one above the other, with a gap in the first word for each insertion and a gap in the second word for each deletion. Columns with two *different* characters correspond to substitutions. In this representation, the number of editing steps is just the number of columns that do *not* contain the same character twice.

```
F O O   D
M O N E Y
```

It's fairly obvious that we can't transform FOOD into MONEY in three steps, so the edit distance between FOOD and MONEY is exactly four. Unfortunately, it's not so easy in general to tell when a sequence of edits is as short as possible. For example, the following alignment shows that the distance between the strings ALGORITHM and ALTRUISTIC is *at most* 6. Is that the best we can do?

```
A L G O R   I   T H M
A L   T R U I S T I C
```

Recursive Structure

To develop a dynamic programming algorithm to compute edit distance, we first need to formulate the problem recursively. Our alignment representation for edit sequences exposes an important recursive structure.[15] Suppose we have the gap representation for the shortest edit sequence for two strings. **If we remove the last column, the remaining columns must represent the shortest edit sequence for the remaining prefixes.** We can easily prove this observation

[15] Many algorithms textbooks use the phrase "optimal substructure" here, but that phrases incorrectly presumes that all dynamic programming can only be used to solve *optimization* problems.

by contradiction: If the prefixes had a shorter edit sequence, gluing the last column back on would give us a shorter edit sequence for the original strings. So once we figure out what should happen in the last column, the Recursion Fairy can figure out the rest of the optimal gap representation.

Said differently, the alignment we are looking for represents a sequence of editing operations, ordered (for no particular reason) from right to left. Solving the edit distance problem requires making a sequence of decisions, one for each column in the output alignment. In the middle of this sequence of decisions, we have already aligned a suffix of one string with a suffix of the other.

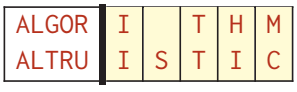

Because the cost of an alignment is just the number of mismatched columns, our remaining decisions don't depend on the editing operations we've already chosen; they only depend on the prefixes we haven't aligned yet.

Thus, for any two input strings $A[1..m]$ and $B[1..n]$, we can formulate the edit distance problem recursively as follows: For any indices i and j, let $\boldsymbol{Edit(i, j)}$ denote the edit distance between the prefixes $A[1..i]$ and $B[1..j]$. We need to compute $Edit(m, n)$.

Recurrence

When i and j are both positive, there are exactly three possibilities for the last column in the optimal alignment of $A[1..i]$ and $B[1..j]$:

- **Insertion:** The last entry in the top row is empty. In this case, the edit distance is equal to $Edit(i, j-1)+1$. The $+1$ is the cost of the final insertion, and the recursive expression gives the minimum cost for the remaining alignment.

- **Deletion:** The last entry in the bottom row is empty. In this case, the edit distance is equal to $Edit(i-1, j)+1$. The $+1$ is the cost of the final deletion, and the recursive expression gives the minimum cost for the remaining alignment.

- **Substitution:** Both rows have characters in the last column. If these two characters are different, then the edit distance is equal to $Edit(i-1, j-1)+1$. If these two characters are equal, the substitution is free, so the edit distance is $Edit(i-1, j-1)$.

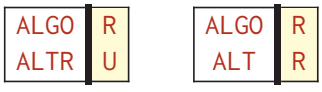

This generic case analysis breaks down if either $i = 0$ or $j = 0$, but those boundary cases are easy to handle directly.

- Transforming the empty string into a string of length j requires j insertions, so $Edit(0, j) = j$.
- Transforming a string of length i into the empty string requires i deletions, so $Edit(i, 0) = i$.

As a sanity check, both of these base cases correctly indicate that the edit distance between the empty string and the empty string is zero!

We conclude that the *Edit* function satisfies the following recurrence:

$$Edit(i,j) = \begin{cases} i & \text{if } j = 0 \\ j & \text{if } i = 0 \\ \min \left\{ \begin{array}{c} Edit(i, j-1) + 1 \\ Edit(i-1, j) + 1 \\ Edit(i-1, j-1) + [A[i] \neq B[j]] \end{array} \right\} & \text{otherwise} \end{cases}$$

Dynamic Programming

Now that we have a recurrence, we can transform it into a dynamic programming algorithm following our usual mechanical recipe.

- **Subproblems:** Each recursive subproblem is identified by two indices $0 \le i \le m$ and $0 \le j \le n$.
- **Memoization structure:** So we can memoize all possible values of $Edit(i, j)$ in a two-dimensional array $Edit[0..m, 0..n]$.
- **Dependencies:** Each entry $Edit[i, j]$ depends only on its three neighboring entries $Edit[i-1, j]$, $Edit[i, j-1]$, and $Edit[i-1, j-1]$.
- **Evaluation order:** If we fill this array in standard row-major order—row by row from top down, each row from left to right—then whenever we reach an entry in the array, all the entries it depends on are already available. (This isn't the *only* evaluation order we could use, but it works, so let's go with it.)

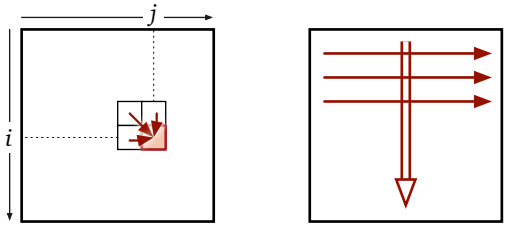

- **Space and time:** The memoization structure uses $O(mn)$ space. We can compute each entry $Edit[i, j]$ in $O(1)$ time once we know its predecessors, so the overall algorithm runs in ***O(mn) time.***

Here is the resulting dynamic programming algorithm:

```
EditDistance(A[1..m], B[1..n]):
  for j ← 0 to n
      Edit[0, j] ← j

  for i ← 1 to m
      Edit[i, 0] ← i
      for j ← 1 to n
          ins ← Edit[i, j − 1] + 1
          del ← Edit[i − 1, j] + 1
          if A[i] = B[j]
              rep ← Edit[i − 1, j − 1]
          else
              rep ← Edit[i − 1, j − 1] + 1
          Edit[i, j] ← min {ins, del, rep}

  return Edit[m, n]
```

This algorithm is most commonly attributed to Robert Wagner and Michael Fischer, who described the algorithm in 1974. However, in full compliance with Stigler's Law of Eponymy, either identical or more general algorithms were independently discovered by Taras Vintsyuk in 1968, V. M. Velichko and N. G. Zagoruyko in 1970, David Sankoff in 1972, Peter Sellers in 1974, and almost certainly several others.[16] Interestingly, *none* of these authors cite either Levenshtein or Ulam!

[16]This algorithm is sometimes also *incorrectly* attributed to Saul Needleman and Christian Wunsch in 1970. "The Needleman-Wunsch algorithm" more commonly refers to the standard dynamic programming algorithm for computing the longest common subsequence of two strings (or equivalently, the edit distance where only insertions and deletions are permitted) in $O(mn)$ time, but that attribution is *also* incorrect! In fact, Needleman and Wunsch's algorithm computes (weighted) longest common subsequences (possibly with gap costs) in $O(m^2n^2)$ time, using a different recurrence. Sankoff explicitly describes his $O(mn)$-time algorithm as an improvement of Needleman and Wunsch's algorithm.

The memoization table for the input strings ALGORITHM and ALTRUISTIC is shown below. Bold numbers indicate places where characters in the two strings are equal. The edit distance between ALGORITHM and ALTRUISTIC is indeed six!

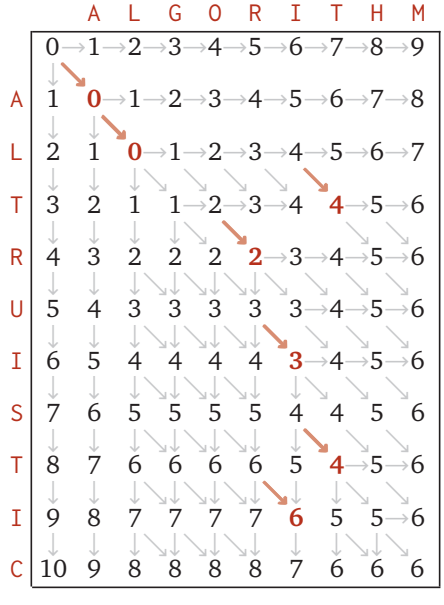

The arrows in this table indicate which predecessor(s) actually define each entry. Each direction of arrow corresponds to a different edit operation: horizontal=deletion, vertical=insertion, and diagonal=substitution. Bold red diagonal arrows indicate "free" substitutions of a letter for itself. Any path of arrows from the top left corner to the bottom right corner of this table represents an optimal edit sequence between the two strings. The example memoization array contains exactly three directed paths from the top left corner to the bottom right corner, each indicating a different sequence of six edits transforming ALGORITHM into ALTRUISTIC, as shown on the next page.

```
A L G O R I   T H M
A L T R U I S T I C

A L G O R   I   T H M
A L   T R U I S T I C

A L G O R   I   T H M
A L T   R U I S T I C
```

Our EDITDISTANCE algorithm does not actually compute or store any arrows in the table, but the arrow(s) leading into any entry in the table can be reconstructed on the fly in $O(1)$ time from the numerical values. Thus, once

we've filled in the table, we can reconstruct the shortest edit sequence in $O(n+m)$ additional time.

3.8 Subset Sum

Recall that the *Subset Sum* problem asks whether any subset of a given array $X[1..n]$ of positive integers sums to a given integer T. In the previous chapter, we developed a recursive Subset Sum algorithm that can be reformulated as follows. Fix the original input array $X[1..n]$ and define the boolean function

$$SS(i,t) = \textsc{True} \text{ if and only if some subset of } X[i..n] \text{ sums to } t.$$

We need to compute $SS(1,T)$. This function satisfies the following recurrence:

$$SS(i,t) = \begin{cases} \textsc{True} & \text{if } t = 0 \\ \textsc{False} & \text{if } t < 0 \text{ or } i > n \\ SS(i+1,t) \;\vee\; SS(i+1,t-X[i]) & \text{otherwise} \end{cases}$$

We can transform this recurrence into a dynamic programming algorithm following the usual boilerplate.

- **Subproblems:** Each subproblem is described by an integer i such that $1 \le i \le n+1$, and an integer $t \le T$. However, subproblems with $t < 0$ are trivial, so it seems rather silly to memoize them.[17] Indeed, we can modify the recurrence so that those subproblems never arise:

$$SS(i,t) = \begin{cases} \textsc{True} & \text{if } t = 0 \\ \textsc{False} & \text{if } i > n \\ SS(i+1,t) & \text{if } t < X[i] \\ SS(i+1,t) \;\vee\; SS(i+1,t-X[i]) & \text{otherwise} \end{cases}$$

- **Data structure:** We can memoize our recurrence into a two-dimensional array $SS[1..n+1, 0..T]$, where $SS[i,t]$ stores the value of $SS(i,t)$.
- **Evaluation order:** Each entry $SS[i,t]$ depends on at most two other entries, both of the form $SS[i+1,\cdot]$. So we can fill the array by considering rows from bottom to top in the outer loop, and considering the elements in each row in arbitrary order in the inner loop.
- **Space and time:** The memoization structure uses $O(nT)$ space. If $S[i+1,t]$ and $SS[i+1,t-X[i]]$ are already known, we can compute $SS[i,t]$ in constant time, so the algorithm runs in $\boldsymbol{O(nT)}$ ***time***.

[17] Yes, I'm breaking my own rule against premature optimization.

Here is the resulting dynamic programming algorithm:

```
FastSubsetSum(X[1..n], T):
  S[n+1,0] ← True
  for t ← 1 to T
      SS[n+1,t] ← False

  for i ← n down to 1
      SS[i,0] = True
      for t ← 1 to X[i]−1
          SS[i,t] ← SS[i+1,t]        ⟨⟨Avoid the case t < 0⟩⟩
      for t ← X[i] to T
          SS[i,t] ← SS[i+1,t] ∨ SS[i+1,t−X[i]]

  return SS[1,T]
```

The worst-case running time $O(nT)$ for this algorithm is a significant improvement over the $O(2^n)$-time recursive backtracking algorithm when T is small.[18] However, if the target sum T is significantly larger than 2^n, this iterative algorithm is actually slower than the naïve recursive algorithm, because it's wasting time solving subproblems that the recursive algorithm never considers. Dynamic programming isn't *always* an improvement![19]

3.9 The Woodcutter's Problem

The final problem we considered in the previous chapter was planning how to cut a long plank of wood into $n+1$ smaller boards at n predetermined locations. The input is an array $Len[1..n]$, where $Len[i]$ is the length the ith board. What is the minimum number of dollars we need to cut our given plank into the desired boards, assuming that it costs ℓ dollars to cut any plank of length ℓ once?

For any indices $i < k$, let $MinCost(i,k)$ denote the minimum total cost to decompose the subboard consisting of boards i through k. We derived the following recurrence for the function $MinCost$:

$$MinCost(i,k) = \begin{cases} 0 & \text{if } i = k \\ \displaystyle\sum_{j=i}^{k} Len[j] + \min_{i<r\le k} \left\{ \begin{array}{l} MinCost(i,r-1) \\ \quad + MinCost(r,k) \end{array} \right\} & \text{otherwise} \end{cases}$$

[18] Even though the subset sum problem is NP-hard, this time bound does *not* imply that P=NP, because T is not necessarily bounded by a polynomial function of the input size.

[19] In the 1967 research memorandum(!) where he proposed memo functions, Donald Michie wrote, "To tabulate values of a function which will not be needed is a waste of space, and to recompute the same values more than once is a waste of time." But in fact, tabulating values of a function that will not be needed is also a waste of *time*!

You can probably guess what we're going to do with this recurrence eventually, but let's rid of that ugly summation first.

For any pair of indices $i \le k$, let $TotLen(i,k)$ denote the total length of boards i through k:

$$TotLen(i,k) := \sum_{j=i}^{k} Len[j]$$

This function satisfies the following simple recurrence:

$$TotLen(i,k) = \begin{cases} Len[i] & \text{if } i = k \\ TotLen(i,k-1) + Len[k] & \text{otherwise} \end{cases}$$

We can compute all possible values of $TotLen(i,k)$ in $O(n^2)$ time using—you guessed it—dynamic programming! The usual mechanical steps give us the following dynamic programming algorithm:

```
InitTotLen(Len[1..n]):
  for i ← 1 to n
      TotLen[i, i−1] ← 0
      for k ← i to n
          TotLen[i,k] ← TotLen[i,k−1] + Len[k]
```

We will use this short algorithm as an initialization subroutine. This initialization allows us to simplify the original *MinCost* recurrence as follows:

$$MinCost(i,k) = \begin{cases} 0 & \text{if } i > k \\ TotLen[i,k] + \min\limits_{i<r\le k} \left\{ \begin{matrix} MinCost(i,r-1) \\ + MinCost(r,k) \end{matrix} \right\} & \text{otherwise} \end{cases}$$

Now let's turn the crank.

- **Subproblems:** The input to each subproblem is an *interval*, specified by two integers i and k, such that $1 \le i \le n+1$ and $0 \le k \le n$.
- **Memoization:** We can store all possible values of *MinCost* in a two-dimensional array $MinCost[1..n+1, 0..n]$. (Only the entries $MinCost[i,k]$ with $k \ge i-1$ will actually be used, but whatever.)
- **Dependencies:** Each entry $MinCost[i,k]$ depends on the entries $MinCost[i, r-1]$ and $MinCost[r+1,k]$, for all r such that $i < r \le k$. In other words, each table entry depends on all entries either directly to the left or directly below.

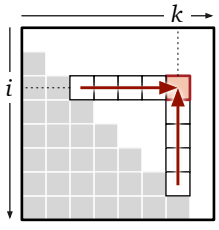

The following subroutine fills the single entry $MinCost[i,k]$, assuming all the entries it depends on have already been computed.

```
COMPUTEMINCOST(i,k):
  MinCost[i,k] ← ∞
  for r ← i+1 to k
      tmp ← MinCost[i,r−1] + MinCost[r+1,k]
      if MinCost[i,k] > tmp
          MinCost[i,k] ← tmp
  MinCost[i,k] ← MinCost[i,k] + TotLen[i,k]
```

- **Evaluation order:** There are at least three different orders that can be used to fill the array. The one that occurs to most students first is to scan through the table one diagonal at a time, starting with the trivial base cases $MinCost[i,i-1]$ and working toward the final answer $MinCost[1,n]$:

```
WOODCUTTER(Len[1..n]):
  INITTOTLEN(Len[1..n])
  for i ← 1 to n+1
      MinCost[i,i−1] ← 0
  for d ← 0 to n−1
      for i ← 1 to n−d        ⟨⟨...or whatever⟩⟩
          COMPUTEMINCOST(i,i+d)
  return MinCost[1,n]
```

But we could also traverse the array row by row from the bottom up, traversing each row from left to right, or column by column from left to right, traversing each columns from the bottom up.

```
WOODCUTTER2(Len[1..n]):
  INITTOTLEN(Len[1..n])
  for i ← n+1 down to 1
      MinCost[i,i−1] ← 0
      for j ← i to n
          COMPUTEMINCOST(i,j)
  return MinCost[1,n]
```

```
WOODCUTTER3(f[1..n]):
  INITTOTLEN(Len[1..n])
  for j ← 0 to n+1
      MinCost[j+1,j] ← 0
      for i ← j down to 1
          COMPUTEMINCOST(i,j)
  return MinCost[1,n]
```

Just as we did for our edit-distance algorithm, we can illustrate these evaluation orders using a double-lined arrow to indicate the outer loop and single-lined arrows to indicate the inner loop. The bidirectional arrows in

the first evaluation order indicate that the order of the inner loops doesn't matter.

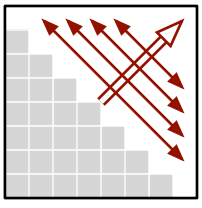
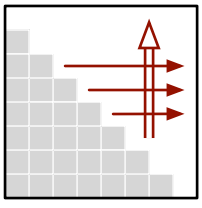
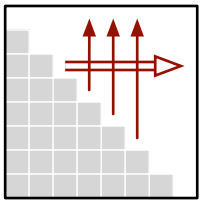

- **Time and space:** The memoization structure uses $O(n^2)$ space. No matter which evaluation order we choose, computing each entry $MinCost[i,k]$ takes $O(n)$ time, so our overall algorithm runs in $\boldsymbol{O(n^3)}$ ***time***.

As usual, we could have predicted the final space and time bounds directly from the original recurrence:

$$MinCost(i,k) = \begin{cases} 0 & \text{if } i > k \\ TotLen[i,k] + \min\limits_{i \le r \le k} \left\{ \begin{array}{l} MinCost(i, r-1) \\ \quad + MinCost(r+1, k) \end{array} \right\} & \text{otherwise} \end{cases}$$

The *MinCost* function has two arguments, each of which can take on roughly n different values, so we probably need a data structure of size $O(n^2)$. On the other hand, there are *three* variables in the body of the recurrence (i, k, and r), each of which can take roughly n different values, so it should take $O(n^3)$ time to compute everything.

3.10 Dynamic Programming on Trees

So far, all of our dynamic programming examples use multidimensional arrays to store the results of recursive subproblems. However, as the next example shows, this is not always the most appropriate data structure to use.

An ***independent set*** in a graph is a subset of the vertices with no edges between them. Finding the largest independent set in an arbitrary graph is extremely hard; in fact, this is one of the canonical NP-hard problems we will study in Chapter **??**. But in some special classes of graphs, we can find largest independent sets quickly. In particular, when the input graph is a *tree* with n vertices, we can actually compute the largest independent set in $O(n)$ time.

Suppose we are given a tree T. Without loss of generality, suppose T is a rooted tree; that is, there is a special node in T called the *root*, and all edges are implicitly directed away from this vertex. (If T is an unrooted tree—a connected acyclic undirected graph—we can choose an arbitrary vertex as the root.) We call vertex w a *descendant* of vertex v if the unique path from w to the root includes v; equivalently, the descendants of v are v itself and the descendants

of the children of v. The *subtree rooted at* v consists of all the descendants of v and the edges between them.

For any node v in T, let $MIS(v)$ denote the size of the largest independent set in the subtree rooted at v. Any independent set in this subtree that excludes v itself is the union of independent sets in the subtrees rooted at the children of v. On the other hand, any independent set that *includes* v necessarily excludes all of v's children, and therefore includes independent sets in the subtrees rooted at v's grandchildren. Thus, the function MIS obeys the following recurrence, where the nonstandard notation $w \downarrow v$ means "w is a child of v":

$$MIS(v) = \max\left\{\sum_{w\downarrow v} MIS(w),\ 1 + \sum_{w\downarrow v}\sum_{x\downarrow w} MIS(x)\right\}$$

We need to compute $MIS(r)$, where r is the root of T.

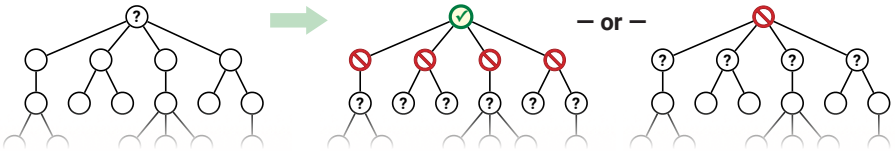

Figure 3.5. Computing the maximum independent set in a tree

What data structure should we use to memoize this recurrence? The most natural choice is ***the tree T itself!*** Specifically, for each vertex v in T, we store the result of $MIS(v)$ in a new field $v.MIS$. (In principle, we *could* use an array instead, but then we'd need pointers back and forth between each node and its corresponding array entry, so why bother?)

What's a good order to consider the subproblems? The subproblem associated with any node v depends on the subproblems associated with the children and grandchildren of v. So we can visit the nodes in any order we like, provided that every vertex is visited before its parent; in particular, we can use a standard ***post-order*** traversal.

What's the running time of the algorithm? The non-recursive time associated with each node v is proportional to the number of children and grandchildren of v; this number can be very different from one vertex to the next. But we can turn the analysis around: Each vertex contributes a constant amount of time to its parent and its grandparent! Because each vertex has at most one parent and at most one grandparent, the algorithm runs in $O(n)$ ***time***.

Here is the resulting dynamic programming algorithm. Yes, it's still recursive, because that's the most natural way to implement a post-order tree traversal.

```
TreeMIS(v):
  skipv ← 0
  for each child w of v
      skipv ← skipv + TreeMIS(w)
  keepv ← 1
  for each grandchild x of v
      keepv ← keepv + x.MIS
  v.MIS ← max{keepv, skipv}
  return v.MIS
```

We can derive an even simpler linear-time algorithm by defining two separate functions over the nodes of T:

- Let $MISyes(v)$ denote the size of the largest independent set of the subtree rooted at v that *includes* v.
- Let $MISno(v)$ denote the size of the largest independent set of the subtree rooted at v that *excludes* v.

Again, we need to compute $\max\{MISyes(r), MISno(r)\}$, where r is the root of T. The first two functions satisfy the following mutual recurrence:

$$MISyes(v) = 1 + \sum_{w \downarrow v} MISno(w)$$

$$MISno(v) = \sum_{w \downarrow v} \max\{MISyes(w), MISno(w)\}$$

Again, we can memoize these functions into the tree itself, by defining two new fields for each vertex. A straightforward post-order tree traversal evaluates both functions at every node in ***O(n) time***. The following algorithm not only memoizes both function values at v, it also returns the larger of those two values.

```
TreeMIS2(v):
  v.MISno ← 0
  v.MISyes ← 1
  for each child w of v
      v.MISno ← v.MISno + TreeMIS2(w)
      v.MISyes ← v.MISyes + w.MISno
  return max{v.MISyes, v.MISno}
```

In the second line of the inner loop, we are using the value $w.MISno$ that was memoized by the recursive call in the previous line.

Exercises

For all of the following exercises—and more generally when developing ***any*** new dynamic programming algorithm—I strongly recommend following the steps

outlined in Section 3.4. In particular, don't even *start* thinking about tables or for-loops until you have a complete recursive solution, including a clear English specification of the recursive subproblems you are actually solving. **First make it work, then make it fast.**

Sequences/Arrays

1. In a previous life, you worked as a cashier in the lost Antarctican colony of Nadiria, spending the better part of your day giving change to your customers. Because paper is a very rare and valuable resource in Antarctica, cashiers were required by law to use the fewest bills possible whenever they gave change. Thanks to the numerological predilections of one of its founders, the currency of Nadiria, called Dream-Dollars, was available in the following denominations: \$1, \$4, \$7, \$13, \$28, \$52, \$91, and \$365.[20]

   ♠(a) The greedy change algorithm repeatedly takes the largest bill that does not exceed the target amount. For example, to make \$122 using the greedy algorithm, we first take a \$91 bill, then a \$28 bill, and finally three \$1 bills. Give an example where this greedy algorithm uses more Dream-Dollar bills than the minimum possible. *[Hint: It may be easier to write a small program than to work this out by hand.]*

   (b) Describe and analyze a recursive algorithm that computes, given an integer k, the minimum number of bills needed to make k Dream-Dollars. (Don't worry about making your algorithm fast; just make sure it's correct.)

   (c) Describe a dynamic programming algorithm that computes, given an integer k, the minimum number of bills needed to make k Dream-Dollars. (This one needs to be fast.)

2. Describe efficient algorithms for the following variants of the text segmentation problem. Assume that you have a subroutine IsWord that takes an array of characters as input and returns True if and only if that string is a "word". Analyze your algorithms by bounding the number of calls to IsWord.

   (a) Given an array $A[1..n]$ of characters, compute the number of partitions of A into words. For example, given the string `ARTISTOIL`, your algorithm should return 2, for the partitions `ARTIST·OIL` and `ART·IS·TOIL`.

   (b) Given two arrays $A[1..n]$ and $B[1..n]$ of characters, decide whether A and B can be partitioned into words at the same indices. For example,

[20] For more details on the history and culture of Nadiria, including images of the various denominations of Dream-Dollars, see http://moneyart.biz/dd/.

the strings BOTHEARTHANDSATURNSPIN and PINSTARTRAPSANDRAGSLAP can be partitioned into words at the same indices as follows:

```
BOT·HEART·HAND·SAT·URNS·PIN
PIN·START·RAPS·AND·RAGS·LAP
```

(c) Given two arrays $A[1..n]$ and $B[1..n]$ of characters, compute the number of different ways that A and B can be partitioned into words at the same indices.

3. Suppose you are given an array $A[1..n]$ of numbers, which may be positive, negative, or zero, and which are ***not*** necessarily integers.

   (a) Describe and analyze an algorithm that finds the largest sum of elements in a contiguous subarray $A[i..j]$.

   (b) Describe and analyze an algorithm that finds the largest *product* of elements in a contiguous subarray $A[i..j]$.

   For example, given the array $[-6, 12, -7, 0, 14, -7, 5]$ as input, your first algorithm should return 19, and your second algorithm should return 504.

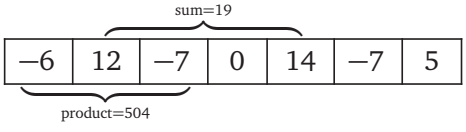

   Given the one-element array $[-374]$ as input, your first algorithm should return 0, and your second algorithm should return 1. (The empty interval is still an interval!) For the sake of analysis, assume that comparing, adding, or multiplying any pair of numbers takes $O(1)$ time.

   [Hint: Part (a) has been a standard computer science interview question since at least the mid-1980s. You can find many correct solutions on the web; the problem even has its own Wikipedia page! But at least in 2016, a significant fraction of the solutions I found on the web for part (b) were either slower than necessary or actually incorrect.]

4. This exercise explores variants of the maximum-subarray problem (Problem 3). In all cases, your input consists of an array $A[1..n]$ of real numbers (which could be positive, negative, or zero) and possibly an additional integer $X \geq 0$.

   (a) *Wrapping around:* Suppose A is a *circular* array. In this setting, a "contiguous subarray" can be either an interval $A[i..j]$ or a suffix followed by a prefix $A[i..n] \cdot A[1..j]$. Describe and analyze an algorithm that finds a contiguous subarray of A with the largest sum.

(b) *Long subarrays only:* Describe and analyze an algorithm that finds a contiguous subarray of A *of length at least* X that has the largest sum. (Assume $X \leq n$.)

(c) *Short subarrays only:* Describe and analyze an algorithm that finds a contiguous subarray of A *of length at most* X that has the largest sum.

(d) *The Price Is Right:* Describe and analyze an algorithm that finds a contiguous subarray of A with the largest sum *less than or equal to* X.

(e) Describe a faster algorithm for Problem 4(d) when every number in the array A is non-negative.

5. This exercise asks you to develop efficient algorithms to find optimal *subsequences* of various kinds. A subsequence is anything obtained from a sequence by extracting a subset of elements, but keeping them in the same order; the elements of the subsequence need not be contiguous in the original sequence. For example, the strings `C`, `DAMN`, `YAIOAI`, and `DYNAMICPROGRAMMING` are all subsequences of the string `DYNAMICPROGRAMMING`.

[Hint: Exactly one of these problems can be solved in $O(n)$ time using a greedy algorithm.]

(a) Let $A[1..m]$ and $B[1..n]$ be two arbitrary arrays. A ***common subsequence*** of A and B is another sequence that is a subsequence of both A and B. Describe an efficient algorithm to compute the length of the *longest* common subsequence of A and B.

(b) Let $A[1..m]$ and $B[1..n]$ be two arbitrary arrays. A ***common supersequence*** of A and B is another sequence that contains both A and B as subsequences. Describe an efficient algorithm to compute the length of the *shortest* common supersequence of A and B.

(c) Call a sequence $X[1..n]$ of numbers ***bitonic*** if there is an index i with $1 < i < n$, such that the prefix $X[1..i]$ is increasing and the suffix $X[i..n]$ is decreasing. Describe an efficient algorithm to compute the length of the longest bitonic subsequence of an arbitrary array A of integers.

(d) Call a sequence $X[1..n]$ of numbers ***oscillating*** if $X[i] < X[i+1]$ for all even i, and $X[i] > X[i+1]$ for all odd i. Describe an efficient algorithm to compute the length of the longest oscillating subsequence of an arbitrary array A of integers.

(e) Describe an efficient algorithm to compute the length of the shortest oscillating supersequence of an arbitrary array A of integers.

(f) Call a sequence $X[1..n]$ of numbers ***convex*** if $2 \cdot X[i] < X[i-1] + X[i+1]$ for all i. Describe an efficient algorithm to compute the length of the longest convex subsequence of an arbitrary array A of integers.

(g) Call a sequence $X[1..n]$ of numbers ***weakly increasing*** if each element is larger than the average of the two previous elements; that is, $2 \cdot X[i] > X[i-1] + X[i-2]$ for all $i > 2$. Describe an efficient algorithm to compute the length of the longest weakly increasing subsequence of an arbitrary array A of integers.

(h) Call a sequence $X[1..n]$ of numbers ***double-increasing*** if $X[i] > X[i-2]$ for all $i > 2$. (In other words, a double-increasing sequence is a perfect shuffle of two increasing sequences.) Describe an efficient algorithm to compute the length of the longest double-increasing subsequence of an arbitrary array A of integers.

(i) Recall that a sequence $X[1..n]$ of numbers is *increasing* if $X[i] < X[i+1]$ for all i. Describe an efficient algorithm to compute the length of the *longest common increasing subsequence* of two given arrays of integers. For example, $\langle 1, 4, 5, 6, 7, 9 \rangle$ is the longest common increasing subsequence of the sequences $\langle 3, 1, 4, 1, 5, 9, 2, 6, 5, 3, 5, 8, 9, 7, 9, 3 \rangle$ and $\langle 1, 4, 1, 4, 2, 1, 3, 5, 6, 2, 3, 7, 3, 0, 9, 5 \rangle$.

6. A *shuffle* of two strings X and Y is formed by interspersing the characters into a new string, keeping the characters of X and Y in the same order. For example, the string BANANAANANAS is a shuffle of the strings BANANA and ANANAS in several different ways.

$$\texttt{BANANA}_{\texttt{ANANAS}} \qquad \texttt{BAN}_{\texttt{ANA}}\texttt{ANA}_{\texttt{NAS}} \qquad \texttt{B}_{\texttt{AN}}\texttt{AN}_{\texttt{A}}\texttt{A}_{\texttt{NA}}\texttt{NA}_{\texttt{S}}$$

Similarly, the strings PRODGYRNAMAMMIINCG and DYPRONGARMAMMICING are both shuffles of DYNAMIC and PROGRAMMING:

$$\texttt{PRO}^{\texttt{D}}\texttt{G}^{\texttt{Y}}\texttt{R}^{\texttt{NAM}}_{\texttt{AMMI}}\texttt{I}^{\texttt{N}}\texttt{C}_{\texttt{G}} \qquad \texttt{DY}^{\texttt{PRO}}\texttt{N}_{\texttt{G}}\texttt{A}_{\texttt{R}}\texttt{M}_{\texttt{AMM}}\texttt{IC}_{\texttt{ING}}$$

(a) Given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, describe and analyze an algorithm to determine whether C is a shuffle of A and B.

(b) A ***smooth*** shuffle of X and Y is a shuffle of X and Y that never uses more than two consecutive symbols of either string. For example,

- $\texttt{PR}_{\texttt{D}}\texttt{O}_{\texttt{Y}}\texttt{G}_{\texttt{NA}}\texttt{RA}_{\texttt{M}}\texttt{M}^{\texttt{M}}\texttt{I}^{\texttt{I}}\texttt{C}^{\texttt{NG}}$ is a smooth shuffle of the strings DYNAMIC and PROGRAMMING.
- $\texttt{DY}^{\texttt{PR}}\texttt{N}^{\texttt{OGR}}\texttt{A}^{\texttt{A}}\texttt{M}^{\texttt{MM}}\texttt{IC}^{\texttt{ING}}$ is a shuffle of DYNAMIC and PROGRAMMING, but it is not a smooth shuffle (because of the substrings OGR and ING).
- $\texttt{XX}_{\texttt{X}}\texttt{X}_{\texttt{X}}\texttt{X}_{\texttt{XX}}\texttt{XX}_{\texttt{XX}}\texttt{XX}_{\texttt{X}}\texttt{X}_{\texttt{X}}\texttt{XX}$ is a smooth shuffle of the strings XXXXXXX and XXXXXXXXXXX.
- There is no smooth shuffle of the strings XXXX and XXXXXXXXXXXX.

Describe and analyze an algorithm to decide, given three strings X, Y, and Z, whether Z is a smooth shuffle of X and Y.

7. For each of the following problems, the input consists of two arrays $X[1..k]$ and $Y[1..n]$ where $k \le n$.

   (a) Describe and analyze an algorithm to decide whether X is a subsequence of Y. For example, the string PPAP is a subsequence of the string PENPINEAPPLEAPPLEPEN.

   (b) Describe and analyze an algorithm to find the smallest number of symbols that can be removed from Y so that X is no longer a subsequence. Equivalently, your algorithm should find the longest subsequence of Y that is *not* a supersequence of X. For example, after removing removing two symbols from the string PENPINEAPPLEAPPLEPEN, the string PPAP is no longer a subsequence.

   ♥(c) Describe and analyze an algorithm to determine whether X occurs as two *disjoint* subsequences of Y. For example, the string PPAP appears as two disjoint subsequences in the string PENPINEAPPLEAPPLEPEN.

   (d) Suppose the input also includes a third array $C[1..n]$ of numbers, which may be positive, negative, or zero, where $C[i]$ is the *cost* of $Y[i]$. Describe and analyze an algorithm to compute the minimum-cost occurrence of X as a subsequence of Y. That is, we want to find an array $I[1..k]$ such that $I[j] < I[j+1]$ and $X[I[j]] = Y[j]$ for every index j, and the total cost $\sum_{j=1}^{k} C[j]$ is as small as possible.

   (e) Describe and analyze an algorithm to compute the total number of (possibly overlapping) occurrences of X as a subsequence of Y. For purposes of analysis, assume that we can add two arbitrary integers in $O(1)$ time. For example, the string PPAP appears exactly 23 times as a subsequence of the string PENPINEAPPLEAPPLEPEN. If all characters in X and Y are equal, your algorithm should return $\binom{n}{k}$.

   (f) What is the running time of your algorithm for part (d) if adding two ℓ-bit integers requires $O(\ell)$ time?

8. Describe and analyze an efficient algorithm to find the length of the longest contiguous substring that appears both forward and backward in an input string $T[1..n]$. The forward and backward substrings must not overlap. Here are several examples:

   - Given the input string ALGORITHM, your algorithm should return 0.
   - Given the input string RECURSION, your algorithm should return 1, for the substring R.
   - Given the input string REDIVIDE, your algorithm should return 3, for the substring EDI. (The forward and backward substrings must not overlap!)

- Given the input string DYNAMICPROGRAMMINGMANYTIMES, your algorithm should return 4, for the substring YNAM. (In particular, it should *not* return 6, for the subsequence YNAMIR).

9. A palindrome is any string that is exactly the same as its reversal, like I, or DEED, or RACECAR, or AMANAPLANACATACANALPANAMA.

   (a) Describe and analyze an algorithm to find the length of the *longest subsequence* of a given string that is also a palindrome.

   For example, the longest palindrome subsequence of the string MAHDYNAMICPROGRAMZLETMESHOWYOUTHEM is MHYMRORMYHM; thus, given that string as input, your algorithm should return 11.

   (b) Describe and analyze an algorithm to find the length of the *shortest supersequence* of a given string that is also a palindrome. For example, the shortest palindrome supersequence of TWENTYONE is TWENTOYOTNEWT, so given the string TWENTYONE as input, your algorithm should return 13.

   (c) Any string can be decomposed into a sequence of palindromes. For example, the string BUBBASEESABANANA can be broken into palindromes in the following ways (and 65 others):

   BUB • BASEESAB • ANANA

   B • U • BB • ASEESA • B • ANANA

   BUB • B • A • SEES • ABA • N • ANA

   B • U • BB • A • S • EE • S • A • B • A • NAN • A

   B • U • B • B • A • S • E • E • S • A • B • A • N • A • N • A

   Describe and analyze an efficient algorithm to find the smallest number of palindromes that make up a given input string. For example, given the input string BUBBASEESABANANA, your algorithm should return 3.

   (d) Describe and analyze an efficient algorithm to find the largest integer k such that a given string can be split into palindromes of length at least k. For example:

   - Given the string PALINDROME, your algorithm should return 1.
   - Given the string BUBBASEESABANANA, your algorithm should return 3, for the partition BUB • BASEESAB • ANANA.
   - Given a string of n identical symbols, your algorithm should return n.

   (e) Describe and analyze an efficient algorithm to find the number of different ways that a given string can be decomposed into palindromes. For example:

   - Given the string PALINDROME, your algorithm should return 1.

- Given the string BUBBASEESABANANA, your algorithm should return 70.
- Given a string of n identical symbols, your algorithm should return 2^{n-1}.

(f) Describe and analyze an efficient algorithm to find the minimum number of palindromes that *cover* a given string. For example, the palindromes ABCBA, CBAABC, and CAC cover the string ABCBAABCAC.

♥(g) A ***metapalindrome*** is a decomposition of a string into a sequence of palindromes, such that the sequence of palindrome lengths is itself a palindrome. For example:

BOB • S • MAM • ASEESA • UKU • L • ELE

is a metapalindrome for the string BOBSMAMASEESAUKULELE, whose length sequence is the palindrome $(3,1,3,6,3,1,3)$. Describe and analyze an efficient algorithm to find the length of the shortest metapalindrome for a given string. For example, given the input string BOBSMAMASEESAUKULELE, your algorithm should return 11.

10. Suppose you are given an array $A[1..n]$ of positive integers. An *increasing back-and-forth subsequence* is an sequence of indices $I[1..\ell]$ with the following properties:

- $1 \le I[j] \le n$ for all j.
- $A[I[j]] < A[I[j+1]]$ for all $j < \ell$.
- If $I[j]$ is even, then $I[j+1] > I[j]$.
- If $I[j]$ is odd, then $I[j+1] < I[j]$.

Less formally, suppose we are given an array of n squares, each containing a positive integer. Suppose we place a token on one of the squares, and then repeatedly move the token left (if it's on an odd-indexed square) or right (if it's on an even-indexed square), always moving from a smaller number to a larger number. Then the sequence of token positions is an increasing back-and-forth subsequence.

Describe an algorithm to compute the length of the longest increasing back-and-forth subsequence of a given array of n integers. For example, given the input array

1	1	8	7	5	6	3	6	4	4	8	3	9	1	2	2	3	9	4	0
1<	2>	3<	4>	5<	6>	7<	8>	9<	10>	11<	12>	13<	14>	15<	16>	17<	18>	19<	20>

your algorithm should return the integer 9, which is the length of the following increasing back-and-forth subsequence:

0	1	2	3	4	6	7	8	9
20>	1<	15<	18>	10>	6>	4>	3<	13<

11. Suppose we want to typeset a paragraph of text onto a piece of paper (or if you insist, a computer screen). The text consists of a sequence of n words, where the ith word has length $\ell[i]$. We want to break the paragraph into several lines of total length exactly L. For example, according to TEX, the program used to typeset these notes, *the paragraph you are reading right now* is approximately 11.94794 cm ≈ 4.7055 inches wide.

    Depending on how the paragraph is broken into lines of text, we must insert different amounts of white space between the words. The paragraph should be fully justified, meaning that the first character on each line starts at the left margin, and *except for the last line*, the last character on each line ends at the right margin. There must be at least one unit of white space between any two words on the same line. See *the paragraph you are reading right now*? Just like that.

    Define the *slop* of a paragraph layout as the sum over all lines, *except the last*, of the cube of the amount of extra white-space in each line, not counting the one unit of required space between each adjacent pair of words. Specifically, if a line contains words i through j, then the slop of that line is defined to be $\left(L - j + i - \sum_{k=i}^{j} \ell[k]\right)^3$. Describe a dynamic programming algorithm to print the paragraph with minimum slop.

12. You and your eight-year-old nephew Elmo decide to play a simple card game. At the beginning of the game, the cards are dealt face up in a long row. Each card is worth a different number of points. After all the cards are dealt, you and Elmo take turns removing either the leftmost or rightmost card from the row, until all the cards are gone. At each turn, you can decide which of the two cards to take. The winner of the game is the player that has collected the most points when the game ends.

    Having never taken an algorithms class, Elmo follows the obvious greedy strategy—when it's his turn, Elmo *always* takes the card with the higher point value. Your task is to find a strategy that will beat Elmo whenever possible. (It might seem mean to beat up on a little kid like this, but Elmo absolutely *hates* it when grown-ups let him win.)

    (a) Prove that you should not also use the greedy strategy. That is, show that there is a game that you can win, but only if you do *not* follow the same greedy strategy as Elmo.

    (b) Describe and analyze an algorithm to determine, given the initial sequence of cards, the maximum number of points that you can collect playing against Elmo.

♣(c) When Elmo was four, he used an even simpler strategy—on his turn, he always chose his next card uniformly at random. That is, if there was more than one card left on his turn, he would take the leftmost card with probability 1/2, and the rightmost card with probability 1/2. Describe an algorithm to determine, given the initial sequence of cards, the maximum *expected* number of points you can collect playing against four-year-old-Elmo.

(d) Five years later, thirteen-year-old Elmo has become a *much* stronger player. Describe and analyze an algorithm to determine, given the initial sequence of cards, the maximum number of points that you can collect playing against a *perfect* opponent.

13. It's almost time to show off your flippin' sweet dancing skills! Tomorrow is the big dance contest you've been training for your entire life, except for that summer you spent with your uncle in Alaska hunting wolverines. You've obtained an advance copy of the list of n songs that the judges will play during the contest, in chronological order. Yesssssssssss!

You know all the songs, all the judges, and your own dancing ability extremely well. For each integer k, you know that if you dance to the kth song on the schedule, you will be awarded exactly $Score[k]$ points, but then you will be physically unable to dance for the next $Wait[k]$ songs (that is, you cannot dance to songs $k+1$ through $k+Wait[k]$). The dancer with the highest total score at the end of the night wins the contest, so you want your total score to be as high as possible.

Describe and analyze an efficient algorithm to compute the maximum total score you can achieve. The input to your sweet algorithm is the pair of arrays $Score[1..n]$ and $Wait[1..n]$.

14. The new swap-puzzle game *Candy Swap Saga XIII* involves n cute animals numbered from 1 to n. Each animal holds one of three types of candy: circus peanuts, Heath bars, and Cioccolateria Gardini chocolate truffles. You also have a candy in your hand; at the start of the game, you have a circus peanut.

To earn points, you visit each of the animals in order from 1 to n. For each animal, you can either keep the candy in your hand or exchange it with the candy the animal is holding.

- If you swap your candy for another candy of the *same* type, you earn one point.
- If you swap your candy for a candy of a *different* type, you lose one point. (Yes, your score can be negative.)

- If you visit an animal and decide not to swap candy, your score does not change.

You *must* visit the animals in order, and once you visit an animal, you can never visit it again.

Describe and analyze an efficient algorithm to compute your maximum possible score. Your input is an array $C[1..n]$, where $C[i]$ is the type of candy that the ith animal is holding.

15. Lenny Rutenbar, the founding dean of the new Maksymilian R. Levchin College of Computer Science, has commissioned a series of snow ramps on the south slope of the Orchard Downs sledding hill[21] and challenged Bill Kudeki, head of the Department of Electrical and Computer Engineering, to a sledding contest. Bill and Lenny will both sled down the hill, each trying to maximize their air time. The winner gets to expand their department/college into both Siebel Center and the new ECE Building; the loser has to move their entire department/college in the Boneyard culvert next to Loomis Lab.

Whenever Lenny or Bill reaches a ramp *while on the ground*, they can either use that ramp to jump through the air, possibly flying over one or more ramps, or sled past that ramp and stay on the ground. Obviously, if someone flies over a ramp, they cannot use that ramp to extend their jump.

(a) Suppose you are given a pair of arrays $Ramp[1..n]$ and $Length[1..n]$, where $Ramp[i]$ is the distance from the top of the hill to the ith ramp, and $Length[i]$ is the distance that any sledder who takes the ith ramp will travel through the air. Describe and analyze an algorithm to determine the maximum total distance that Lenny or Bill can spend in the air.

(b) The university lawyers heard about Lenny and Bill's little bet and immediately objected. To protect the university from either lawsuits or sky-rocketing insurance rates, they impose an upper bound on the number of jumps that either sledder can take. Describe and analyze an algorithm to determine the maximum total distance that Lenny or Bill can spend in the air *with at most k jumps*, given the original arrays $Ramp[1..n]$ and $Length[1..n]$ and the integer k as input.

♥(c) When the lawyers realized that imposing their restriction didn't immediately shut down the contest, they added a new restriction: No ramp can be used more than once! Disgusted by the legal interference, Lenny and Bill give up on their bet and decide to cooperate to put on a good show for the spectators. Describe and analyze an algorithm to determine the maximum total distance that Lenny and Bill can spend in the air, each

[21] The north slope is faster, but too short for an interesting contest.

taking at most k jumps (so at most $2k$ jumps total), and with each ramp used at most once.

16. Farmers Boggis, Bunce, and Bean have set up an obstacle course for Mr. Fox. The course consists of a long row of booths, each with a number painted on the front with bright red paint. Formally, Mr. Fox is given an array $A[1..n]$, where $A[i]$ is the number painted on the front of the ith booth. Each number $A[i]$ could be positive, negative, or zero. Everyone agrees with the following rules:

    - At each booth, Mr. Fox ***must*** say either "Ring!" or "Ding!".
    - If Mr. Fox says "Ring!" at the ith booth, he earns a reward of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox pays a penalty of $-A[i]$ chickens.)
    - If Mr. Fox says "Ding!" at the ith booth, he pays a penalty of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox earns a reward of $-A[i]$ chickens.)
    - Mr. Fox is forbidden to say the same word more than three times in a row. For example, if he says "Ring!" at booths 6, 7, and 8, then he must say "Ding!" at booth 9.
    - All accounts will be settled at the end, after Mr. Fox visits every booth and the umpire calls "Hot box!" Mr. Fox does not actually have to carry chickens through the obstacle course.
    - Finally, if Mr. Fox violates any of the rules, or if he ends the obstacle course owing the farmers chickens, the farmers will shoot him.

    Describe and analyze an algorithm to compute, the largest number of chickens that Mr. Fox can earn by running the obstacle course, given the array $A[1..n]$ of numbers as input. *[Hint: Watch out for the burning pine cone!]*

17. ***Dance Dance Revolution*** is a dance video game, first introduced in Japan by Konami in 1998. Players stand on a platform marked with four arrows, pointing forward, back, left, and right, arranged in a cross pattern. During play, the game plays a song and scrolls a sequence of n arrows (←, ↑, ↓, or →) from the bottom to the top of the screen. At the precise moment each arrow reaches the top of the screen, the player must step on the corresponding arrow on the dance platform. (The arrows are timed so that you'll step with the beat of the song.)

    You are playing a variant of this game called "Vogue Vogue Revolution", where the goal is to play perfectly but move as little as possible. When an arrow reaches the top of the screen, if one of your feet is already on the correct arrow, you are awarded one style point for maintaining your current pose. If neither foot is on the right arrow, you must move one (and *only* one)

foot from its current location to the correct arrow on the platform. If you ever step on the wrong arrow, or fail to step on the correct arrow, or move more than one foot at a time, or move either foot when you are already standing on the correct arrow, all your style points are taken away and you lose the game.

How should you move your feet to maximize your total number of style points? For purposes of this problem, assume you always start with your left foot on ←, and your right foot on →, and that you've memorized the entire sequence of arrows. For example, if the sequence is ↑↑↓↓←→←→, you can earn 5 style points by moving your feet as shown below:

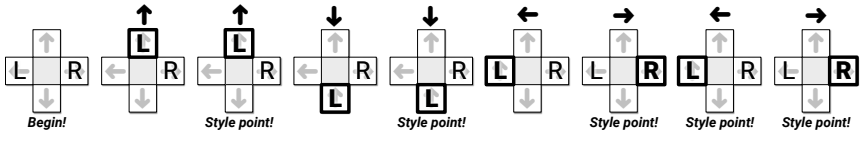

(a) ***Prove*** that for *any* sequence of n arrows, it is possible to earn at least $n/4 - 1$ style points.

(b) Describe an efficient algorithm to find the maximum number of style points you can earn during a given VVR routine. The input to your algorithm is an array $Arrow[1..n]$ containing the sequence of arrows.

18. Consider the following solitaire form of Scrabble. We begin with a fixed, finite sequence of tiles; each tile has both a letter and a numerical value. At the start of the game, we draw the first seven tiles from the sequence and put them into our hand. In each turn, we form an English word from some or all of the tiles in our hand, place those tiles on the table, and receive the total value of those tiles as points. (If no English word can be formed from the tiles in our hand, the game immediately ends.) Then we repeatedly draw the next tile from the start of the sequence until either (a) we have seven tiles in our hand, or (b) the sequence is empty. (Sorry, no double/triple word/letter scores, bingos, blanks, or passing.) Our goal is to obtain as many points as possible.

For example, consider the following sequence of 20 tiles:

I_2	N_2	X_8	A_1	N_2	A_1	D_3	U_5	D_3	I_2	D_3	K_8	U_5	B_4	L_2	A_1	K_8	H_5	A_1	N_2

Given this sequence of tiles at the beginning of the game, we can earn 68 points as follows:

- We initially draw I_2 N_2 X_8 A_1 N_2 A_1 D_3.
- Play the word N_2 A_1 I_2 A_1 D_3 for 9 points, leaving N_2 X_8 in hand.
- Draw the next five tiles U_5 D_3 I_2 D_3 K_8.

- Play the word $\boxed{\mathtt{U}_5}\boxed{\mathtt{N}_2}\boxed{\mathtt{D}_3}\boxed{\mathtt{I}_2}\boxed{\mathtt{D}_3}$ for 15 points, leaving $\boxed{\mathtt{K}_8}\boxed{\mathtt{X}_8}$ in hand.
- Draw the next five tiles $\boxed{\mathtt{U}_5}\boxed{\mathtt{B}_4}\boxed{\mathtt{L}_2}\boxed{\mathtt{A}_1}\boxed{\mathtt{K}_8}$.
- Play the word $\boxed{\mathtt{B}_4}\boxed{\mathtt{U}_5}\boxed{\mathtt{L}_2}\boxed{\mathtt{K}_8}$ for 19 points, leaving $\boxed{\mathtt{K}_8}\boxed{\mathtt{X}_8}\boxed{\mathtt{A}_1}$ in hand.
- Draw the last three tiles $\boxed{\mathtt{H}_5}\boxed{\mathtt{A}_1}\boxed{\mathtt{N}_2}$.
- Play the word $\boxed{\mathtt{A}_1}\boxed{\mathtt{N}_2}\boxed{\mathtt{K}_8}\boxed{\mathtt{H}_5}$ for 16 points, leaving $\boxed{\mathtt{X}_8}\boxed{\mathtt{A}_1}$ in hand.
- Play the word $\boxed{\mathtt{A}_1}\boxed{\mathtt{X}_8}$ for 9 points, emptying our hand and ending the game.

(a) Suppose the sequence of tiles is represented by two arrays $Letter[1..n]$, containing a sequence of letters between A and Z, and $Value[\mathtt{A}..\mathtt{Z}]$, where $Value[\ell]$ is the value of any tile with letter ℓ. Design and analyze an efficient algorithm to compute the maximum number of points that can be earned from the given sequence of tiles.

(b) Now suppose two tiles with the same letter might have different values. Now the tile sequence is represented by two arrays $Letter[1..n]$ and $Value[1..n]$, where $Value[i]$ is the value of the ith tile. Design and analyze an efficient algorithm to compute the maximum number of points that can be earned from the given sequence of tiles.

In both problems, the output is a single number: the maximum possible score. Assume (because it's true) that you can find all English words that can be made from any set of at most seven tiles, along with the point values of those words, in $O(1)$ time.

19. Suppose we are given a set L of n line segments in the plane, where each segment has one endpoint on the line $y = 0$ and one endpoint on the line $y = 1$, and all $2n$ endpoints are distinct.

(a) Describe and analyze an algorithm to compute the largest subset of L in which no pair of segments intersects.

(b) Describe and analyze an algorithm to compute the largest subset of L in which ***every*** pair of segments intersects.

Now suppose we are given a set L of n line segments in the plane, where both endpoints of each segment lie on the unit circle $x^2 + y^2 = 1$, and all $2n$ endpoints are distinct.

(c) Describe and analyze an algorithm to compute the largest subset of L in which no pair of segments intersects.

(d) Describe and analyze an algorithm to compute the largest subset of L in which ***every*** pair of segments intersects.

20. Let P be a set of n points evenly distributed on the unit circle, and let S be a set of m line segments with endpoints in P. The endpoints of the m

segments are *not* necessarily distinct; n could be significantly smaller than $2m$.

(a) Describe an algorithm to find the size of the largest subset of segments in S such that every pair is disjoint. Two segments are disjoint if they do not intersect even at their endpoints.

(b) Describe an algorithm to find the size of the largest subset of segments in S such that every pair is interior-disjoint. Two segments are interior-disjoint if their intersection is either empty or an endpoint of both segments.

(c) Describe an algorithm to find the size of the largest subset of segments in S such that every pair intersects.

(d) Describe an algorithm to find the size of the largest subset of segments in S such that every pair crosses. Two segments cross if they intersect but not at their endpoints.

For full credit, all four algorithms should run in $O(mn)$ time.

21. You are driving a bus along a highway, full of rowdy, hyper, thirsty students and a soda fountain machine. Each minute that a student is on your bus, that student drinks one ounce of soda. Your goal is to drop the students off quickly, so that the total amount of soda consumed by all students is as small as possible.

    You know how many students will get off of the bus at each exit. Your bus begins somewhere along the highway (probably not at either end) and moves at a constant speed of 37.4 miles per hour. You must drive the bus along the highway; however, you may drive forward to one exit then backward to an exit in the opposite direction, switching as often as you like. (You can stop the bus, drop off students, and turn around instantaneously.)

    Describe an efficient algorithm to drop the students off so that they drink as little soda as possible. Your input consists of the bus route (a list of the exits, together with the travel time between successive exits), the number of students you will drop off at each exit, and the current location of your bus (which you may assume is an exit).

22. Let's define a *summary* of two strings A and B to be a concatenation of substrings of the following form:

    - ▲SNA indicates a substring SNA of only the first string A.
    - ♦FOO indicates a common substring FOO of both strings.
    - ▼BAR indicates a substring BAR of only the second string B.

A summary is *valid* if we can recover the original strings A and B by concatenating the appropriate substrings of the summary in order and discarding the delimiters ▲, ♦, and ▼. Each regular character has length 1, and each delimiter ▲, ♦, or ▼ has some fixed non-negative length Δ. The *length* of a summary is the sum of the lengths of its symbols.

For example, each of the following strings is a valid summary of the strings KITTEN and KNITTING:

- ♦K▼N♦ITT▲E▼I♦N▼G has length $9 + 7\Delta$.
- ♦K▼N♦ITT▲EN▼ING has length $10 + 5\Delta$.
- ♦K▲ITTEN▼NITTING has length $13 + 3\Delta$.
- ▲KITTEN▼KNITTING has length $14 + 2\Delta$.

Describe and analyze an algorithm that computes the length of the shortest summary of two given strings $A[1..m]$ and $B[1..n]$. The delimiter length Δ is also part of the input to your algorithm. For example:

- Given strings KITTEN and KNITTING and $\Delta = 0$, your algorithm should return 9.
- Given strings KITTEN and KNITTING and $\Delta = 1$, your algorithm should return 15.
- Given strings KITTEN and KNITTING and $\Delta = 2$, your algorithm should return 18.

23. *Vankin's Mile* is a solitaire game played on an $n \times n$ square grid. The player starts by placing a token on any square of the grid. Then on each turn, the player moves the token either one square to the right or one square down. The game ends when player moves the token off the edge of the board. Each square of the grid has a numerical value, which could be positive, negative, or zero. The player starts with a score of zero; whenever the token lands on a square, the player adds its value to his score. The object of the game is to score as many points as possible.

For example, given the grid below, the player can score $8-6+7-3+4 = 10$ points by placing the initial token on the 8 in the second row, and then moving down, down, right, down, down. (This is *not* the best possible score for this grid of numbers.)

−1	7	−8	10	−5
−4	−9	**8** ⇓	−6	0
5	−2	**−6** ⇓	−6	7
−7	4	**7** ⇒	**−3** ⇓	−3
7	1	−6	**4** ⇓	−9

(a) Describe and analyze an efficient algorithm to compute the maximum possible score for a game of Vankin's Mile, given the $n \times n$ array of values as input.

(b) In the European version of this game, appropriately called *Vankin's Kilometer*, the player can move the token either one square down, one square right, *or one square left* in each turn. However, to prevent infinite scores, the token cannot land on the same square more than once. Describe and analyze an efficient algorithm to compute the maximum possible score for a game of Vankin's Kilometer, given the $n \times n$ array of values as input.[22]

(c) A variant called *Vankin's Niknav* adds an additional constraint to Vankin's Mile: The sequence of values that the token touches must be a palindrome. Describe and analyze an efficient algorithm to compute the maximum possible score for an instance of Vankin's Niknav, given the $n \times n$ array of values as input.

24. Suppose you are given an $m \times n$ bitmap as an array $M[1..n, 1..n]$ of 0s and 1s. A *solid block* in M is a subarray of the form $M[i..i', j..j']$ in which all bits are equal. A solid block is square if it has the same number of rows and columns.

(a) Describe an algorithm to find the maximum area of a solid *square* block in M in $O(n^2)$ time.

(b) Describe an algorithm to find the maximum area of a solid block in M in $O(n^3)$ time.

(c) Describe an algorithm to find the maximum area of a solid block in M in $O(n^2 \log n)$ time. *[Hint: Divide and conquer.]*

♥(d) Describe an algorithm to find the maximum area of a solid block in M in $O(n^2)$ time.

25. Suppose you are given an array $M[1..n, 1..n]$ of numbers, which may be positive, negative, or zero, and which are *not* necessarily integers. Describe an algorithm to find the largest sum of elements in any rectangular subarray of the form $M[i..i', j..j']$. For full credit, your algorithm should run in $O(n^3)$ time. *[Hint: See problem 3.]*

26. Describe and analyze an algorithm that finds the maximum-area rectangular pattern that appears more than once in a given bitmap. Specifically, given a two-dimensional array $M[1..n, 1..n]$ of bits as input, your algorithm

[22] If we also allowed *upward* movement, the resulting game (Vankin's Fathom?) would be NP-hard.

should output the area of the largest repeated rectangular pattern in M. For example, given the bitmap shown on the left in the figure below, your algorithm should return the integer 195, which is the area of the 15×13 doggo. (Although it doesn't happen in this example, the two copies of the repeated pattern might overlap.)

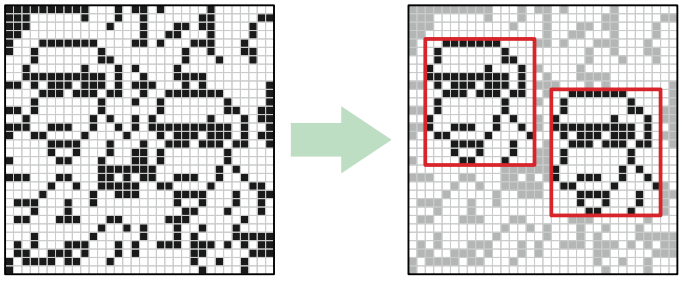

(a) For full credit, describe an algorithm that runs in $O(n^5)$ time.

♥(b) For extra credit, describe an algorithm that runs in $O(n^4)$ time.

♣♥(c) For extra extra credit, describe an algorithm that runs in $O(n^3 \operatorname{polylog} n)$ time.

27. Let P be a set of points in the plane in *convex position*. Intuitively, if a rubber band were wrapped around the points, then every point would touch the rubber band. More formally, for any point p in P, there is a line that separates p from the other points in P. Moreover, suppose the points are indexed $P[1], P[2], \ldots, P[n]$ in counterclockwise order around the "rubber band", starting with the leftmost point $P[1]$.

This problem asks you to solve a special case of the traveling salesman problem, where the salesman must visit every point in P, and the cost of moving from one point $p \in P$ to another point $q \in P$ is the Euclidean distance $|pq|$.

(a) Describe a simple algorithm to compute the shortest *cyclic* tour of P.

(b) A *simple* tour is one that never crosses itself. Prove that the shortest tour of P must be simple.

(c) Describe and analyze an efficient algorithm to compute the shortest tour of P that starts at the leftmost point $P[1]$ and ends at the rightmost point $P[r]$.

(d) Describe and analyze an efficient algorithm to compute the shortest tour of P, with no restrictions on the endpoints.

♥28. Describe and analyze an algorithm to solve the traveling salesman problem in $O(2^n \operatorname{poly}(n))$ time. Given an undirected n-vertex graph G with weighted edges, your algorithm should return the weight of the lightest cycle in G

that visits every vertex exactly once, or ∞ if G has no such cycles. *[Hint: The obvious recursive backtracking algorithm takes $O(n!)$ time.]*

29. Let $W = \{w_1, w_2, \ldots, w_n\}$ be a finite set of strings over some fixed alphabet Σ. An *edit center* for W is a string $C \in \Sigma^*$ such that the maximum edit distance from C to any string in W is as small as possible. The *edit radius* of W is the maximum edit distance from an edit center to a string in W. A set of strings may have several edit centers, but its edit radius is unique.

$$\mathit{EditRadius}(W) := \min_{C \in \Sigma^*} \max_{w \in W} \mathit{Edit}(w, C)$$

$$\mathit{EditCenter}(W) := \arg\min_{C \in \Sigma^*} \max_{w \in W} \mathit{Edit}(w, C)$$

(a) Describe and analyze an efficient algorithm to compute the edit radius of three given strings.

♣♥(b) Describe and analyze an efficient algorithm to approximate the edit radius of an arbitrary set of strings within a factor of 2. (Computing the *exact* edit radius is NP-hard unless the number of strings is fixed.)

♥30. Let $D[1..n]$ be an array of digits, each an integer between 0 and 9. A ***digital subsequence*** of D is a sequence of positive integers composed in the usual way from disjoint substrings of D. For example, the sequence $3, 4, 5, 6, 8, 9, 32, 38, 46, 64, 83, 279$ is a digital subsequence of the first several digits of π:

$$\underline{3}, 1, \underline{4}, 1, \underline{5}, 9, 2, \underline{6}, 5, 3, 5, \underline{8}, \underline{9}, 7, 9, \underline{3,2}, \underline{3,8}, \underline{4,6}, 2, \underline{6,4}, 3, 3, \underline{8,3}, \underline{2,7,9}$$

The *length* of a digital subsequence is the number of integers it contains, *not* the number of digits; the preceding example has length 12. As usual, a digital subsequence is ***increasing*** if each number is larger than its predecessor.

Describe and analyze an efficient algorithm to compute the longest increasing digital subsequence of D. *[Hint: Be careful about your computational assumptions. How long does it take to compare two k-digit numbers?]*

For full credit, your algorithm should run in $O(n^4)$ time; faster algorithms are worth extra credit. The fastest algorithm I know for this problem runs in $O(n^{3/2} \log n)$ time; achieving this bound requires several tricks, both in the design of the algorithm and in its analysis, but nothing outside the scope of this class.[23]

[23] With more advanced techniques, I believe the running time can be reduced to $O(n^{3/2} \log\log n)$, but I haven't worked through the details.

♥31. Consider the following variant of the classical Tower of Hanoi problem. As usual, there are n disks with distinct sizes, placed on three pegs numbered 0, 1, and 2. Initially, all n disks are on peg 0, sorted by size from smallest on top to largest on bottom. Our goal is to move all the disks to peg 2. In a single step, we can move the highest disk on any peg to a different peg, provided we satisfy two constraints. First, we must never place a smaller disk on top of a larger disk. Second—and this is the non-standard part—***we must never move a disk directly from peg 0 to peg 2.***

Describe and analyze an algorithm to compute the exact number of moves required to move all n disks from peg 0 to peg 2, subject to the stated restrictions. For full credit, your algorithm should use only $O(\log n)$ arithmetic operations in the worst case. For the sake of analysis, assume that adding or multiplying two k-digit numbers requires $O(k)$ time. *[Hint: Matrices!]*

Splitting Sequences/Arrays

32. A ***basic arithmetic expression*** is composed of characters from the set $\{1, +, \times\}$ and parentheses. Almost every integer can be represented by more than one basic arithmetic expression. For example, all of the following basic arithmetic expression represent the integer 14:

$$1+1+1+1+1+1+1+1+1+1+1+1+1+1+1+1$$
$$((1+1)\times(1+1+1+1+1))+((1+1)\times(1+1))$$
$$(1+1)\times(1+1+1+1+1+1+1)$$
$$(1+1)\times(((1+1+1)\times(1+1))+1)$$

Describe and analyze an algorithm to compute, given an integer n as input, the minimum number of 1s in a basic arithmetic expression whose value is equal to n. The number of parentheses doesn't matter, just the number of 1s. For example, when $n = 14$, your algorithm should return 8, for the final expression above. The running time of your algorithm should be bounded by a small polynomial function of n.

33. Suppose you are given a sequence of integers separated by $+$ and $-$ signs; for example:

$$1+3-2-5+1-6+7$$

You can change the value of this expression by adding parentheses in different places. For example:

$$1+3-2-5+1-6+7=-1$$

$$(1+3-(2-5))+(1-6)+7=9$$
$$(1+(3-2))-(5+1)-(6+7)=-17$$

Describe and analyze an algorithm to compute, given a list of integers separated by $+$ and $-$ signs, the maximum possible value the expression can take by adding parentheses. Parentheses must be used only to group additions and subtractions; in particular, do not use them to create implicit multiplication as in $1+3(-2)(-5)+1-6+7=33$.

34. Suppose you are given a sequence of integers separated by $+$ and $\times$ signs; for example:

$$1+3\times2\times0+1\times6+7$$

You can change the value of this expression by adding parentheses in different places. For example:

$$(1+(3\times2))\times0+(1\times6)+7=13$$
$$((1+(3\times2\times0)+1)\times6)+7=19$$
$$(1+3)\times2\times(0+1)\times(6+7)=104$$

(a) Describe and analyze an algorithm to compute the maximum possible value the given expression can take by adding parentheses, assuming all integers in the input are *positive*. *[Hint: This is easy.]*

(b) Describe and analyze an algorithm to compute the maximum possible value the given expression can take by adding parentheses, assuming all integers in the input are *non-negative*.

(c) Describe and analyze an algorithm to compute the maximum possible value the given expression can take by adding parentheses, with no restrictions on the input numbers.

Assume any arithmetic operation takes $O(1)$ time.

35. After graduating from Sham-Poobanana University, you decide to interview for a position at the Wall Street bank ***Long Live Boole***. The managing director of the bank, Eloob Egroeg, poses a 'solve-or-die' problems to each new employee, which they must solve within 24 hours. Those who fail to solve the problem are fired immediately!

Entering the bank for the first time, you notice that the employee offices are organized in a straight row, with a large T or F printed on the door of each office. Furthermore, between each adjacent pair of offices, there is a board marked by one of the symbols $\wedge, \vee$, or $\oplus$. When you ask about these arcane symbols, Eloob confirms that T and F represent the boolean values TRUE and FALSE, and the symbols on the boards represent the standard

boolean operators And, Or, and Xor. He also explains that these letters and symbols describe whether certain combinations of employees can work together successfully. At the start of any new project, Eloob hierarchically clusters his employees by adding parentheses to the sequence of symbols, to obtain an unambiguous boolean expression. The project is successful if this parenthesized boolean expression evaluates to T.

For example, if the bank has three employees, and the sequence of symbols on and between their doors is $T \wedge F \oplus T$, there is exactly one successful parenthesization scheme: $(T \wedge (F \oplus T))$. However, if the list of door symbols is $F \wedge T \oplus F$, there is no way to add parentheses to make the project successful.

Eloob finally poses your solve-or-die interview question: Describe an algorithm to decide whether a given sequence of symbols can be parenthesized so that the resulting boolean expression evaluates to T. Your input is an array $S[0..2n]$, where $S[i] \in \{T, F\}$ when i is even, and $S[i] \in \{\vee, \wedge, \oplus\}$ when i is odd.

36. Every year, as part of its annual meeting, the Antarctican Snail Lovers of Upper Glacierville hold a Round Table Mating Race. Several high-quality breeding snails are placed at the edge of a round table. The snails are numbered in order around the table from 1 to n. During the race, each snail wanders around the table, leaving a trail of slime behind it. The snails have been specially trained never to fall off the edge of the table or to cross a slime trail, even their own. If two snails meet, they are declared a breeding pair, removed from the table, and whisked away to a romantic hole in the ground to make little baby snails. Note that some snails may never find a mate, even if the race goes on forever.

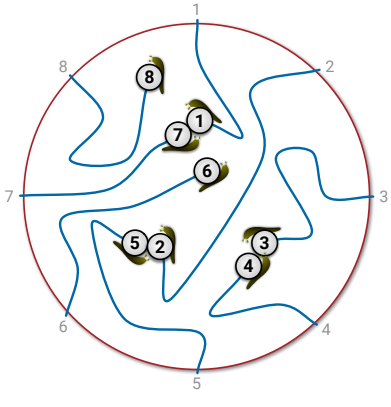

Figure 3.6. The end of a typical Antarctican SLUG race. Snails 6 and 8 never find mates. The organizers must pay $M[3,4] + M[2,5] + M[1,7]$.

For every pair of snails, the Antarctican SLUG race organizers have posted a monetary reward, to be paid to the owners if that pair of snails meets during the Mating Race. Specifically, there is a two-dimensional array $M[1..n, 1..n]$ posted on the wall behind the Round Table, where $M[i, j] = M[j, i]$ is the reward to be paid if snails i and j meet.

Describe and analyze an algorithm to compute the maximum total reward that the organizers could be forced to pay, given the array M as input.

37. You have mined a large slab of marble from a quarry. For simplicity, suppose the marble slab is a rectangle measuring n inches in height and m inches in width. You want to cut the slab into smaller rectangles of various sizes—some for kitchen counter tops, some for large sculpture projects, others for memorial headstones. You have a marble saw that can make either horizontal or vertical cuts across any rectangular slab. At any time, you can query the spot price $P[x, y]$ of an x-inch by y-inch marble rectangle, for any positive integers x and y. These prices depend on customer demand, and people who buy marble counter tops are weird, so don't make any assumptions about them; in particular, larger rectangles may have significantly smaller spot prices. Given the array of spot prices and the integers m and n as input, describe an algorithm to compute how to subdivide an $n \times m$ marble slab to maximize your profit.

38. This problem asks you to design efficient algorithms to construct optimal binary search trees that satisfy additional balance constraints. Your input consists of a sorted array $A[1..n]$ of search keys and an array $f[1..n]$ of frequency counts, where $f[i]$ is the number of searches for $A[i]$. This is exactly the same cost function as described in Section **??**. But now your task is to compute an optimal tree that satisfies some additional constraints.

    (a) ***AVL trees*** were the earliest self-balancing balanced binary search trees, first described in 1962 by Georgy Adelson-Velsky and Evgenii Landis. An AVL tree is a binary search tree where for every node v, the height of the left subtree of v and the height of the right subtree of v differ by at most one.

    Describe and analyze an efficient algorithm to construct an optimal AVL tree for a given set of search keys and frequencies.

    (b) ***Symmetric binary B-trees*** are another self-balancing binary trees, first described by Rudolf Bayer in 1972; these are better known by the name ***red-black trees***, after a somewhat simpler reformulation by Leo Guibas and Bob Sedgwick in 1978. A red-black tree is a binary search tree with the following additional constraints:

- Every node is either red or black.
- Every red node has a black parent.
- Every root-to-leaf path contains the same number of black nodes.

Describe and analyze an efficient algorithm to construct an optimal red-black tree for a given set of search keys and frequencies.

(c) ***AA trees*** were proposed by proposed by Arne Andersson in 1993 and slightly simplified (and named) by Mark Allen Weiss in 2000. AA trees are also known as *left-leaning red-black trees*, after a symmetric reformulation (with different rebalancing algorithms) by Bob Sedgewick in 2006. An AA tree is a red-black tree with one additional constraint:

- No left child is red.[24]

Describe and analyze an efficient algorithm to construct an optimal AA tree for a given set of search keys and frequencies.

39. Suppose you are given an $m \times n$ bitmap as an array $M[1..m, 1..n]$ of 0s and 1s. A *solid block* in M is a subarray of the form $M[i..i', j..i']$ in which all bits are equal. Suppose you want to decompose M into as few disjoint blocks as possible.

One natural recursive partitioning strategy is called a *guillotine subdivision*. If the entire bitmap M is a solid block, there is nothing to do. Otherwise, we cut M into two smaller bitmaps along a horizontal or vertical line, and then recursively decompose the two smaller bitmaps into solid blocks.

Any guillotine subdivision can be represented as a binary tree, where each internal node stores the position and orientation of a cut, and each leaf stores a single but 0 or 1 indicting the contents of the corresponding block. The *size* of a guillotine subdivision is the number of leaves in the corresponding binary tree (that is, the final number of solid blocks), and the *depth* of a guillotine subdivision is the depth of the corresponding binary tree.

(a) Describe and analyze an algorithm to compute a guillotine subdivision of M of minimum possible size.

(b) Show that a guillotine subdivision does *not* always yield a partition into the smallest number of solid blocks.

(c) Describe and analyze an algorithm to compute a guillotine subdivision for M with the smallest possible depth.

[24] Sedgwick's reformulation requires that no *right* child is red. Whatever. Andersson and Sedgewick are strangely silent about whether to measure angles clockwise or counterclockwise, whether Pluto is a planet, whether "lower rank" means "better" or "worse", and whether it's better to fight a hundred duck-sized horses or a single horse-sized duck.

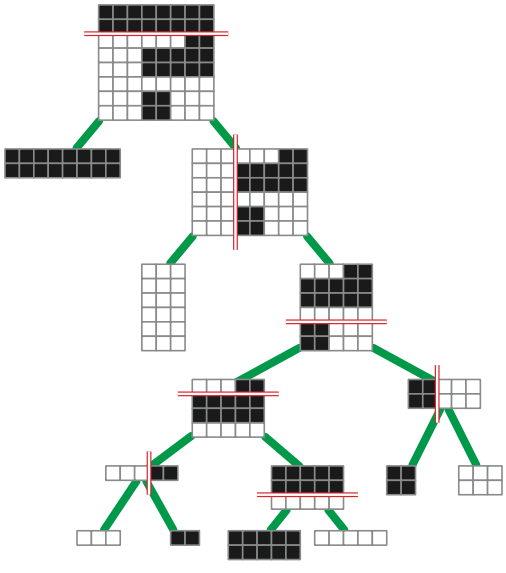

Figure 3.7. A guillotine subdivision with size 8 and depth 5.

(d) Describe and analyze an algorithm to determine $M[i, j]$, given the tree representing a guillotine decomposition for M and two indices i and j.

(e) Define the *depth* of a pixel $M[i, j]$ in a guillotine subdivision to be the depth of the leaf that contains that pixel. Describe and analyze an algorithm to compute a guillotine subdivision for M such that the sum of the depths of the pixels as small as possible.

(f) Describe and analyze an algorithm to compute a guillotine subdivision for M such that the sum of the depths of the *black* pixels as small as possible.

♠40. Congratulations! You've been hired by the giant online bookstore DeNile ("Not just a river in Egypt!") to optimize their warehouse robots. Each book that DeNile sells has a unique ISBN (International Standard Book Number), which is just a numerical value. Each of DeNile's warehouses contains a long row of bins, each containing multiple copies of a single book. These bins are arranged in sorted order by ISBN; each bin's ISBN is printed on the front of the bin in machine-readable form. Books are retrieved from these bins by robots, which run along rails parallel to the row of bins.

DeNile does not maintain a list of which bins contain which ISBN numbers; that would be too simple! Instead, to retrieve a desired book, the robot must first find that book's bin using a binary search. Because the search requires physical motion by the robot, we can no longer assume that each step of the binary search requires $O(1)$ time. Specifically:

- The robot always starts at the "0th bin" (where the books are loaded into boxes to ship to customers).
- Moving the robot from the ith bin to the jth bin requires $\alpha|i-j|$ seconds for some constant α.
- The robot must be directly in front of a bin in order to read that bin's ISBN. Reading an ISBN requires β seconds, for some constant β.
- Reversing the robot's direction of motion (from increasing to decreasing or vice versa) requires γ additional seconds, for some constant γ.
- When the robot finds the target bin, it extracts one book from that bin and returns to "the 0th bin".

Design and analyze an algorithm to compute a binary search tree over the bins that minimizes the total time the robot spends searching for books. Your input is an array $f[1..n]$ of integers, where $f[i]$ is the number of times that the robot will be asked to retrieve a book from the ith bin, along with the time parameters α, β, and γ.

♠41. A standard method to improve the cache performance of search trees is to pack more search keys and subtrees into each node. A ***B-tree*** is a rooted tree in which each internal node stores up to B keys and pointers to up to $B+1$ children, each the root of a smaller B-tree. Specifically, each node v stores three fields:

- a positive integer $v.d \le B$,
- a *sorted* array $v.key[1..v.d]$, and
- an array $v.child[0..v.d]$ of child pointers.

In particular, the number of child pointers is always exactly one more than the number of keys.[25]

Each pointer $v.child[i]$ is either NULL or a pointer to the root of a B-tree whose keys are all larger than $v.key[i]$ and smaller than $v.key[i+1]$. In particular, all keys in the leftmost subtree $v.child[0]$ are smaller than $v.key[1]$, and all keys in the rightmost subtree $v.child[v.d]$ are larger than $v.key[v.d]$.

[25]Normally, B-trees are required to satisfy two additional constraints, which guarantee a worst-case search cost of $O(\log_B n)$: Every leaf must have exactly the same depth, and every node except possibly the root must contain at least $B/2$ keys. However, in this problem, we are not interested in optimizing the *worst-case* search cost, but rather the *total* cost of a sequence of searches, so we will not impose these additional constraints.

Intuitively, you should have the following picture in mind:

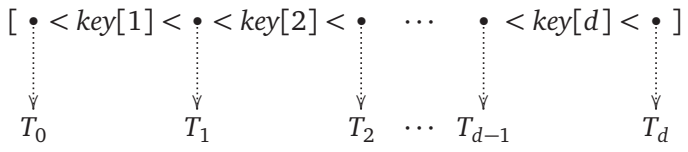

Here T_i is the subtree pointed to by $child[i]$.

The ***cost*** of searching for a key x in a B-tree is the number of nodes in the path from the root to the node containing x as one of its keys. A 1-tree is just a standard binary search tree.

Fix an arbitrary positive integer $B > 0$. (I suggest $B = 8$.) Suppose your are given a sorted array $A[1, \ldots, n]$ of search keys and a corresponding array $F[1, \ldots, n]$ of frequency counts, where $F[i]$ is the number of times that we will search for $A[i]$. Your task is to describe and analyze an efficient algorithm to find a B-tree that minimizes the total cost of searching for the given keys with the given frequencies.

(a) Describe a polynomial-time algorithm for the special case $B = 2$.

(b) Describe an algorithm for arbitrary B that runs in $O(n^{B+c})$ time for some fixed integer c.

♥(c) Describe an algorithm for arbitrary B that runs in $O(n^c)$ time for some fixed integer c that does *not* depend on B.

42. A string w of parentheses `(` and `)` and brackets `[` and `]` is ***balanced*** if it satisfies one of the following conditions:

- w is the empty string.
- $w = (x)$ for some balanced string x
- $w = [x]$ for some balanced string x
- $w = xy$ for some balanced strings x and y

For example, the string

$$w = \texttt{([()][]())[()()]()}$$

is balanced, because $w = xy$, where

$$x = \texttt{([()] [] ())} \qquad \text{and} \qquad y = \texttt{[() ()] ()}.$$

(a) Describe and analyze an algorithm to determine whether a given string of parentheses and brackets is balanced.

(b) Describe and analyze an algorithm to compute the length of a longest balanced subsequence of a given string of parentheses and brackets.

(c) Describe and analyze an algorithm to compute the length of a shortest balanced supersequence of a given string of parentheses and brackets.

♥(d) Describe and analyze an algorithm to compute the minimum edit distance from a given string of parentheses and brackets to a balanced string of parentheses and brackets.

♥(e) Describe and analyze an algorithm to compute the longest common balanced subsequence of two given strings of parentheses and brackets.

♥(f) Describe and analyze an algorithm to compute the longest palindromic balanced subsequence of a given string of parentheses and brackets.

♥(g) Describe and analyze an algorithm to compute the longest common palindromic balanced subsequence (whew!) of two given strings of parentheses and brackets.

For each problem, your input is an array $w[1..n]$, where $w[i] \in \{(,), [,]\}$ for every index i. (You may prefer to use different symbols instead of parentheses and brackets—for example, L, R, l, r or $\triangleleft, \triangleright, \blacktriangleleft, \blacktriangleright$—but please tell your grader what symbols you're using!)

♥43. Congratulations! Your research team has just been awarded a $50M multi-year project, jointly funded by DARPA, Google, and McDonald's, to produce DWIM: The first compiler to read programmers' minds! Your proposal and your numerous press releases all promise that DWIM will automatically correct errors in any given piece of code, while modifying that code as little as possible. Unfortunately, now it's time to start actually making the damn thing work.

As a warmup exercise, you decide to tackle the following necessary subproblem. Recall that the *edit distance* between two strings is the minimum number of single-character insertions, deletions, and replacements required to transform one string into the other. An *arithmetic expression* is a string w such that

- w is a string of one or more decimal digits,
- $w = (x)$ for some arithmetic expression x, or
- $w = x \diamond y$ for some arithmetic expressions x and y and some binary operator $\diamond$.

Suppose you are given a string of tokens from the alphabet $\{\#, \diamond, (,)\}$, where # represents a decimal digit and $\diamond$ represents a binary operator. Describe and analyze an algorithm to compute the minimum edit distance from the given string to an arithmetic expression.

44. Ribonucleic acid (RNA) molecules are long chains of millions of nucleotides or *bases* of four different types: adenine (A), cytosine (C), guanine (G), and

uracil (U). The *sequence* of an RNA molecule is a string $base[1..n]$, where each character $base[i] \in \{\texttt{A}, \texttt{C}, \texttt{G}, \texttt{U}\}$ corresponds to a base. In addition to the chemical bonds between adjacent bases in the sequence, hydrogen bonds can form between certain pairs of bases. The set of bonded base pairs is called the *secondary structure* of the RNA molecule.

We say that two base pairs (i, j) and (i', j') with $i < j$ and $i' < j'$ ***overlap*** if $i < i' < j < j'$ or $i' < i < j' < j$. In practice, most base pairs are non-overlapping. Overlapping base pairs create so-called *pseudoknots* in the secondary structure, which are essential for some RNA functions, but are more difficult to predict.

Suppose we want to predict the best possible secondary structure for a given RNA sequence. We will adopt a drastically simplified model of secondary structure:

- Each base can bond with at most one other base.
- Only A–U pairs and C–G pairs can bond.
- Pairs of the form $(i, i+1)$ and $(i, i+2)$ cannot bond.
- Bonded base pairs cannot overlap.

The last (and least realistic) restriction allows us to visualize RNA secondary structure as a sort of fat tree, as shown below.

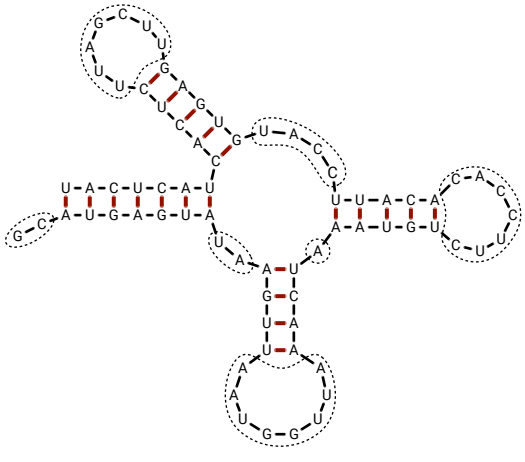

Figure 3.8. Example RNA secondary structure with 21 bonded base pairs, indicated by heavy red lines. Gaps are indicated by dotted curves. This structure has score $2^2 + 2^2 + 8^2 + 1^2 + 7^2 + 4^2 + 7^2 = 187$.

(a) Describe and analyze an algorithm that computes the maximum possible *number* of bonded base pairs in a secondary structure for a given RNA sequence.

(b) A *gap* in a secondary structure is a maximal substring of unpaired bases. Large gaps lead to chemical instabilities, so secondary structures with

smaller gaps are more likely. To account for this preference, let's define the *score* of a secondary structure to be the sum of the *squares* of the gap lengths; see Figure 3.8. (This score function is utterly fictional; real RNA structure prediction requires *much* more complicated scoring functions.)

Describe and analyze an algorithm that computes the minimum possible score of a secondary structure for a given RNA sequence.

♣45. (a) Describe and analyze an efficient algorithm to determine, given a string w and a regular expression R, whether $w \in L(R)$.

(b) *Generalized* regular expressions allow the binary operator $\cap$ (intersection) and the unary operator $\neg$ (complement), in addition to the usual $\bullet$ (concatenation), $+$ (or), and * (Kleene closure) operators. NFA constructions and Kleene's theorem imply that any generalized regular expression E represents a regular language $L(E)$.

Describe and analyze an efficient algorithm to determine, given a string w and a generalized regular expression E, whether $w \in L(E)$.

In both problems, assume that you are actually given a parse tree for the (generalized) regular expression, not just a string.

Trees and Subtrees

46. You've just been appointed as the new organizer of the first annual mandatory holiday party at Giggle (a subsidiary of Abugida). Giggle employees are organized into a strict hierarchy—a tree with the company president at the root. The all-knowing oracles in Human Resources have assigned a real number to each employee measuring how "fun" the employee is. To keep things social, there is one restriction on the guest list: an employee cannot attend the party if their immediate supervisor is also present. On the other hand, the president of the company *must* attend the party, even though she has a negative fun rating; it's her company, after all. Give an algorithm that makes a guest list for the party that maximizes the sum of the "fun" ratings of the guests.

47. Since so few people came to last year's holiday party, the president of Giggle decides to give each employee a present instead this year. Specifically, each employee must receive on the three gifts: (1) an all-expenses-paid six-week vacation anywhere in the world, (2) an all-the-pancakes-you-can-sort breakfast for two at Jumping Jack Flash's Flapjack Stack Shack, or (3) a burning paper bag full of dog poop. Corporate regulations prohibit any employee from receiving exactly the same gift as his/her direct supervisor. Any employee who receives a better gift than his/her direct supervisor will almost certainly be fired in a fit of jealousy.

As Giggle's official party czar, it's *your* job to decide which gift each employee receives. Describe an algorithm to distribute gifts so that the minimum number of people are fired. Yes, you may send the president a flaming bag of dog poop.

More formally, you are given a rooted tree T, representing the company hierarchy, and you want to label the nodes of T with integers 1, 2, or 3, so that every node has a different label from its parent. The *cost* of an labeling is the number of nodes with smaller labels than their parents. See Figure 3.9 for an example. Describe and analyze an algorithm to compute the minimum-cost labeling of T.

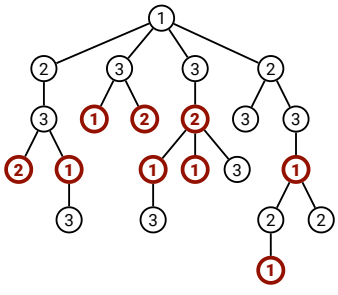

Figure 3.9. A tree labeling with cost 9. The nine bold nodes have smaller labels than their parents. This is *not* the optimal labeling for this tree.

48. After the Flaming Dog Poop Holiday Debacle, you were strongly encouraged to seek other employment, and so you left Giggle for rival company Twitbook. Unfortunately, the new president of Twitbook just decided to imitate Giggle by throwing her own holiday party, and in light of your past experience, appointed you as the official party organizer. The president demands that you invite exactly k employees, including the president herself, and everyone who is invited is required to attend. Yeah, that'll be fun.

    Just like at Giggle, employees at Twitbook are organized into a strict hierarchy: a tree with the company president at the root. The all-knowing oracles in Human Resources have assigned a real number to each employee indicating the *awkwardness* of inviting both that employee and their immediate supervisor; a negative value indicates that the employee and their supervisor actually like each other. Your goal is to choose a subset of exactly k employees to invite, so that the total awkwardness of the resulting party is as small as possible. For example, if the guest list does not include both an employee and their immediate supervisor, the total awkwardness is zero. The input to your algorithm is the tree T, the integer k, and the awkwardness of each node in T.

(a) Describe an algorithm that computes the total awkwardness of the least awkward subset of k employees, assuming the company hierarchy is described by a *binary* tree. That is, assume that each employee directly supervises at most two others.

♥(b) Describe an algorithm that computes the total awkwardness of the least awkward subset of k employees, with no restrictions on the company hierarchy.

49. Suppose we need to broadcast a message to all the nodes in a rooted tree. Initially, only the root node knows the message. In a single round, any node that knows the message can forward it to at most one of its children. See Figure 3.10 for an example.

(a) Design an algorithm to compute the minimum number of rounds required to broadcast the message to all nodes in a *binary* tree.

♦(b) Design an algorithm to compute the minimum number of rounds required to broadcast the message to all nodes in an *arbitrary* rooted tree. *[Hint: You may find techniques in the next chapter useful to prove your algorithm is correct, even though it's not a greedy algorithm.]*

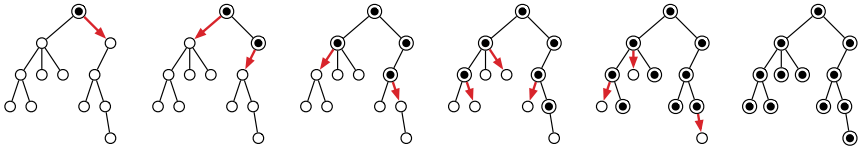

Figure 3.10. A message being distributed through a tree in five rounds.

50. One day, Alex got tired of climbing in a gym and decided to take a very large group of climber friends outside to climb. The climbing area where they went, had a huge wide boulder, not very tall, with various marked hand and foot holds. Alex quickly determined an "allowed" set of moves that her group of friends can perform to get from one hold to another.

The overall system of holds can be described by a rooted tree T with n vertices, where each vertex corresponds to a hold and each edge corresponds to an allowed move between holds. The climbing paths converge as they go up the boulder, leading to a unique hold at the summit, represented by the root of T.[26]

Alex and her friends (who are all excellent climbers) decided to play a game, where as many climbers as possible are simultaneously on the boulder and each climber needs to perform a sequence of *exactly* k moves. Each

[26] Q: Why do computer science professors think trees have their roots at the top?
A: Because they've never been outside!

climber can choose an arbitrary hold to start from, and all moves must move away from the ground. Thus, each climber traces out a path of k edges in the tree T, all directed toward the root. However, no two climbers are allowed to touch the same hold; the paths followed by different climbers cannot intersect at all.

Describe and analyze an efficient algorithm to compute the maximum number of climbers that can play this game. More formally, you are given a rooted tree T and an integer k, and you want to find the largest possible number of disjoint paths in T, where each path has length k. Do ***not*** assume that T is a binary tree. For example, given the tree T below and $k = 3$ as input, your algorithm should return the integer 8.

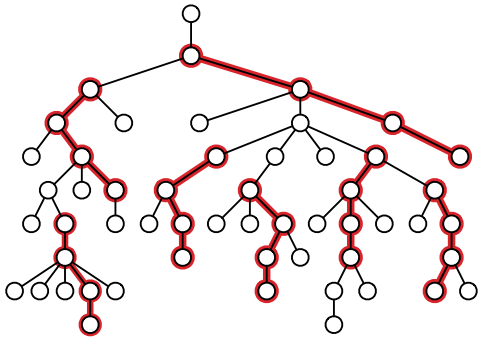

Figure 3.11. Seven disjoint paths of length $k = 3$. This is *not* the largest such set of paths in this tree.

51. Let T be a rooted binary tree with n vertices, and let $k \le n$ be a positive integer. We would like to mark k vertices in T so that every vertex has a nearby marked ancestor. More formally, we define the *clustering cost* of any subset K of vertices as

$$cost(K) = \max_v cost(v, K),$$

where the maximum is taken over all vertices v in the tree, and $cost(v, K)$ is the distance from v to its nearest ancestor in K:

$$cost(v, K) = \begin{cases} 0 & \text{if } v \in K \\ \infty & \text{if } v \text{ is the root of } T \text{ and } v \notin K \\ 1 + cost(parent(v)) & \text{otherwise} \end{cases}$$

In particular, $cost(K) = \infty$ if K excludes the root of T.

♥(a) Describe a dynamic programming algorithm to compute, given the tree T and an integer k, the minimum clustering cost of any subset of k vertices in T. For full credit, your algorithm should run in $O(n^2k^2)$ time.

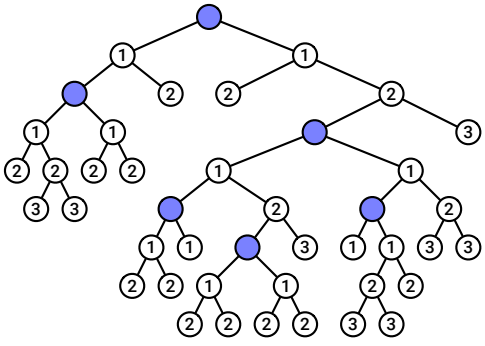

Figure 3.12. A subset of five vertices in a binary tree, with clustering cost 3.

(b) Describe a dynamic programming algorithm to compute, given the tree T and an integer r, the size of the smallest subset of vertices whose clustering cost is at most r. For full credit, your algorithm should run in $O(nr)$ time.

(c) Show that your solution for part (b) implies an algorithm for part (a) that runs in $O(n^2 \log n)$ time.

52. This question asks you to find efficient algorithms to compute the ***largest common rooted subtree*** of two given rooted trees. Recall that a *rooted tree* is a connected acyclic graph with a designated node called the root. A rooted subtree of a rooted tree consists of an arbitrary node and all its descendants. The precise definition of "common" depends on which pairs of rooted trees we consider isomorphic.

(a) Recall that a *binary* tree is a rooted tree in which every node has a (possibly empty) *left* subtree and a (possibly empty) *right* subtree. Two binary trees are isomorphic if and only if they are both empty, or their left subtrees are isomorphic and their right subtrees are isomorphic. Describe an algorithm to find the largest common *binary* subtree of two given *binary* trees.

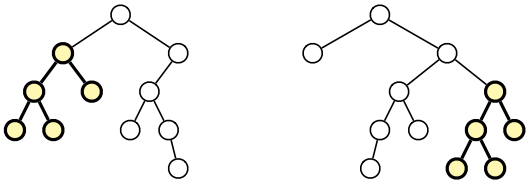

Figure 3.13. Two binary trees, with their largest common (rooted) subtree emphasized.

(b) In an *ordered* rooted tree, each node has a *sequence* of children, which are the roots of ordered rooted subtrees. Two ordered rooted trees are isomorphic if they are both empty, or if their ith subtrees are isomorphic

for every index i. Describe an algorithm to find the largest common ordered subtree of two ordered trees T_1 and T_2.

♦♥(c) In an *unordered* rooted tree, each node has an unordered *set* of children, which are the roots of unordered rooted subtrees. Two unordered rooted trees are isomorphic if they are both empty, or the subtrees of each root *can be ordered so that* their ith subtrees are isomorphic for every index i. Describe an algorithm to find the largest common unordered subtree of two unordered trees T_1 and T_2.

53. This question asks you to find efficient algorithms to compute optimal subtrees in *unrooted* trees—connected acyclic undirected graphs. A *subtree* of an unrooted tree is any connected subgraph.

(a) Suppose you are given an unrooted tree T with weights on its *edges*, which may be positive, negative, or zero. Describe an algorithm to find a *path* in T with maximum total weight.

(b) Suppose you are given an unrooted tree T with weights on its *vertices*, which may be positive, negative, or zero. Describe an algorithm to find a *subtree* of T with maximum total weight. *[This was a 2016 Google interview question.]*

(c) Let T_1 and T_2 be arbitrary *ordered* unrooted trees, meaning that the neighbors of every node have a well-defined cyclic order. Describe an algorithm to find the largest common *ordered* subtree of T_1 and T_2.

♦♥(d) Let T_1 and T_2 be arbitrary *unordered* unrooted trees. Describe an algorithm to find the largest common *unordered* subtree of T_1 and T_2.

54. ***Rooted minors*** of rooted trees are a natural generalization of subsequences. A rooted minor of a rooted tree T is any tree obtained by *contracting* one or more edges. When we contract an edge $u \to v$, where u is the parent of v, the children of v become new children of u and then v is deleted. In particular, the root of T is also the root of every rooted minor of T.

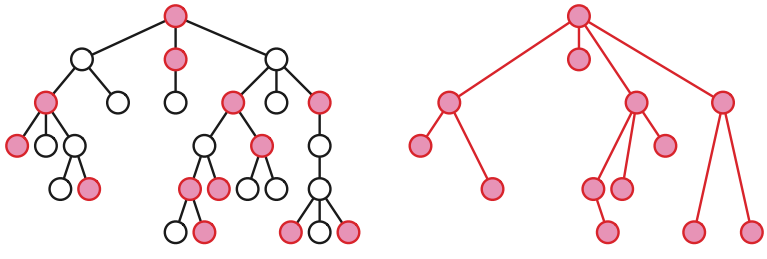

Figure 3.14. A rooted tree and one of its rooted minors.

(a) Let T be a rooted tree with labeled nodes. We say that T is *boring* if, for each node x, all children of x have the same label; children of different nodes may have different labels. Describe an algorithm to find the largest boring rooted minor of a given labeled rooted tree.

(b) Suppose we are given a rooted tree T whose nodes are labeled with numbers. Describe an algorithm to find the largest *heap-ordered rooted minor* of T. That is, your algorithm should return the largest rooted minor M such that every node in M has a smaller label than its children in M.

(c) Suppose we are given a *binary* tree T whose nodes are labeled with numbers. Describe an algorithm to find the largest *binary-search-ordered rooted minor* of T. That is, your algorithm should return a rooted minor M such that every node in M has at most two children, and an inorder traversal of M is an increasing subsequence of an inorder traversal of T.

(d) Recall that a rooted tree is *ordered* if the children of each node have a well-defined left-to-right order. Describe an algorithm to find the largest binary-search-ordered minor of an *arbitrary* ordered tree T whose nodes are labeled with numbers. Again, the left-to-right order of nodes in M should be consistent with their order in T.

♥(e) Describe an algorithm to find the largest common *ordered* rooted minor of two *ordered* labeled rooted trees.

(f) *[broken problem removed]*

Last revised 5:25am 10/12/25