

*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.)<sup>1</sup>

<sup>1</sup>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,<sup>2</sup> it took about a millennium for that rule to be stated explicitly. In the 7th century CE, another Indian scholar named Virahāṅka 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, Virahāṅka’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 Virahāṅka, in Leonardo of Pisa’s 1202 treatise *Liber Abaci*, one of the most influential early European works on “algorithm”. In full compliance with Stigler’s Law of Eponymy,<sup>3</sup> the modern *Fibonacci numbers* are defined using Virahāṅka’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*.

<sup>2</sup>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).

<sup>3</sup>“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  $T(n) = 2F_{n+1} - 1$ , 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<sup>4</sup> 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 **RECFIBO**(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  $0 \leq k < n$ . 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.

<sup>4</sup>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.<sup>5</sup>


```

MEMFIBO( $n$ ):
  if  $n = 0$ 
    return 0
  else if  $n = 1$ 
    return 1
  else
    if  $F[n]$  is undefined
       $F[n] \leftarrow \text{MEMFIBO}(n-1) + \text{MEMFIBO}(n-2)$ 
    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!

<sup>5</sup>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( $n$ ):
   $F[0] \leftarrow 0$ 
   $F[1] \leftarrow 1$ 
  for  $i \leftarrow 2$  to  $n$ 
     $F[i] \leftarrow F[i-1] + F[i-2]$ 
  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 Virahāṅka and later Sanskrit prosodists in the 12th century, and again by Fibonacci at the turn of the 13th century!<sup>6</sup>

<sup>6</sup>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

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.<sup>7</sup> 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-WWII America.<sup>8</sup> 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:

---

Bletchley 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.

<sup>7</sup>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.

<sup>8</sup>. . . 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  $\begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$  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  $n$ th Fibonacci number, we only need to compute the  $n$ th power of the matrix  $\begin{bmatrix} 0 & 1 \\ 1 & 1 \end{bmatrix}$ . If we use repeated squaring,<sup>9</sup> computing the  $n$ th power of something requires only  $O(\log n)$  multiplications. Here, because “something” is a  $2 \times 2$  matrix, that means  $O(\log n)$   $2 \times 2$  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:

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

<sup>9</sup>as suggested by Piṅgala for powers of 2 elsewhere in *Chandaḥśāstra*