

6 Dynamic Programming

Dynamic programming can be thought of as an optimization technique for particular classes of backtracking algorithms where subproblems are repeatedly solved. Note that the term *dynamic* in dynamic programming should not be confused with dynamic programming languages, like Scheme or Lisp. Nor should the term *programming* be confused with the act of writing computer programs. In the context of algorithms, dynamic programming always refers to the technique of filling in a table with values computed from other table values. (It's dynamic because the values in the table are filled in by the algorithm based on other values of the table, and it's programming in the sense of setting things in a table, like how television programming is concerned with when to broadcast what shows.)

6.1 Fibonacci Numbers

Before presenting the dynamic programming technique, it will be useful to first show a related technique, called **memoization**, on a toy example: The Fibonacci numbers. What we want is a routine to compute the n th Fibonacci number:

```
// fib -- compute Fibonacci(n)
function fib(integer n): integer
```

By definition, the n th Fibonacci number, denoted F_n is

$$F_0 = 0$$

$$F_1 = 1$$

$$F_n = F_{n-1} + F_{n-2}$$

How would one create a good algorithm for finding the n th Fibonacci-number? Let's begin with the naive algorithm, which codes the mathematical definition:

```
// fib -- compute Fibonacci(n)
function fib(integer n): integer
  assert (n >= 0)
  if n == 0: return 0 fi
  if n == 1: return 1 fi
```

```

return fib(n - 1) + fib(n - 2)
end

```

This code sample is also available in Ada^a

^a http://en.wikibooks.org/wiki/Ada_Programming%2FAlgorithms%23Simple_Implementation

Note that this is a toy example because there is already a mathematically closed form for F_n :

$$F(n) = \frac{\phi^n - (1 - \phi)^n}{\sqrt{5}}$$

where:

$$\phi = \frac{1 + \sqrt{5}}{2}$$

This latter equation is known as the Golden Ratio¹. Thus, a program could efficiently calculate F_n for even very large n . However, it's instructive to understand what's so inefficient about the current algorithm.

To analyze the running time of `fib` we should look at a call tree for something even as small as the sixth Fibonacci number:

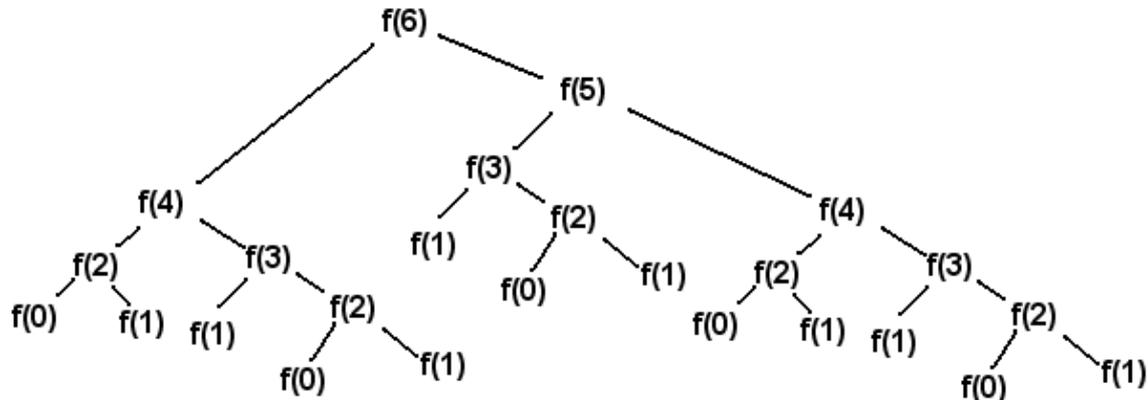


Figure 5

Every leaf of the call tree has the value 0 or 1, and the sum of these values is the final result. So, for any n , the number of leaves in the call tree is actually F_n itself! The closed form thus tells us that the number of leaves in `fib(n)` is approximately equal to

$$\left(\frac{1 + \sqrt{5}}{2}\right)^n \approx 1.618^n = 2^{\lg(1.618^n)} = 2^{n \lg(1.618)} \approx 2^{0.69n}.$$

¹ http://en.wikipedia.org/wiki/golden_ratio

(Note the algebraic manipulation used above to make the base of the exponent the number 2.) This means that there are far too many leaves, particularly considering the repeated patterns found in the call tree above.

One optimization we can make is to save a result in a table once it's already been computed, so that the same result needs to be computed only once. The optimization process is called memoization and conforms to the following methodology:

Memoization Methodology

```
# Start with a backtracking algorithm# Look up the problem in a table; if there's a
valid entry for it, return that value# Otherwise, compute the problem recursively, and
then store the result in the table before returning the value
```

Consider the solution presented in the backtracking chapter for the Longest Common Subsequence problem. In the execution of that algorithm, many common subproblems were computed repeatedly. As an optimization, we can compute these subproblems once and then store the result to read back later. A recursive memoization algorithm can be turned "bottom-up" into an iterative algorithm that fills in a table of solutions to subproblems. Some of the subproblems solved might not be needed by the end result (and that is where dynamic programming differs from memoization), but dynamic programming can be very efficient because the iterative version can better use the cache and have less call overhead. Asymptotically, dynamic programming and memoization have the same complexity.

So how would a fibonacci program using memoization work? Consider the following program ($f[n]$ contains the n th Fibonacci-number if has been calculated, -1 otherwise):

```
function fib(integer n): integer
  if n == 0 or n == 1:
    return n
  else-if f[n] != -1:
    return f[n]
  else
    f[n] = fib(n - 1) + fib(n - 2)
    return f[n]
  fi
end
```

This code sample is also available in Ada^a

^a http://en.wikibooks.org/wiki/Ada_Programming%2FAlgorithms%23Cached_Implementation

The code should be pretty obvious. If the value of $\text{fib}(n)$ already has been calculated it's stored in $f[n]$ and then returned instead of calculating it again. That means all the copies of the sub-call trees are removed from the calculation.

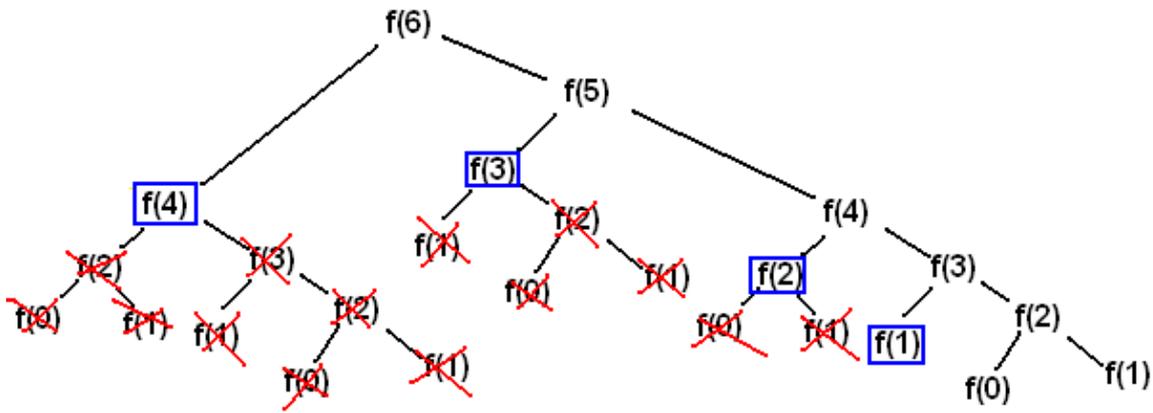


Figure 6

The values in the blue boxes are values that already have been calculated and the calls can thus be skipped. It is thus a lot faster than the straight-forward recursive algorithm. Since every value less than n is calculated once, and only once, the first time you execute it, the asymptotic running time is $O(n)$. Any other calls to it will take $O(1)$ since the values have been precalculated (assuming each subsequent call's argument is less than n).

The algorithm does consume a lot of memory. When we calculate $\text{fib}(n)$, the values $\text{fib}(0)$ to $\text{fib}(n)$ are stored in main memory. Can this be improved? Yes it can, although the $O(1)$ running time of subsequent calls are obviously lost since the values aren't stored. Since the value of $\text{fib}(n)$ only depends on $\text{fib}(n-1)$ and $\text{fib}(n-2)$ we can discard the other values by going bottom-up. If we want to calculate $\text{fib}(n)$, we first calculate $\text{fib}(2) = \text{fib}(0) + \text{fib}(1)$. Then we can calculate $\text{fib}(3)$ by adding $\text{fib}(1)$ and $\text{fib}(2)$. After that, $\text{fib}(0)$ and $\text{fib}(1)$ can be discarded, since we don't need them to calculate any more values. From $\text{fib}(2)$ and $\text{fib}(3)$ we calculate $\text{fib}(4)$ and discard $\text{fib}(2)$, then we calculate $\text{fib}(5)$ and discard $\text{fib}(3)$, etc. etc. The code goes something like this:

```
function fib(integer n): integer
  if n == 0 or n == 1:
    return n
  fi

  let u := 0
  let v := 1

  for i := 2 to n:
    let t := u + v
    u := v
    v := t
  repeat

  return v
end
```

This code sample is also available in Ada^a

^a http://en.wikibooks.org/wiki/Ada_Programming%2FAlgorithms%23Memory_Optimized_Implementation

We can modify the code to store the values in an array for subsequent calls, but the point is that we don't *have* to. This method is typical for dynamic programming. First we identify what subproblems need to be solved in order to solve the entire problem, and then we calculate the values bottom-up using an iterative process.

6.2 Longest Common Subsequence (DP version)

This will remind us of the backtracking version and then improve it via memoization. Finally, the recursive algorithm will be made iterative and be full-fledged DP. [TODO: write this section]

6.3 Matrix Chain Multiplication

Suppose that you need to multiply a series of n matrices M_1, \dots, M_n together to form a product matrix P :

$$P = M_1 \cdot M_2 \cdots M_{n-1} \cdot M_n$$

This will require $n - 1$ multiplications, but what is the fastest way we can form this product? Matrix multiplication is associative, that is,

$$(A \cdot B) \cdot C = A \cdot (B \cdot C)$$

for any A, B, C , and so we have some choice in what multiplication we perform first. (Note that matrix multiplication is *not* commutative, that is, it does not hold in general that $A \cdot B = B \cdot A$.)

Because you can only multiply two matrices at a time the product $M_1 \cdot M_2 \cdot M_3 \cdot M_4$ can be paranthesized in these ways:

$$((M_1 M_2) M_3) M_4$$

$$(M_1 (M_2 M_3)) M_4$$

$$M_1 ((M_2 M_3) M_4)$$

$$(M_1 M_2) (M_3 M_4)$$

$$M_1 (M_2 (M_3 M_4))$$

Two matrices M_1 and M_2 can be multiplied if the number of columns in M_1 equals the number of rows in M_2 . The number of rows in their product will equal the number rows in M_1 and the number of columns will equal the number of columns in M_2 . That is, if the dimensions of M_1 is $a \times b$ and M_2 has dimensions $b \times c$ their product will have dimensions $a \times c$.

To multiply two matrices with each other we use a function called matrix-multiply that takes two matrices and returns their product. We will leave implementation of this function alone for the moment as it is not the focus of this chapter (how to multiply two matrices in the fastest way has been under intensive study for several years [TODO: propose this topic for the *Advanced* book]). The time this function takes to multiply two matrices of size $a \times b$ and $b \times c$ is proportional to the number of scalar multiplications, which is proportional to abc . Thus, paranthezation matters: Say that we have three matrices M_1 , M_2 and M_3 . M_1 has dimensions 5×100 , M_2 has dimensions 100×100 and M_3 has dimensions 100×50 . Let's paranthezise them in the two possible ways and see which way requires the least amount of multiplications. The two ways are

$((M_1 M_2) M_3)$, and

$(M_1 (M_2 M_3))$.

To form the product in the first way requires 75000 scalar multiplications ($5 \times 100 \times 100 = 50000$ to form product $(M_1 M_2)$ and another $5 \times 100 \times 50 = 25000$ for the last multiplications.) This might seem like a lot, but in comparison to the 525000 scalar multiplications required by the second parenthesization ($50 \times 100 \times 100 = 500000$ plus $5 \times 50 \times 100 = 25000$) it is miniscule! You can see why determining the parenthesization is important: imagine what would happen if we needed to multiply 50 matrices!

6.3.1 Forming a Recursive Solution

Note that we concentrate on finding a how many scalar multiplications are needed instead of the actual order. This is because once we have found a working algorithm to find the amount it is trivial to create an algorithm for the actual parenthesization. It will, however, be discussed in the end.

So how would an algorithm for the optimum parenthesization look? By the chapter title you might expect that a dynamic programming method is in order (not to give the answer away or anything). So how would a dynamic programming method work? Because dynamic programming algorithms are based on optimal substructure, what would the optimal substructure in this problem be?

Suppose that the optimal way to parenthesize

$$M_1 M_2 \dots M_n$$

splits the product at k :

$$(M_1 M_2 \dots M_k) (M_{k+1} M_{k+2} \dots M_n)$$

Then the optimal solution contains the optimal solutions to the two subproblems

$$(M_1 \dots M_k)$$

$$(M_{k+1} \dots M_n)$$

That is, just in accordance with the fundamental principle of dynamic programming, the solution to the problem depends on the solution of smaller sub-problems.

Let's say that it takes $c(n)$ scalar multiplications to multiply matrices M_n and M_{n+1} , and $f(m, n)$ is the number of scalar multiplications to be performed in an optimal parenthesization of the matrices $M_m \dots M_n$. The definition of $f(m, n)$ is the first step toward a solution.

When $n - m = 1$, the formulation is trivial; it is just $c(m)$. But what is it when the distance is larger? Using the observation above, we can derive a formulation. Suppose an optimal solution to the problem divides the matrices at matrices k and $k+1$ (i.e. $(M_m \dots M_k)(M_{k+1} \dots M_n)$) then the number of scalar multiplications are.

$$f(m, k) + f(k + 1, n) + c(k)$$

That is, the amount of time to form the first product, the amount of time it takes to form the second product, and the amount of time it takes to multiply them together. But what is this optimal value k ? The answer is, of course, the value that makes the above formula assume its minimum value. We can thus form the complete definition for the function:

$$f(m, n) = \begin{cases} \min_{m \leq k < n} f(m, k) + f(k + 1, n) + c(k) & \text{if } n - m > 1 \\ 0 & \text{if } n = m \end{cases}$$

A straight-forward recursive solution to this would look something like this (*the language is Wikicode²*):

```
function f(m, n) {
  if m == n
    return 0

  let minCost := ∞

  for k := m to n - 1 {
    v := f(m, k) + f(k + 1, n) + c(k)
    if v < minCost
      minCost := v
  }
  return minCost
}
```

² <http://en.wikipedia.org/wiki/Wikipedia:Wikicode>

This rather simple solution is, unfortunately, not a very good one. It spends mountains of time recomputing data and its running time is exponential.

Using the same adaptation as above we get:

```
function f(m, n) {  
    if m == n  
        return 0  
  
    else-if f[m,n] != -1:  
        return f[m,n]  
    fi  
  
    let minCost := ∞  
  
    for k := m to n - 1 {  
        v := f(m, k) + f(k + 1, n) + c(k)  
        if v < minCost  
            minCost := v  
    }  
    f[m,n]=minCost  
    return minCost  
}
```

6.4 Parsing Any Context-Free Grammar

Note that special types of context-free grammars can be parsed much more efficiently than this technique, but in terms of generality, the DP method is the only way to go.