

CSC373

Week 3: Dynamic Programming

Nisarg Shah

Recap

- Greedy Algorithms
 - Interval scheduling
 - Interval partitioning
 - Minimizing lateness
 - Huffman encoding
 - ...

5.4 Warning: Greed is Stupid

If we're very very very very lucky, we can bypass all the recurrences and tables and so forth, and solve the problem using a *greedy* algorithm. The general greedy strategy is look for the best first step, take it, and then continue. While this approach seems very natural, it almost never works; optimization problems that can be solved correctly by a greedy algorithm are *very* rare. Nevertheless, for many problems that should be solved by dynamic programming, many students' first intuition is to apply a greedy strategy.

For example, a greedy algorithm for the edit distance problem might look for the longest common substring of the two strings, match up those substrings (since those substitutions don't cost anything), and then recursively look for the edit distances between the left halves and right halves of the strings. If there is no common substring—that is, if the two strings have no characters in common—the edit distance is clearly the length of the larger string. If this sounds like a stupid hack to you, pat yourself on the back. It isn't even *close* to the correct solution.

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.⁶

Jeff Erickson on greedy algorithms...

The 1950s were not good years for mathematical research. 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'. The RAND Corporation was employed by the Air Force, and the Air Force had Wilson as its boss, essentially. Hence, 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?

— Richard Bellman, on the origin of his term 'dynamic programming' (1984)



Richard Bellman's quote from Jeff Erickson's book

Dynamic Programming

- **Outline**

- Breaking the problem down into simpler subproblems, solve each subproblem just once, and store their solutions.
- The next time the same subproblem occurs, instead of recomputing its solution, simply look up its previously computed solution.
- Hopefully, we save a lot of computation at the expense of modest increase in storage space.
- Also called “**memoization**”

- **How is this different from divide & conquer?**

Weighted Interval Scheduling


- **Problem**

- Job j starts at time s_j and finishes at time f_j
- Each job j has a weight w_j
- Two jobs are compatible if they don't overlap
- **Goal:** find a set S of mutually compatible jobs with highest total weight $\sum_{j \in S} w_j$

- Recall: If all $w_j = 1$, then this is simply the interval scheduling problem from last week
 - Greedy algorithm based on earliest finish time ordering was optimal for this case

Recall: Interval Scheduling

- What if we simply try to use it again?
 - Fails spectacularly!




Weighted Interval Scheduling

- What if we use other orderings?
 - By weight: choose jobs with highest w_j first
 - Maximum weight per time: choose jobs with highest $w_j / (f_j - s_j)$ first
 - ...
- None of them work!
 - They're arbitrarily worse than the optimal solution
 - In fact, under a certain formalization, “no greedy algorithm” can produce any “decent approximation” in the worst case (beyond this course!)

Weighted Interval Scheduling

- **Convention**

- Jobs are sorted by finish time: $f_1 \leq f_2 \leq \dots \leq f_n$
- $p[j] =$ largest index $i < j$ such that job i is compatible with job j (i.e. $f_i < s_j$)



Among jobs before job j , the ones compatible with it are precisely $1 \dots i$

E.g.

$$\begin{aligned} p[8] &= 1, \\ p[7] &= 3, \\ p[2] &= 0 \end{aligned}$$

Weighted Interval Scheduling

- **The DP approach**

- Let OPT be an optimal solution

- **Two cases** regarding job n :

- **Option 1: Job n is in OPT**

- Can't use incompatible jobs $\{p[n] + 1, \dots, n - 1\}$
- Must select optimal subset of jobs from $\{1, \dots, p[n]\}$

- **Option 2: Job n is not in OPT**

- Must select optimal subset of jobs from $\{1, \dots, n - 1\}$

- OPT is best of both

- **Note:** In both cases, knowing how to solve any prefix of our ordering is enough solve the overall problem

Weighted Interval Scheduling

- **The DP approach**

- $OPT(j)$ = maximum value from compatible jobs in $\{1, \dots, j\}$

- Base case: $OPT(0) = 0$

- **Two cases** regarding job j :

- **Job j is selected:** optimal value is $v_j + OPT(p[j])$

- **Job j is not selected:** optimal value is $OPT(j - 1)$

- $OPT(j)$ is best of both worlds

- **Bellman equation:**

$$OPT(j) = \begin{cases} 0 & \text{if } j = 0 \\ \max\{OPT(j - 1), v_j + OPT(p[j])\} & \text{if } j > 0 \end{cases}$$

Brute Force Solution

BRUTE-FORCE ($n, s_1, \dots, s_n, f_1, \dots, f_n, w_1, \dots, w_n$)

Sort jobs by finish time and renumber so that $f_1 \leq f_2 \leq \dots \leq f_n$.

Compute $p[1], p[2], \dots, p[n]$ via binary search.

RETURN COMPUTE-OPT(n).

COMPUTE-OPT(j)

IF ($j = 0$)

 RETURN 0.

ELSE

 RETURN $\max \{ \text{COMPUTE-OPT}(j-1), w_j + \text{COMPUTE-OPT}(p[j]) \}$.

Brute Force Solution

COMPUTE-OPT(j)

IF ($j = 0$)

 RETURN 0.

ELSE

 RETURN $\max \{ \text{COMPUTE-OPT}(j-1), w_j + \text{COMPUTE-OPT}(p[j]) \}$.

- Q: Worst-case running time of COMPUTE-OPT(n)?
 - a) $\Theta(n)$
 - b) $\Theta(n \log n)$
 - c) $\Theta(1.618^n)$
 - d) $\Theta(2^n)$

Brute Force Solution

- Brute force running time
 - It is possible that $p(j) = j - 1$ for each j
 - Then, we call $\text{COMPUTE-OPT}(j - 1)$ twice in $\text{COMPUTE-OPT}(j)$
 - So this might take 2^n steps
 - But we can just check if j is compatible with $j - 1$, and if so, only execute the part where we select j
 - Now the worst case is where $p(j) = j - 2$ for each j
 - **Running time:** $T(n) = T(n - 1) + T(n - 2)$
 - Fibonacci, golden ratio, ... 😊
 - $T(n) = O(\varphi^n)$, where $\varphi \approx 1.618$

Dynamic Programming

- **Why is the runtime high?**
 - Some solutions are being computed many, many times
 - E.g. if $p[5] = 3$, then `COMPUTE-OPT(5)` might call `COMPUTE-OPT(4)` and `COMPUTE-OPT(3)`
 - But `COMPUTE-OPT(4)` might in turn call `COMPUTE-OPT(3)`
- **Memoization** trick
 - Simply remember what you've already computed, and re-use the answer if needed in future

Dynamic Program: Top-Down

- Let's store $\text{COMPUTE-OPT}(j)$ in $M[j]$

$\text{TOP-DOWN}(n, s_1, \dots, s_n, f_1, \dots, f_n, w_1, \dots, w_n)$

Sort jobs by finish time and renumber so that $f_1 \leq f_2 \leq \dots \leq f_n$.

Compute $p[1], p[2], \dots, p[n]$ via binary search.

$M[0] \leftarrow 0$.  global array

RETURN $\text{M-COMPUTE-OPT}(n)$.

$\text{M-COMPUTE-OPT}(j)$

IF ($M[j]$ is uninitialized)

$M[j] \leftarrow \max \{ \text{M-COMPUTE-OPT}(j-1), w_j + \text{M-COMPUTE-OPT}(p[j]) \}$.

RETURN $M[j]$.

Dynamic Program: Top-Down

- **Claim:** This memoized version takes $O(n \log n)$ time
 - Sorting jobs takes $O(n \log n)$
 - It also takes $O(n \log n)$ to do n binary searches to compute $p(j)$ for each j
 - M-Compute-OPT(j) is called *at most once* for each j
 - Each such call takes $O(1)$ time, not considering the time taken by any subroutine calls
 - So M-Compute-OPT(n) takes only $O(n)$ time
 - Overall time is $O(n \log n)$

Dynamic Program: Bottom-Up

- Find an order in which to call the functions so that the sub-solutions are ready when needed

BOTTOM-UP($n, s_1, \dots, s_n, f_1, \dots, f_n, w_1, \dots, w_n$)

Sort jobs by finish time and renumber so that $f_1 \leq f_2 \leq \dots \leq f_n$.

Compute $p[1], p[2], \dots, p[n]$.

$M[0] \leftarrow 0$.

previously computed values

FOR $j = 1$ **TO** n

$M[j] \leftarrow \max \{ M[j-1], w_j + M[p[j]] \}$.

Top-Down vs Bottom-Up

- **Top-Down may be preferred...**
 - ...when not all sub-solutions need to be computed on some inputs
 - ...because one does not need to think of the “right order” in which to compute sub-solutions
- **Bottom-Up may be preferred...**
 - ...when all sub-solutions will anyway need to be computed
 - ...because it is sometimes faster as it prevents recursive call overheads and unnecessary random memory accesses

Optimal Solution

- This approach gave us the optimal value
- **What about the actual solution (subset of jobs)?**
 - Typically, this is done by **maintaining the optimal value and an optimal solution** for each subproblem
 - So, we compute two quantities:

$$OPT(j) = \begin{cases} 0 & \text{if } j = 0 \\ \max\{OPT(j-1), v_j + OPT(p[j])\} & \text{if } j > 0 \end{cases}$$

$$S(j) = \begin{cases} \emptyset & \text{if } j = 0 \\ S(j-1) & \text{if } j > 0 \wedge OPT(j-1) \geq v_j + OPT(p[j]) \\ \{j\} \cup S(p[j]) & \text{if } j > 0 \wedge OPT(j-1) < v_j + OPT(p[j]) \end{cases}$$

Optimal Solution

$$OPT(j) = \begin{cases} 0 & \text{if } j = 0 \\ \max\{OPT(j-1), v_j + OPT(p[j])\} & \text{if } j > 0 \end{cases}$$

$$S(j) = \begin{cases} \emptyset & \text{if } j = 0 \\ S(j-1) & \text{if } j > 0 \wedge OPT(j-1) \geq v_j + OPT(p[j]) \\ \{j\} \cup S(p[j]) & \text{if } j > 0 \wedge OPT(j-1) < v_j + OPT(p[j]) \end{cases}$$

This works with both top-down (memoization) and bottom-up approaches.

In this problem, we can do something simpler: just compute OPT first, and later compute S using only OPT .


Optimal Substructure Property

- Dynamic programming applies well to problems that have **optimal substructure property**
 - Optimal solution to a problem contains (or can be computed easily given) optimal solution to subproblems.
- **Recall: divide-and-conquer also uses this property**
 - You can think of divide-and-conquer as a special case of dynamic programming, where the two (or more) subproblems you need to solve don't "overlap"
 - So there's no need for memoization
 - In dynamic programming, one of the subproblems may in turn require solution to the other subproblem...

Knapsack Problem

- **Problem**

- n items: item i provides value $v_i > 0$ and has weight $w_i > 0$
- Knapsack has weight capacity W
- Assumption: W , each v_i , and each w_i is an integer
- **Goal:** pack the knapsack with a collection of items with highest total value given that their total weight is at most W



i	v_i	w_i
1	\$1	1 kg
2	\$6	2 kg
3	\$18	5 kg
4	\$22	6 kg
5	\$28	7 kg

knapsack instance
(weight limit $W = 11$)

A First Attempt

- Let $OPT(w)$ = maximum value we can pack with a knapsack of capacity w
 - **Goal:** Compute $OPT(W)$
 - **Claim?** $OPT(w)$ must use at least one job j with weight $\leq w$ and then optimally pack the remaining capacity of $w - w_j$
 - Let $w^* = \min_j w_j$
 - $OPT(w) = \begin{cases} 0 & \text{if } w < w^* \\ \max_{j:w_j \leq w} v_j + OPT(w - w_j) & \text{if } w \geq w^* \end{cases}$
- **Why is this wrong?**
 - It might use an item more than once!

A Refined Attempt

- $OPT(i, w)$ = maximum value we can pack using only items $1, \dots, i$ given capacity w
 - **Goal:** Compute $OPT(n, W)$
- Consider item i
 - If $w_i > w$, then we can't choose i . Just use $OPT(i - 1, w)$
 - If $w_i \leq w$, there are two cases:
 - If we choose i , the best is $v_i + OPT(i - 1, w - w_i)$
 - If we don't choose i , the best is $OPT(i - 1, w)$

$$OPT(i, w) = \begin{cases} 0 & \text{if } i = 0 \\ OPT(i - 1, w) & \text{if } w_i > w \\ \max \{ OPT(i - 1, w), v_i + OPT(i - 1, w - w_i) \} & \text{otherwise} \end{cases}$$

Running Time

- Consider possible evaluations $OPT(i, w)$
 - $i \in \{1, \dots, n\}$
 - $w \in \{1, \dots, W\}$ (recall weights and capacity are integers)
 - There are $O(n \cdot W)$ possible evaluations of OPT
 - Each is evaluated at most once (memoization)
 - Each takes $O(1)$ time to evaluate
 - So the total running time is $O(n \cdot W)$
- Q: Is this polynomial in the input size?
 - A: No! But it's pseudo-polynomial.

What if...?

- Note that this algorithm runs in polynomial time when W is polynomially bounded in the length of the input
- Q: What if instead of W, w_1, \dots, w_n being small integers, we were told that v_1, \dots, v_n are small integers?
 - Then we can use a different dynamic programming approach!

A Different DP

- $OPT(i, v)$ = minimum capacity needed to pack a total value of at least v using items $1, \dots, i$
 - **Goal:** Compute $\max\{v \in \{1, \dots, V\} : OPT(i, v) \leq W\}$
- Consider item i
 - If we choose i , we need capacity $w_i + OPT(i - 1, v - v_i)$
 - If we don't choose i , we need capacity $OPT(i - 1, v)$

$$OPT(i, v) = \begin{cases} 0 & \text{if } v \leq 0 \\ \infty & \text{if } v > 0, i = 0 \\ \min \left\{ w_i + OPT(i - 1, v - v_i), \right. \\ \quad \left. OPT(i - 1, v) \right\} & \text{if } v > 0, i > 0 \end{cases}$$

A Different DP

- $OPT(i, v)$ = minimum capacity needed to pack a total value of at least v using items $1, \dots, i$
 - **Goal:** Compute $\max\{v \in \{1, \dots, V\} : OPT(i, v) \leq W\}$
- This approach has running time $O(n \cdot V)$, where $V = v_1 + \dots + v_n$
- So we can get $O(n \cdot W)$ or $O(n \cdot V)$
- Can we remove the dependence on both V and W ?
 - Not likely. Knapsack problem is NP-complete (we'll see later).

Looking Ahead: FPTAS

- While we cannot hope to solve the problem exactly in time $O(\text{poly}(n, \log W, \log V))$...
 - For any $\epsilon > 0$, we can get a value that is within $1 + \epsilon$ multiplicative factor of the optimal value in time $O\left(\text{poly}\left(n, \log W, \log V, \frac{1}{\epsilon}\right)\right)$
 - Such algorithms are known as fully polynomial-time approximation scheme (FPTAS)
 - **Core idea behind FPTAS for knapsack:**
 - Approximate all weights and values up to the desired precision
 - Solve knapsack on approximate input using DP

Single-Source Shortest Paths

- **Problem**


- **Input:** A directed graph $G = (V, E)$ with edge lengths ℓ_{vw} on each edge (v, w) , and a source vertex s
- **Goal:** Compute the length of the shortest path from s to every vertex t

- When $\ell_{vw} \geq 0$ for each (v, w) ...

- Dijkstra's algorithm can be used for this purpose
- But it fails when some edge lengths can be negative
- What do we do in this case?

Single-Source Shortest Paths

- Cycle length = sum of lengths of edges in the cycle
- If there is a negative length cycle, shortest paths are not even well defined...
 - You can traverse the cycle arbitrarily many times to get arbitrarily “short” paths




Single-Source Shortest Paths

- But if there are no negative cycles...
 - Shortest paths are well-defined even when some of the edge lengths may be negative
- **Claim:** With no negative cycles, there is always a shortest path from any vertex to any other vertex that is **simple**
 - Consider the shortest $s \rightsquigarrow t$ path with the fewest edges among all shortest $s \rightsquigarrow t$ paths
 - If it has a cycle, removing the cycle creates a path with fewer edges that is no longer than the original path


Optimal Substructure Property

- Consider a simple shortest $s \rightsquigarrow t$ path P
 - It could be just a single edge
 - But if P has more than one edges, consider u which immediately precedes t in the path
 - If $s \rightsquigarrow t$ is shortest, $s \rightsquigarrow u$ must be shortest as well and it must use one fewer edge than the $s \rightsquigarrow t$ path



Optimal Substructure Property

- $OPT(t, i)$ = shortest path from s to t using at most i edges
- **Then:**
 - Either this path uses at most $i - 1$ edges $\Rightarrow OPT(t, i - 1)$
 - Or it uses i edges $\Rightarrow \min_u OPT(u, i - 1) + \ell_{ut}$



Optimal Substructure Property

- $OPT(t, i)$ = shortest path from s to t using at most i edges
- **Then:**
 - Either this path uses at most $i - 1$ edges $\Rightarrow OPT(t, i - 1)$
 - Or it uses i edges $\Rightarrow \min_u OPT(u, i - 1) + \ell_{ut}$

$$OPT(t, i) = \begin{cases} 0 & i = 0 \vee t = s \\ \infty & i = 0 \wedge t \neq s \\ \min \left\{ OPT(t, i - 1), \min_u OPT(u, i - 1) + \ell_{ut} \right\} & \text{otherwise} \end{cases}$$

- **Running time:** $O(n^2)$ calls, each takes $O(n)$ time $\Rightarrow O(n^3)$
- **Q:** What do you need to store to also get the actual paths?

Side Notes

- **Bellman-Ford-Moore algorithm**
 - Improvement over this DP
 - Running time remains $O(m \cdot n)$ for n vertices, m edges
 - But the space complexity reduces to $O(m + n)$

year	worst case	discovered by
1955	$O(n^4)$	Shimbel
1956	$O(m n^2 W)$	Ford
1958	$O(m n)$	Bellman, Moore
1983	$O(n^{3/4} m \log W)$	Gabow
1989	$O(m n^{1/2} \log(nW))$	Gabow-Tarjan
1993	$O(m n^{1/2} \log W)$	Goldberg
2005	$O(n^{2.38} W)$	Sankowski, Yuster-Zwick
2016	$\tilde{O}(n^{10/7} \log W)$	Cohen-Mądry-Sankowski-Vladu
20xx	???	

single-source shortest paths with weights between $-W$ and W

Maximum Length Paths?

- Can we use a similar DP to compute maximum length paths from s to all other vertices?
- This is well defined when there are no positive cycles, in which case, yes.
- What if there are positive cycles, but we want maximum length *simple* paths?

Maximum Length Paths?

- **What goes wrong?**
 - Our DP doesn't work because its path from s to t might use a path from s to u and edge from u to t
 - But path from s to u might in turn go through t
 - The path may no longer remain simple
- In fact, maximum length simple path is **NP-hard**
 - Hamiltonian path problem (i.e. is there a path of length $n - 1$ in a given undirected graph?) is a special case

All-Pairs Shortest Paths

- **Problem**

- **Input:** A directed graph $G = (V, E)$ with edge lengths ℓ_{vw} on each edge (v, w) and no negative cycles
- **Goal:** Compute the length of the shortest path from all vertices s to all other vertices t

- **Simple idea:**

- Run single-source shortest paths from each source s
- Running time is $O(n^4)$
- Actually, we can do this in $O(n^3)$ as well

All-Pairs Shortest Paths

- **Problem**

- **Input:** A directed graph $G = (V, E)$ with edge lengths ℓ_{vw} on each edge (v, w) and no negative cycles
- **Goal:** Compute the length of the shortest path from all vertices s to all other vertices t

- $OPT(u, v, k) =$ length of shortest simple path from u to v in which intermediate nodes from $\{1, \dots, k\}$
- **Exercise:** Write down the recursion formula of OPT such that given subsolutions, it requires $O(1)$ time
- **Running time:** $O(n^3)$ calls, $O(1)$ per call $\Rightarrow O(n^3)$

Chain Matrix Product

- **Problem**

- **Input:** Matrices M_1, \dots, M_n where the dimension of M_i is $d_{i-1} \times d_i$
- **Goal:** Compute $M_1 \cdot M_2 \cdot \dots \cdot M_n$

- But matrix multiplication is **associative**

- $A \cdot (B \cdot C) = (A \cdot B) \cdot C$
- So isn't the optimal solution going to call the algorithm for multiplying two matrices exactly $n - 1$ times?
- **Insight:** the time it takes to multiply two matrices depends on their dimensions

Chain Matrix Product

- Assume

- We use the brute force approach for matrix multiplication
- So multiplying $p \times q$ and $q \times r$ matrices requires $p \cdot q \cdot r$ operations

- Example

- M_1 is 5×10 , M_2 is 10×100 , and M_3 is 100×50
- $(M_1 \cdot M_2) \cdot M_3$ requires $5 \cdot 10 \cdot 100 + 5 \cdot 100 \cdot 50 = 30000$ ops
- $M_1 \cdot (M_2 \cdot M_3)$ requires $10 \cdot 100 \cdot 50 + 5 \cdot 10 \cdot 50 = 52500$ ops

Chain Matrix Product

- **Note**

- Our input is simply the dimensions d_0, d_1, \dots, d_n and not the actual matrices

- **Why is DP right for this problem?**

- Optimal substructure property

- Think of the final product computed, say $A \cdot B$

- A is the product of some prefix, B is the product of the remaining suffix

- For the overall optimal computation, each of A and B should be computed optimally

Chain Matrix Product

- $OPT(i, j) = \text{min ops required to compute } M_i \cdot \dots \cdot M_j$
 - Here, $1 \leq i \leq j \leq n$
 - **Q:** Why do we not just care about prefixes and suffices?
 - $(M_1 \cdot (M_2 \cdot M_3 \cdot M_4)) \cdot M_5 \Rightarrow \text{need to know optimal solution for } M_2 \cdot M_3 \cdot M_4$

$$OPT(i, j) = \begin{cases} 0 & i = j \\ \min\{OPT(i, k) + OPT(k + 1, j) + d_{i-1}d_kd_j : i \leq k < j\} & \text{if } i < j \end{cases}$$


- **Running time:** $O(n^2)$ calls, $O(n)$ time per call $\Rightarrow O(n^3)$

Chain Matrix Product

This slide is not in the scope of the course

- Can we do better?

- Surprisingly, yes. But not by a DP algorithm (that I know of)
- Hu & Shing (1981) developed $O(n \log n)$ time algorithm by reducing chain matrix product to the problem of “optimally” triangulating a regular polygon



Polygon
representation of
(AB)C

Polygon
representation of
A(BC)

Source: Wikipedia

Example

- A is 10×30 , B is 30×5 , C is 5×60
- The cost of each triangle is the product of its vertices
- Want to minimize total cost of all triangles