

Dynamic Programming determines optimal strategies among a range of possibilities typically putting together ‘smaller’ solutions. In some cases it is little more than a careful enumeration of the possibilities but can be organized to save effort by only computing the answer to a small problem once rather than many times.

In Graph Theory, the dynamic programming idea yields an effective algorithm for shortest (or indeed longest) paths or walks in a graph $G = (V, E)$. Let $w(i, j)$ denote the length of the edge (i, j) . Let $d^k(i, j)$ denote the length of the shortest i - j -walk from i to j that uses at most k edges. The base case can either be

$$d^0(i, i) = 0 \text{ and } d^0(i, j) = \infty \text{ for } i \neq j$$

or

$$d^1(i, j) = \begin{cases} w(i, j) & \text{for } (i, j) \in E \\ \infty & \text{for } (i, j) \notin E \end{cases}$$

The recurrence becomes for $k \geq 2$:

$$d^k(i, j) = \min_p \{d^{k-1}(i, p) + w(p, j) : (p, j) \in E\}$$

The parameter k representing the number of arcs is crucial. In our applications, the number of arcs will be replaced by the number of time periods. Typically we will determine the optimal solutions for i time periods to go from the optimal solutions for $i - 1$ time periods to go.

Some notation is helpful, rather fearsome at first glance, but useful in its generality.

Stages: these are the time periods. Typically they are indexed as $i = n, n - 1, \dots, 1, 0$ where i measures the number of stages left to go before the end of the process. This typically $i = 0$ refers to the final stage.

States Each state is a set of information describing the process at each stage. For many of our problems this can be a single integer.

s_i The state occupied by the process at stage i .

$D_i(s_i)$ The set of decisions that are allowed to be taken when the process is in state s_i at stage i .

d_i is the decision taken when the process is at stage i .

t_i is the transition function, defined on the decisions and states available at stage i . We have the equation

$$s_{i-1} = t_i(d_i, s_i)$$

that gives the outcome of the process at stage $i - 1$ (with $i - 1$ stages to go) after decision d_i is taken when in state s_i . The function depends only on d_i, s_i and not on how the process got to state s_i with i stages to go. We need this for dynamic programming to work.

We need some notion of cost or revenue.

$f_i(d_i, s_i)$ is the immediate or short term revenue resulting from taking decision d_i while in state s_i with i stages to go.

$v_i(s_i)$ denotes the long term revenue that can be derived (by selecting decisions optimally) in the remaining i stages given that we are in state s_i with i stages to go.

For most of our problems the function $f_i(d_i, s_i)$ is quite straightforward.

The first problem we consider is the planned building of power plants to meet projected demand. This comes from Bradley, Hax, and Magnanti. The relevant computations are in a file

dynamicbuild.xls (on the first sheet) located on our website. We are asked to plan the construction over a number of years. The construction costs for a power plant depend on the year of construction and we are allowed to construct up to three plants in a given year. Moreover if we are doing any construction in a given year, we incur an extra administrative cost of \$1.5M. For us stages are years and states give the number of plants constructed.

stage	year	cumulative demand	cost per plant
6 →			
	1981	1	5.4
5 →			
	1981	2	5.6
4 →			
	1982	4	5.8
3 →			
	1983	6	5.7
2 →			
	1984	7	5.5
1 →			
	1985	8	5.2

Thus for example at stage 4 we must have $s_4 \geq 2$ and the possible decisions are those that result in $s_3 \geq 4$. The decision d_i refers to the number of plants constructed in that year so that $t_i(d_i, s_i) = d_i + s_i$. The ‘value’ (or in this case ‘cost’) is

$$f_4(d_4, s_4) = \begin{cases} 0 & d_4 = 0 \\ 5.8d_4 + 1.5 & d_4 > 0 \end{cases}$$

You note that $v_6(0)$ represents the solution, namely the cost of being in state 0 (i.e. with no power plants constructed) with 6 years to go.

What follows is the general algorithm (quite easy!) whose main drawback is the potential for too many states and too many stages.

Bellman’s Equations.

$$v_i(s_i) = \max\{f_i(d_i, s_i) + v_{i-1}(t_i(d_i, s_i)) : d_i \in D_i(s_i)\}.$$

We are maximizing the cumulative revenue (or minimizing the cost)

$$f_n(d_n, s_n) + f_{n-1}(d_{n-1}, s_{n-1}) + \dots + f_0(d_0, s_0) : d_i \in D_i(s_i) \text{ for } i = n, n-1, \dots, 1.$$

This simple formulation adapts to many situations and also generalizations. One simple generalization considers the discount rate which while unfamiliar to students is a business practice. Roughly speaking each enterprise can estimate its own discount rate as $e^{-\mu} = \beta < 1$ so that \$1 one year from now is worth $\$e^{-\mu} = \beta$ today (i.e. slightly less). One can view this as a consequence of interest rates but it occurs as a result of the investment opportunities of a business. Thus the discount rate varies depending on the business.

Bellman’s Equations with discounted future returns.

$$v_i(s_i) = \max\{f_i(d_i, s_i) + e^{-\mu}v_{i-1}(t_i(d_i, s_i)) : d_i \in D_i(s_i)\}.$$

We are maximizing the cumulative revenue (or minimizing the cost)

$$f_n(d_n, s_n) + e^{-\mu} f_{n-1}(d_{n-1}, s_{n-1}) + \cdots + e^{-n\mu} f_0(d_0, s_0) : d_i \in D_i(s_i) \text{ for } i = n, n-1, \dots, 1.$$

Again in the file `dynamicbuild.xls` (on the second sheet) we can see the power plant example worked out with discounting. The file is set up so that you could adjust the discount rate but note that the column corresponding to the actual optimal decisions at each stage are in fact hand entered.