Lesson 5. Introduction to Dynamic Programming

1 The knapsack problem

Example 1. You are a thief deciding which precious metals to steal from a vault:

	Metal	Weight (kg)	Value
1	Gold	3	11
2	Silver	2	7
3	Platinum	4	12

You have a knapsack that can hold at most 8kg. If you decide to take a particular metal, you must take all of it. Which items should you take to maximize the value of your theft?

- This example problem is pretty small, so we can easily solve it by inspection
- Maximum total value: 23
- Items that give the maximum total value: gold + platinum (1 + 3)
- We can also formulate this problem as a longest path problem:

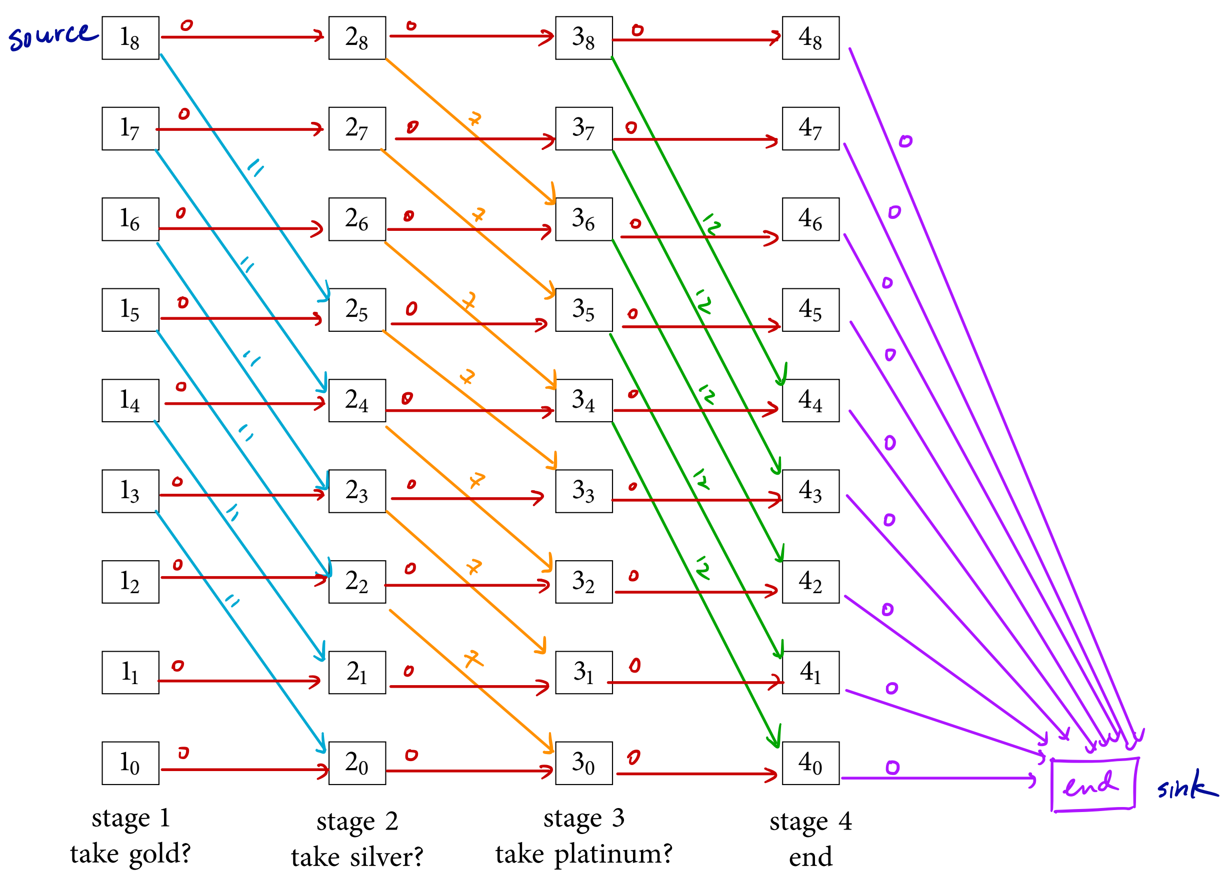

- We consider filling up our knapsack in **stages**
- In stage $t = 1, 2, 3$, we decide whether to take metal t
- The last stage (stage 4) represents the end of our decision process
- Node t_n represents

  there are n kgs left at stage t, when metals t, t+1, … are remaining

- The edges represent the **decisions** we can make
- Suppose we are deciding whether to take metal 2 (silver), and we have 5 kgs of space left in our knapsack
- Two possible decisions:
  1. Take metal 2
     - This is represented by the edge $(2_5, 3_3)$
     - This decision has a value of 7, so we use this as the length of this edge
  2. Don't take metal 2
     - This is represented by the edge $(2_5, 3_5)$
     - This decision has a value of 0, so we use this as the length of this edge
- We can complete the rest of the digraph in a similar fashion
- **Key observation.** Finding an optimal solution to the knapsack problem is equivalent to finding the longest path in this graph from node 1_8 to some stage 4 node
  - In this example, the longest path is $1_8 \to 2_5 \to 3_5 \to 4_1$ with a length of 23
  - The longest path length tells us:

    the maximum total value of items we can take in our knapsack

  - The nodes and edges in the longest path tell us:

    $(1_8, 2_5)$ take gold

    $(2_5, 3_5)$ don't take silver

    $(3_5, 4_1)$ take platinum

- To reformulate this as a shortest path problem:
  - Negate all edge lengths
  - Connect all stage 4 nodes to an "end" node with edges of length 0
  - Find the shortest path from node 1_8 to the finish node

2 Dynamic programming

- A **dynamic program** (DP) is a mathematical model that captures situations where decisions are made sequentially in order to optimize some objective
- In particular:
  - DPs divide problems into **stages** with a **decision** required at each stage
  - Each stage has a number of **states** – the possible conditions of the system at that stage
  - A decision at each stage transforms the current state into a state in the next stage with some associated **cost** or **reward**
- DPs come in several different flavors, and can be described in various ways
- For now, we will think of a DP as a specially-structured shortest/longest path problem

Dynamic program – shortest/longest path representation

- **Stages** $t = 1, 2, \dots, T$ and **states** $n = 0, 1, 2, \dots, N$
- Directed graph:

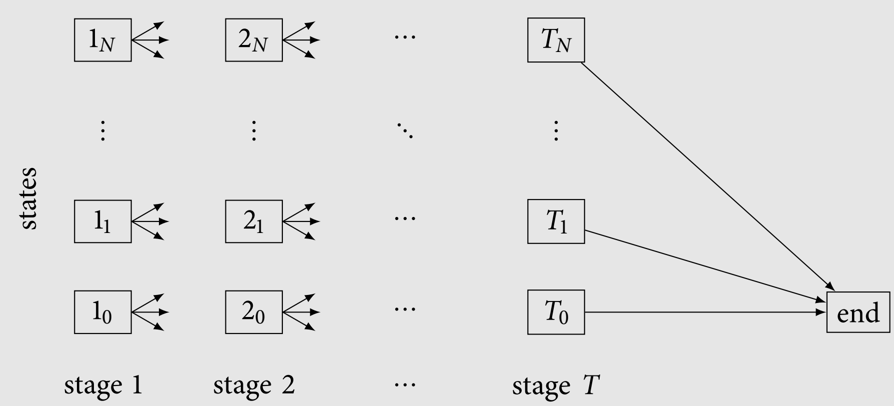

  - Node $t_n \longleftrightarrow$ being in state n at stage t
    - Nodes for the tth stage are put in the tth column
  - Edge $(t_n, (t+1)_m) \longleftrightarrow$ the **decision** to go to state m from state n at stage t
    - Length of this edge = **cost** or **reward** of making this decision
    - An edge must connect a node in the tth column to a node in the $(t+1)$st column
  - Nodes for the last stage are connected to an "end" node
    - Typically: all nodes in last stage are connected to the end node with edge lengths of 0
- Shortest/longest path problem:
  - Source node = one of the first stage nodes (depends on the problem)
  - Target node = end node
  - Edge lengths correspond to rewards $\Longrightarrow$ Find the longest path from source to target
  - Edge lengths correspond to costs $\Longrightarrow$ Find the shortest path from source to target

Example 2. The Simplexville Police Department wants to determine how to assign patrol cars to each precinct in Simplexville. Each precinct can be assigned 0, 1, or 2 patrol cars. The number of crimes in each precinct depends on the number of patrol cars assigned to each precinct:

	Number of patrol cars assigned to precinct		
Precinct	0	1	2
1	14	10	7
2	25	19	16
3	20	14	11

The department has 5 patrol cars. The department's goal is to minimize the total number of crimes across all 3 precincts. Formulate this problem as a dynamic program by giving its shortest/longest path representation.

Stage t represents assigning patrol cars to precinct t.

Node t_n represents having n patrol cars left at stage t, with precincts $t, t+1, \ldots$ remaining.

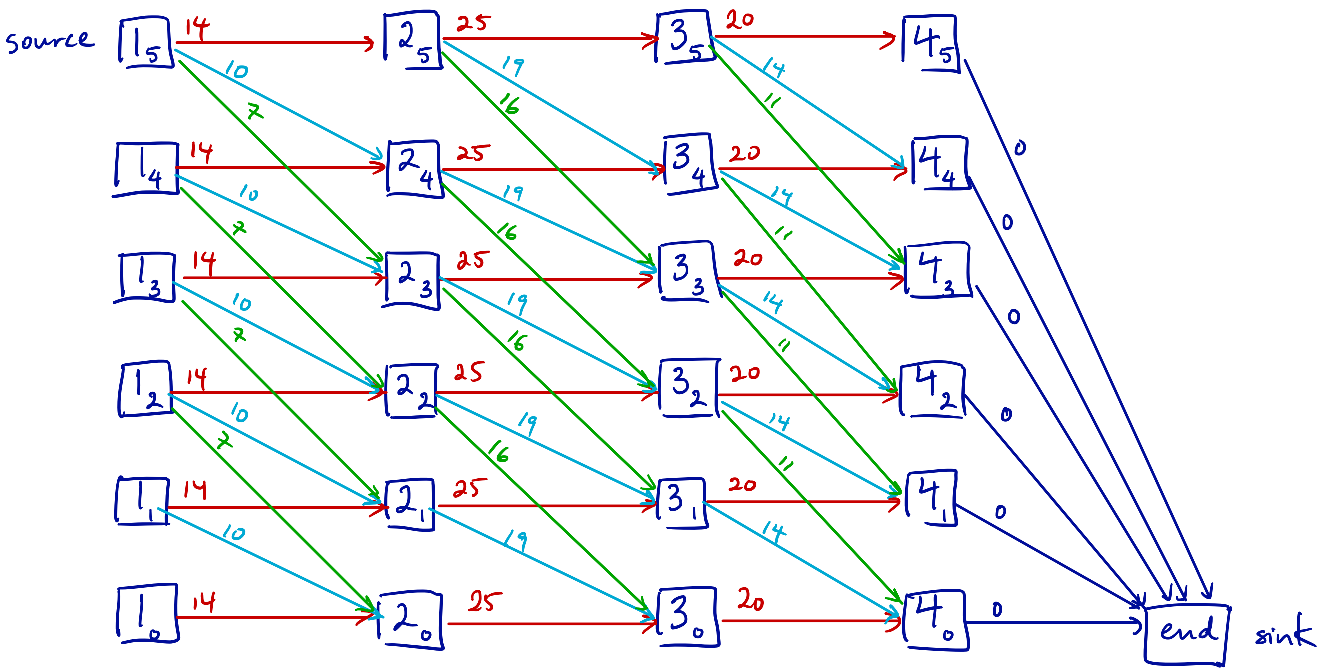

We want to find the shortest path from source to sink in the above graph.

It turns out:

shortest path length $= 37$

shortest path nodes $= 1_5 \to 2_3 \to 3_2 \to 4_0 \to$ end

$\Rightarrow$ Minimum total crimes $= 37$

Patrol assignments:

2 cars to precinct 1

1 cars to precinct 2

2 cars to precinct 3

Example 3. The Dijkstra Brewing Company is planning production of its new limited run beer, Primal Pilsner. The company must supply 1 batch next month, then 2 and 4 in successive months. Each month in which the company produces the beer requires a factory setup cost of $5,000. Each batch of beer costs $2,000 to produce. Batches can be held in inventory at a cost of $1,000 per batch per month. Capacity limitations allow a maximum of 3 batches to be produced during each month. In addition, the size of the company's warehouse restricts the ending inventory for each month to at most 3 batches. The company has no initial inventory.

The company wants to find a production plan that will meet all demands on time and minimizes its total production and holding costs over the next 3 months. Formulate this problem as a dynamic program by giving its shortest/longest path representation.

Let stage t represent month t $(t=1,2,3)$ or the end of the decision making process $(t=4)$

Let node t_n represent having n batches in inventory at the beginning of month t $(n=0,1,2,3)$

Find a shortest path:

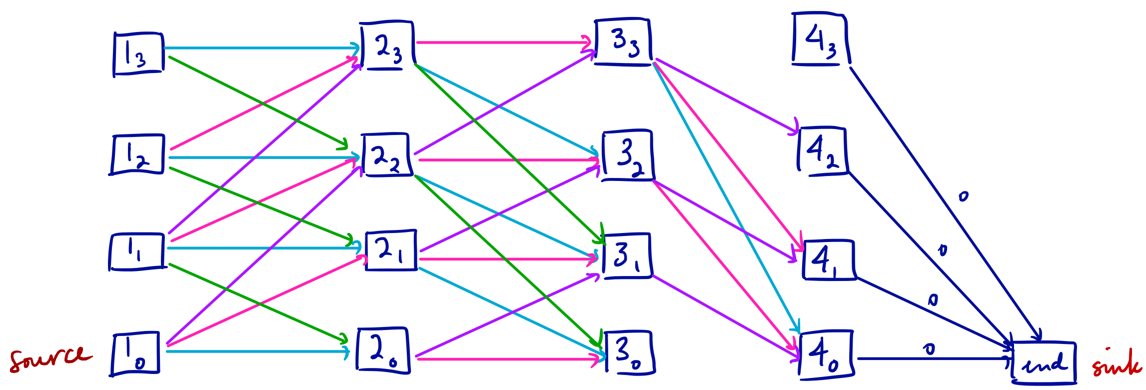

Stage	Production	Edges	Edge length
1	0	$(1_n, 2_{n-1})$ $n=1,2,3$	$1(n-1)$
1	1	$(1_n, 2_n)$ $n=0,1,2,3$	$5+2(1)+1(n)$
1	2	$(1_n, 2_{n+1})$ $n=0,1,2$	$5+2(2)+1(n+1)$
1	3	$(1_n, 2_{n+2})$ $n=0,1$	$5+2(3)+1(n+2)$
2	0	$(2_n, 3_{n-2})$ $n=2,3$	$1(n-2)$
2	1	$(2_n, 3_{n-1})$ $n=1,2,3$	$5+2(1)+1(n-1)$
2	2	$(2_n, 3_n)$ $n=0,1,2,3$	$5+2(2)+1(n)$
2	3	$(2_n, 3_{n+1})$ $n=0,1,2$	$5+2(3)+1(n+1)$
3	0	not possible!	
3	1	$(3_n, 4_{n-3})$ $n=3$	$5+2(1)+1(n-3)$
3	2	$(3_n, 4_{n-2})$ $n=2,3$	$5+2(2)+1(n-2)$
3	3	$(3_n, 4_{n-1})$ $n=1,2,3$	$5+2(3)+1(n-1)$

Shortest path length $= 30$
$=$ minimum total production/holding cost over the next 3 mos.

Shortest path $=$
$1_0 \to 2_0 \to 3_1 \to 4_0$

$\Rightarrow$ Produce 1 in mo. 1
3 in mo. 2
3 in mo. 3

Example 4. The State of Maryland is conducting research on reducing traffic on the I-270 corridor. Three research teams are currently trying three different approaches to the problem. It has been estimated that the probability that the respective teams – call them 1, 2, and 3 – will fail is 0.40, 0.60, and 0.80, respectively.

In order to decrease probability of failure, the state wants to assign two additional researchers to the project. The following table gives the estimated probability that the respective teams will fail when 0, 1, or 2 additional researchers are added to that team:

	Probability of failure		
Number of new researchers	Team 1	Team 2	Team 3
0	0.40	0.60	0.80
1	0.20	0.40	0.50
2	0.15	0.20	0.30

The state wants to determine how to allocate the two additional researchers in order to minimize the probability that all three teams will fail. Formulate this problem as a dynamic program by giving its shortest/longest path representation.

Let stage t represent the decision to add researchers to team t $(t = 1, 2, 3)$ or the end of the decision-making process $(t = 4)$

Let node t_n represent having n researchers left to add at stage t $(n = 0, 1, 2)$

Find the shortest path:

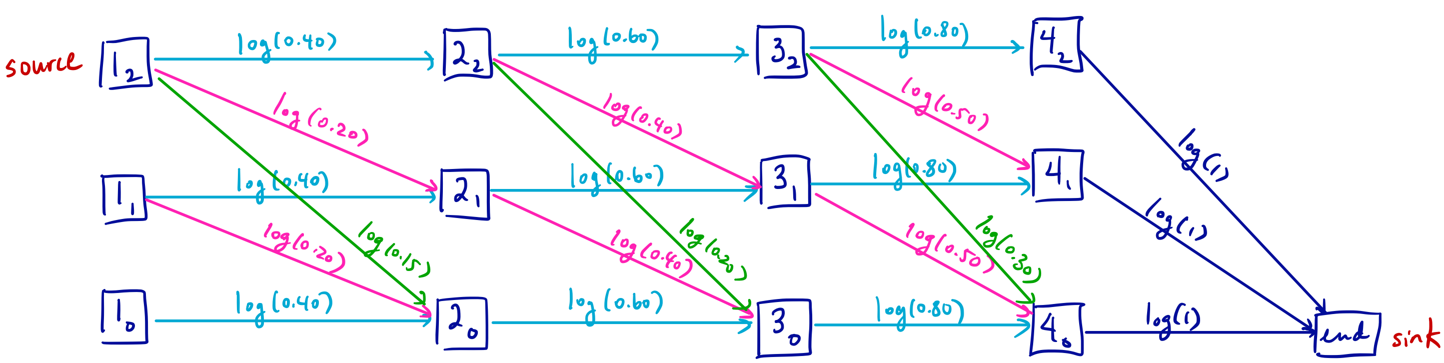

Shortest path length ≈ -2.813 $\Rightarrow$ Minimum probability of all 3 teams failing $= e^{-2.813} = 0.06$

Shortest path: $1_2 \to 2_1 \to 3_1 \to 4_0$ $\Rightarrow$ Assign 1 researcher to Team 1, 1 researcher to Team 3

Products of costs and rewards

- Consider Example 2
- Let
$$p_i = \text{probability that team } i \text{ fails} \quad \text{for } i = 1, 2, 3$$
- We want to minimize the probability that all teams fail, or:

$$\text{minimize} \quad p_1 p_2 p_3$$

- We know how to find a shortest path – a path with the smallest sum of edge lengths
- What if wanted to find a path with the smallest product of edge lengths?
- Recall that the logarithm function is monotonic increasing:

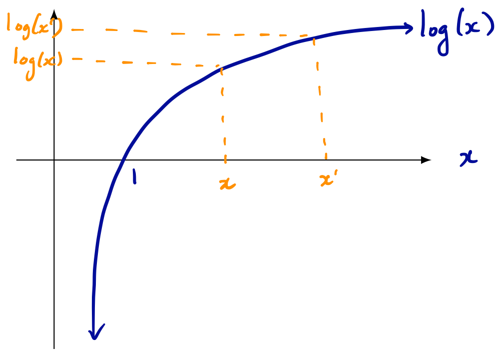

$$x < x' \iff \log(x) < \log(x')$$

- Therefore:

$$\text{minimize} \quad p_1 p_2 p_3 \iff \text{minimize} \quad \log(p_1 p_2 p_3)$$
$$\iff \text{minimize} \quad \log(p_1) + \log(p_2) + \log(p_3)$$

- So instead of setting the edge lengths equal to the probabilities, we can set the edge lengths equal to the logarithms of the probabilities
- Just make sure to invert the logarithms when interpreting the shortest path length

Example 5. To graduate from Simplexville University, Angie needs to pass at least one of the three courses she is taking this semester: literature, finance, and statistics. Angie's busy schedule of extracurricular activities allows her to spend only 4 hours per week on studying. Angie's probability of passing each course depends on the number of hours she spends studying for the course:

	Probability of passing course		
Hours of studying per week	Literature	Finance	Statistics
0	0.20	0.25	0.10
1	0.30	0.30	0.30
2	0.35	0.33	0.40
3	0.38	0.35	0.44
4	0.40	0.38	0.50

Angie wants to maximize the probability that she passes at least one of these three courses. Formulate this problem as a dynamic program by giving its shortest/longest path representation.

Hint. Why is maximizing the probability of passing at least one course equivalent to minimizing the probability of failing all three courses? $\Pr\{\#\text{ courses passed} \geq 1\} = 1 - \Pr\{\#\text{ courses passed} = 0\} = 1 - \Pr\{\text{fail all 3 courses}\}$

Let stage t represent assigning hours/wk to course t $(t = 1, 2, 3)$ or the end of the decision making process $(t = 4)$

Let node t_n represent having n hours/wk left to assign at stage t $(n = 0, 1, 2, 3, 4)$

Find the shortest path:

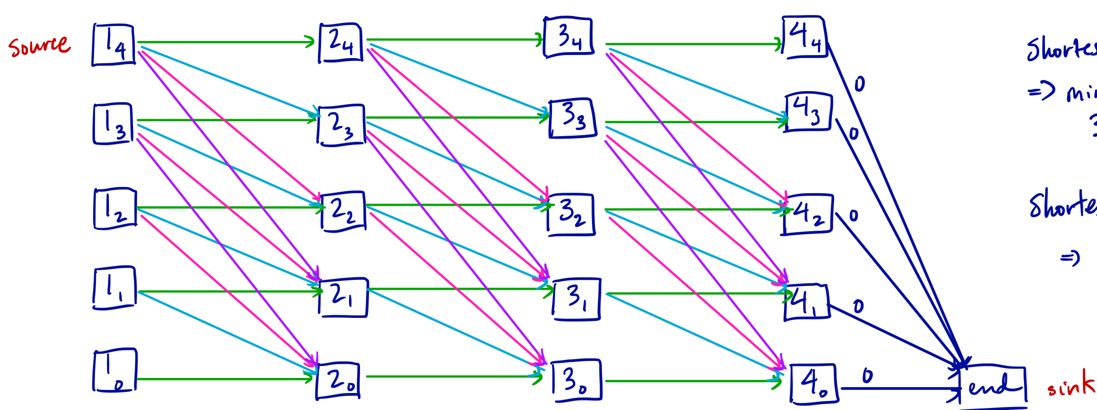

Shortest path length ≈ -1.23
$\Rightarrow$ minimum prob. of failing all 3 courses $= e^{-1.23} = 0.293$

Shortest path: $1_4 \rightarrow 2_2 \rightarrow 3_2 \rightarrow 4_0$
$\Rightarrow$ 2 hours to literature
2 hours to statistics

Course/Stage	Hrs/wk	Edges	Edge lengths
1	0	$(1_n, 2_n)$ $n = 0, 1, \dots, 4$	$\log(0.80)$
1	1	$(1_n, 2_{n-1})$ $n = 1, \dots, 4$	$\log(0.70)$
1	2	$(1_n, 2_{n-2})$ $n = 2, \dots, 4$	$\log(0.65)$
1	3	$(1_n, 2_{n-3})$ $n = 3, 4$	$\log(0.62)$
1	4	$(1_n, 2_{n-4})$ $n = 4$	$\log(0.60)$
2	0	$(2_n, 3_n)$ $n = 0, 1, \dots, 4$	$\log(0.75)$
2	1	$(2_n, 3_{n-1})$ $n = 1, \dots, 4$	$\log(0.70)$
2	2	$(2_n, 3_{n-2})$ $n = 2, \dots, 4$	$\log(0.67)$
2	3	$(2_n, 3_{n-3})$ $n = 3, 4$	$\log(0.65)$
2	4	$(2_n, 3_{n-4})$ $n = 4$	$\log(0.62)$

Course/Stage	Hrs/wk	Edges	Edge lengths
3	0	$(3_n, 4_n)$ $n = 0, 1, \dots, 4$	$\log(0.90)$
3	1	$(3_n, 4_{n-1})$ $n = 1, \dots, 4$	$\log(0.70)$
3	2	$(3_n, 4_{n-2})$ $n = 2, \dots, 4$	$\log(0.60)$
3	3	$(3_n, 4_{n-3})$ $n = 3, 4$	$\log(0.56)$
3	4	$(3_n, 4_{n-4})$ $n = 4$	$\log(0.50)$