Dynamic Programming

Lecture 13 (5/31/2017)

- A Forest Thinning Example -

Projected yield (m3/ha) at age 20 as function of action taken at age 10

Age 10				
Beginning volume	Volume thinned	Residual volume	Ten-year growth	Volume at age 20
260	0	260	390	650
260	50	210	325	535
260	100	160	250	410

Projected yield (m3/ha) at age 30 as function of beginning volume at age 20 and action taken at age 20

Age 20				
Beginning volume	Volume thinned	Residual volume	Ten-year growth	Volume at age 30
	0	650	200	850
650	150	500	250	750
	200	450	200	650
	0	535	215	750
535	100	435	215	650
	175	360	140	500
	0	410	190	600
410	75	335	165	500
	150	260	140	400

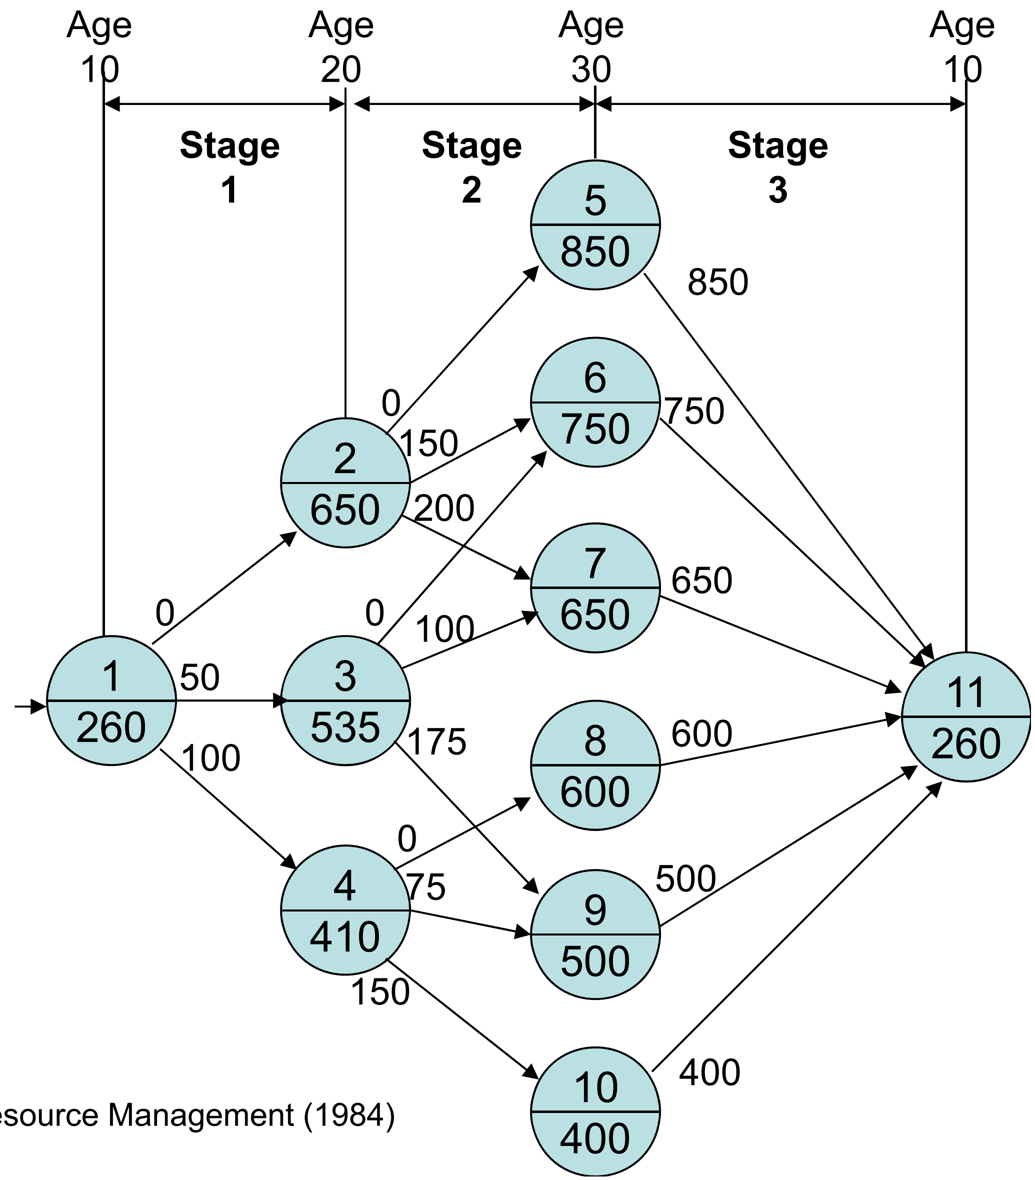

Source: Dykstra's Mathematical Programming for Natural Resource Management (1984)

Dynamic Programming Cont.

- Stages, states and actions (decisions)
- Backward and forward recursion

$x_i =$ decision variable: immediate destination node at stage i;

$f_i(s, x_i) =$ the maximum total volume harvested during all remaining stages, given that the stand has reached the state corresponding to node s at stage i and the immediate destination node is x_i;

$x_i^* =$ the value of x_i that maximizes $f_i(s, x_i)$; and

$f_i^*(s) =$ the maximum value of $f_i(s, x_i)$, i.e., $f_i^*(s) = f_i(s, x_i^*)$.

Solving Dynamic Programs

- Recursive relation at stage *i* (Bellman Equation):

$$f_i^*(s) = \max_{x_i} \left\{ H_{s,x_i} + f_{i+1}^*(x_i) \right\},$$

where H_{s,x_i} is the volume harvested by thinning at the beginning of stage i in order to move the stand from state s at the beginning of stage i to state x_i at the end of stage i.

Dynamic Programming

- Structural requirements of DP
  - The problem can be divided into stages
  - Each stage has a finite set of associated states (discrete state DP)
  - The impact of a policy decision to transform a state in a given stage to another state in the subsequent stage is deterministic
  - Principle of optimality: given the current state of the system, the optimal policy for the remaining stages is independent of any prior policy adopted

Examples of DP

- The Floyd-Warshall Algorithm (used in the Bucket formulation of ARM):

Let $c(i, j)$ denote the length (or weight) of the path between node i and j; and

Let $s(i, j, k)$ denote the length (or total weight) of the shortest path between nodes i and j going through intermediate node 1, 2, ...,k.

Then, the following recursion will give the length of the shortest paths between all pairs of nodes in a graph:

$$s(i, j, k) = \begin{cases} c(i, j) \text{ if } k = 1 \\ \min\left\{ s(i, j, k-1),\ s(i, k, k-1) + s(k, j, k-1) \right\} \text{ otherwise} \end{cases}$$

Examples of DP cont.

- Minimizing the risk of losing an endangered species (non-linear DP):

Probability of Project Failure

	Funding level (state)		
Project (stage)	$0	$1M	$2M
1	50%	30%	20%
2	70%	50%	30%
3	80%	50%	40%

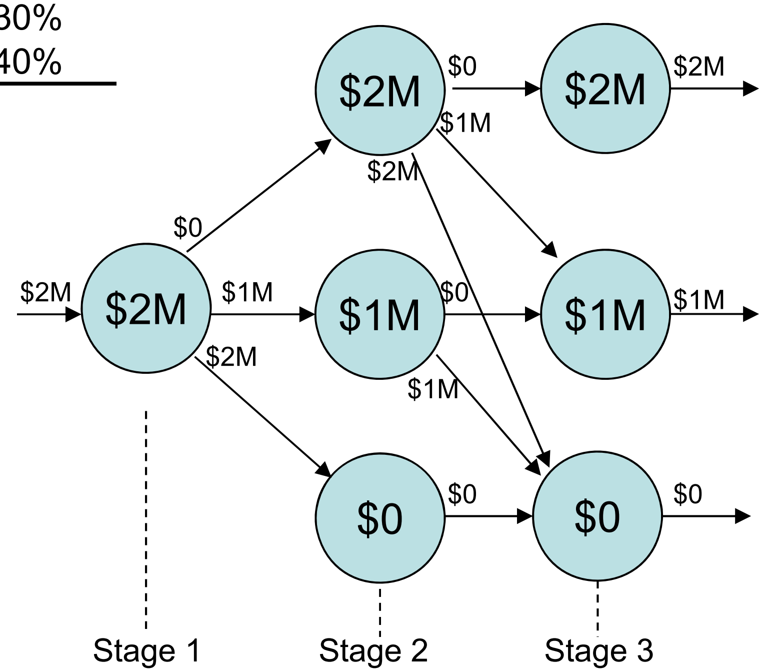

Source: Buongiorno and Gilles' (2003) Decision Methods for Forest Resource Management

Minimizing the risk of losing an endangered species (DP example)

- Stages (t=1,2,3) represent $ allocation decisions for Project 1, 2 and 3;
- States (i=0,1,2) represent the budgets available at each stage t;
- Decisions (j=0,1,2) represent the budgets available at stage t+1;
- $p_t(i, j)$ denotes the probability of failure of Project t if decision j is made (i.e., i-j is spent on Project t); and
- $V_t^*(i)$ is the smallest probability of total failure from stage t onward, starting in state i and making the best decision j*.

Then, the recursive relation can be stated as:

$$V_t^*(i) = \min_j \left\{ p_t(i, j) \cdot V_{t+1}^*(j) \right\}$$

Markov Chains
(based on Buongiorno and Gilles 2003)

Volume by stand states

State i	Volume (m3/ha)
L	<400
M	400-700
H	700<

20-yr Transition probabilities w/o management

Start	End State j		
State i	L	M	H
L	40%	60%	0%
M	0%	30%	70%
H	5%	5%	90%

Transition probability matrix *P*

p_t = probability distribution of stand states in period t

$$p_t = p_{t\text{-}1}P \qquad \forall t$$

- Vector p_t converges to a vector of steady-state probabilities p^*.
- p^* is independent of p_0!

Markov Chains cont.

- Berger-Parker Landscape Index

$$BP_t = \frac{p_{Lt} + p_{Mt} + p_{Ht}}{\max(p_{Lt}, p_{Mt}, p_{Ht})}$$

,where p_{it} is the probability that the stand is in state $i = L,\ M$ or H in period t.

- Mean Residence Times

$$m_i = \frac{D}{(1 - p_{ii})}$$

,where D is the length of each period, and p_{ii} is the probability that a stand in state i in the beginning of period t stays there till the end of period t.

- Mean Recurrence Times

$$m_{ii} = \frac{D}{\pi_i}$$

,where π_i is the steady-state probability of state i.

Markov Chains cont.

- Forest dynamics (expected revenues or biodiversity with vs. w/o management:

20-yr Transition probabilities w/o management

Start State i	End State j L	M	H
L	40%	60%	0%
M	0%	30%	70%
H	5%	5%	90%

20-yr Transition probabilities w/ management

Start State i	End State j L	M	H
L	40%	60%	0%
M	0%	30%	70%
H	40%	60%	0%

Expected long-term biodiversity: $B = \pi_L B_L + \pi_M B_M + \pi_H B_H$

Expected long-term periodic income: $R = \pi_L R_L + \pi_M R_M + \pi_H R_H$

Markov Chains cont.

- Present value of expected returns

$$V_{i,t+1} = R_i + \frac{1}{(1+r)^{20}}\left(p_{iL}V_{Lt} + p_{iM}V_{Mt} + p_{iH}V_{Ht} \right)$$

where $V_{i,t}$ is the present value of the expected return from a stand in state in state $i = L,\ M$ or H managed with a specific harvest policy with t periods to go before the end of the planning horizon. R_i is the immediate return from managing a stand in state i with the given harvest policy. p_{iL}, p_{iM}, p_{iH} are the probabilities that a stand in state i moves to state L, M or H, respectively.

Markov Decision Processes

20-yr Transition probabilities depending on the harvest decision

No Cut			
Start State i	End State j: L	M	H
L	40%	60%	0%
M	0%	30%	70%
H	5%	5%	90%

Cut			
Start State i	End State j: L	M	H
L	40%	60%	0%
M	40%	60%	0%
H	40%	60%	0%

$$V_{i,t+1}^{*} = \max_{j}\left\{ R_{ij} + \frac{1}{(1+r)^{20}}\left(p_{iLj}V_{Lt}^{*} + p_{iMj}V_{Mt}^{*} + p_{iHj}V_{Ht}^{*} \right) \right\}$$

where $V_{i,t+1}^{*}$ is the highest present value of the expected return from a stand in state in state $i = L,\ M$ or H managed with a specific harvest policy with t periods to go before the end of the planning horizon. R_{ij} is the immediate return from managing a stand in state i with harvest policy j.

$p_{iLj}, p_{iMj}, p_{iHj}$ are the probabilities that a stand in state i moves to state L, M or H, respectively if the harvest policy is j.