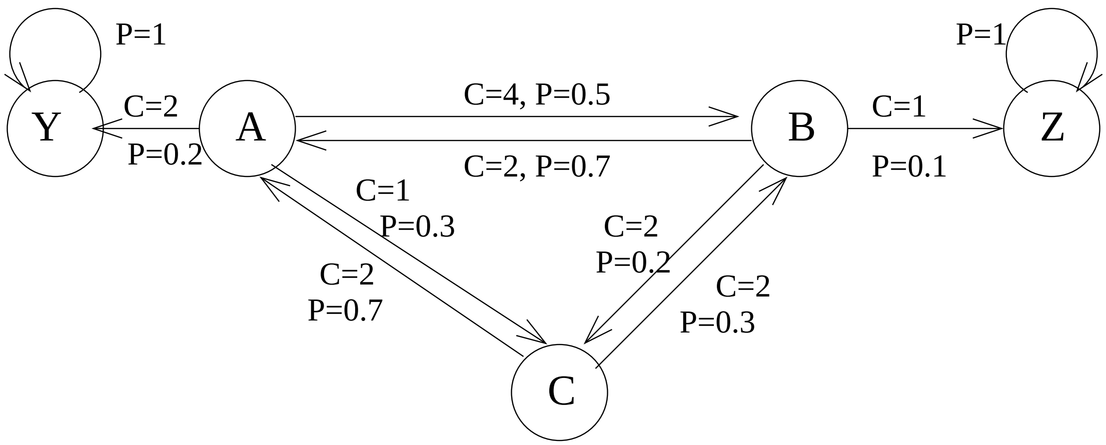

Figure 1: Markov process with final states

Suppose that you have a Markov process, in which some of the states $F_1 \dots F_k$ have the following properties.

- Each F_i transitions with probability 1 to itself.
- From any state S in the Markov process, there is a path of non-zero probability to one of the F_i.

The states $F_1 \dots F_k$ are called *final states*; once you have reached a final state, you stay there. Moreover, starting from any state S in the process, with probability 1 you will eventually end in a final state, because if you wander through the process long enough, you will eventually go down one of the paths that leads to a final state. For example, in the Markov model shown in figure 10.1 in the textbook, the state $S6$ is the unique final state. In figure 1 above, Y and Z are final states.

There are two natural problems that arise in a Markov model with final states:

A. From any given starting state S, what is the expected time until you reach a final state?

B. Suppose the process has more than one final state $F_1 \dots F_k$. If you start in state S, what is the probability that you will end in F_1, in F_2 and so on?

The two problems actually have essentially the same solution. For problem (A), we can additionally allow different arcs to have different costs, and ask for the expected cost to reach a final state. The expected time is then just the special case where all arcs have cost 1.

For instance, figure 1 shows a Markov process with five states: Y, Z are final states and A, B, C are not. We will show how to solve problems (A) and (B) in this particular case; the generalization should be obvious.

For problem (A) we define three variables: a is the expected cost to reach a final state starting from state A, b is the expected cost to reach a final state starting from state B, and c is the expected cost to reach a final state starting from state C. We can then reason as follows. Suppose that at a given time we are in state A. Then:

There is 0.2 probability that we will transition to Y, with total cost 2.

There is 0.5 probability that we will transition to B; the transition will cost 4, and the expected cost of the rest of the path to the final state is just b.

There is 0.3 probability that we will transition to C; the transition will cost 1, and the expected cost of the rest of the path to the final state is just c.

Therefore $a = 0.2 \cdot 2 + 0.5 \cdot (4 + b) + 0.3 \cdot (1 + c)$. Likewise $b = 0.1 \cdot 1 + 0.7 \cdot (2 + a) + 0.2 \cdot (2 + b)$ and $c = 0.7 \cdot (2 + a) + 0.3 \cdot (2 + b)$. We thus have the following system of linear equations:

$$\begin{bmatrix} 1.0 & -0.5 & -0.3 \\ -0.7 & 1.0 & -0.2 \\ -0.7 & -0.3 & 1.0 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 2.7 \\ 1.9 \\ 2.0 \end{bmatrix}$$

with the solution $a = 18.05, b = 18.57, c = 20.21$.

We can reduce problem (B) to problem (A) by the following trick. Consider a new process in which all the transition probabilities are the same as before but the costs are different. Specifically, the arcs that lead to state Y all have cost 1; and all the other arcs have cost 0. Therefore, any path that ends up in state Y has total cost 1 and any path that ends up in state Z has total cost 0. Therefore the probability that you will end up in state Y is exactly equal to the expected cost of reaching a final state in this new process. This gives us the new system of linear equations:

$$\begin{bmatrix} 1.0 & -0.5 & -0.3 \\ -0.7 & 1.0 & -0.2 \\ -0.7 & -0.3 & 1.0 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0.2 \\ 0.0 \\ 0.0 \end{bmatrix}$$

with the solution $a = 0.76, b = 0.68, c = 0.74$.