IEOR 3106: Introduction to Operations Research: Stochastic Models

Class Lecture Notes: Tuesday, October 15, 2013

Continuous-Time Markov Chains, Ross Chapter 6

1. Pooh Bear and the Three Honey Trees.

These notes are extracted from the longer set of notes on continuous-time Markov chains, which we shall follow closely in class.

A bear of little brain named Pooh is fond of honey. Bees producing honey are located in three trees: tree A, tree B and tree C. Tending to be somewhat forgetful, Pooh goes back and forth among these three honey trees randomly (in a Markovian manner) as follows: From A, Pooh goes next to B or C with probability 1/2 each; from B, Pooh goes next to A with probability 3/4, and to C with probability 1/4; from C, Pooh always goes next to A. Pooh stays a random time at each tree. Pooh stays at each tree an exponential length of time, with the mean being 5 hours at tree A or B, but with mean 4 hours at tree C. (For simplicity, assume that the travel times can be ignored. And assume that Pooh does this continuously throughout all time, without a break.)

(a) Construct a CTMC enabling you to find the limiting proportion of time that Pooh spends at each honey tree.

(b) What is the average number of trips per day Pooh makes from tree B to tree A?

ANSWERS:

(a) Find the limiting fraction of time that Pooh spends at each tree.

These problems are part of a longer set of lecture notes, which have been posted on the web page. The focus here is on different ways to model. We are thus focusing on Section 3 of the notes.

For this problem formulation, it is natural to use the SMP (semi-Markov process) formulation of a CTMC (continuous-time Markov chain), involving the embedded DTMC and the mean holding times in each state. (See Sections 6.2 and 7.6 of Ross.) We thus define the embedded transition matrix P directly and the mean holding times $1/\nu_i$ directly.

Note that this problem is formulated directly in terms of the DTMC, describing the random motion at successive transitions, so it is natural to use this initial modelling approach. Here the transition matrix for the DTMC is

$$P = \begin{array}{c} A \\ B \\ C \end{array} \begin{pmatrix} 0 & 1/2 & 1/2 \\ 3/4 & 0 & 1/4 \\ 1 & 0 & 0 \end{pmatrix} .$$

In the displayed transition matrix P, we have only labelled the rows. The columns are assumed to be labelled in the same order.

In general, the steady-state probability is

$$\alpha_i = \frac{\pi_i(1/\nu_i)}{\sum_j \pi_j(1/\nu_j)}$$

In this case, the steady state probability vector of the discrete-time Markov chain is obtained by solving $\pi = \pi P$, yielding

$$\pi = (\frac{8}{17}, \frac{4}{17}, \frac{5}{17}) \ .$$

Then the final steady-state distribution, accounting for the random holding times is

$$\alpha = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4}) \ .$$

You could alternatively work with the infinitesimal transition rate matrix Q. If we want to define the infinitesimal transition matrix Q (Ross uses lower case q), then we can do so by setting

$$Q_{i,j} = \nu_i P_{i,j} \quad \text{for} \quad i \neq j \ .$$

As usual, the diagonal elements $Q_{i,i}$ are set equal to minus the i^{th} row sum; i.e.,

$$Q_{i,i} = -\sum_{j:j\neq i} Q_{i,j} \ .$$

But we do not need the Q matrix to solve for the steady-state distribution. We can use the SMP representation. If we do define Q, then we can alternatively obtain the steady-state probability vector α above by solving the equation $\alpha Q = 0$, where the elements α_i are required to sum to 1. Note that this is just a system of linear equations, just like $\pi = \pi P$. You should work to understand why we here have 0 instead of α for the vector.

(b) What is the average number of trips Pooh makes per day from tree B to tree A?

The long-run fraction of time spent at B is $1/4$, by part (a). Thus, on average, Pooh spends 6 hours per day at tree B. When at tree B, the rate of trips from B is $1/5$ per hour (the reciprocal of 5 hours), and thus, on average, Pooh makes $6/5 = 1.2$ trips per day from tree B. However, $3/4$ of the trips from tree B are to tree A, so the average number of trips per day from B to A is $(6/5) \times (3/4) = (18/20) = 0.9$.

---

2. Copier Breakdown and Repair.

Consider two copier machines that are maintained by a single repairman. Machine i functions for an exponentially distributed amount of time with mean $1/\gamma_i$, and thus rate γ_i, before it breaks down. The repair times for copier i are exponential with mean $1/\beta_i$, and thus rate β_i, but the repairman can only work on one machine at a time. Assume that the machines are repaired in the order in which they fail. Suppose that we wish to construct a CTMC model of this system, with the goal of finding the long-run proportions of time that each copier is working and the repairman is busy. How can we proceed?

(a) Let $\{X(t) : t \geq 0\}$ be a stochastic process, where $X(t)$ represents the number of working machines at time t. Is $\{X(t) : t \geq 0\}$ a Markov process?

(b) Formulate a CTMC describing the evolution of the system.

(c) Suppose that $\gamma_1 = 1$, $\beta_1 = 2$, $\gamma_2 = 3$ and $\beta_2 = 4$. Find the stationary distribution.

(d) Now suppose, instead, that machine 1 is much more important than machine 2, so that the repairman will always service machine 1 if it is down, regardless of the state of machine 2. Formulate a CTMC for this modified problem and find the stationary distribution.

ANSWERS:

(a) Let $\{X(t) : t \geq 0\}$ be a stochastic process, where $X(t)$ represents the number of working machines at time t. Is $\{X(t) : t \geq 0\}$ a Markov process?

This process is not a Markov process. To be a Markov process, we need the conditional distribution of a future state, given a present state and past states to depend only upon the present state; i.e., we need

$$P(X(t) = j|X(s) = i, X(u), 0 \leq u \leq s) = P(X(t) = j|X(s) = i)$$

for all s and t with $0 \leq s < t$, and for all i and j. Here, however, the Markov property does *not* hold: When both machines are down, the next transition depends on which of the two machines failed first.

(b) Formulate a CTMC describing the evolution of the system.

However, we can use 5 states with the states being: 0 for no copiers failed, 1 for copier 1 is failed (and copier 2 is working), 2 for copier 2 is failed (and copier 1 is working), $(1, 2)$ for both copiers down (failed) with copier 1 having failed first and being repaired, and $(2, 1)$ for both copiers down with copier 2 having failed first and being repaired. (Of course, these states could be relabelled 0, 1, 2, 3 and 4, but we do not do that.)

From the problem specification, it is natural to work with transition rates, where these transition rates are obtained directly from the originally-specified failure rates and repair rates (the rates of the exponential random variables). In Figure 1 we display a **rate diagram** showing the possible transitions with these 5 states together with the appropriate rates. It can be helpful to construct such rate diagrams as part of the modelling process.

From Figure 1, we see that there are 8 possible transitions. The 8 possible transitions should clearly have transition rates

$$Q_{0,1} = \gamma_1, Q_{0,2} = \gamma_2, Q_{1,0} = \beta_1, Q_{1,(1,2)} = \gamma_2, Q_{2,0} = \beta_2, Q_{2,(2,1)} = \gamma_1, Q_{(1,2),2} = \beta_1, Q_{(2,1),1} = \beta_2 \ .$$

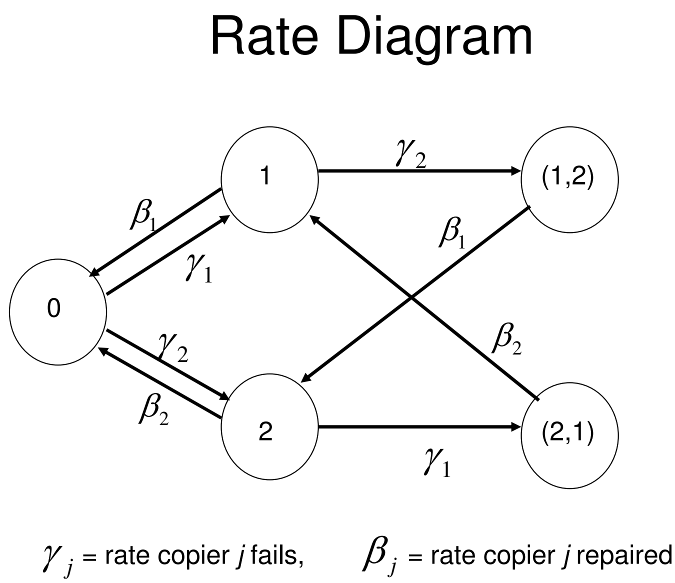

Figure 1: A rate diagram showing the transition rates among the 5 states in Problem 2, involving copier breakdown and repair.

We can thus define the transition-rate matrix Q. For this purpose, recall that the diagonal entries are minus the sum of the non-diagonal row elements, i.e.,

$$Q_{i,i} = -\sum_{j, j \neq i} Q_{i,j} \quad \text{for all} \quad i \ .$$

That can be explained by the fact that the rate matrix Q is defined to be the derivative (from above) of the transition matrix $P(t)$ at $t = 0$; see the CTMC notes.

In other words, the rate matrix should be

$$Q = \begin{array}{c} 0 \\ 1 \\ 2 \\ (1,2) \\ (2,1) \end{array} \left(\begin{array}{ccccc} -(\gamma_1 + \gamma_2) & \gamma_1 & \gamma_2 & 0 & 0 \\ \beta_1 & -(\gamma_2 + \beta_1) & 0 & \gamma_2 & 0 \\ \beta_2 & 0 & -(\gamma_1 + \beta_2) & 0 & \gamma_1 \\ 0 & 0 & \beta_1 & -\beta_1 & 0 \\ 0 & \beta_2 & 0 & 0 & -\beta_2 \end{array} \right) .$$

(c) Suppose that $\gamma_1 = 1$, $\beta_1 = 2$, $\gamma_2 = 3$ and $\beta_2 = 4$. Find the stationary distribution.

We first substitute the specified numbers for the rates γ_i and β_i in the rate matrix Q above, obtaining

$$Q = \begin{array}{c} 0 \\ 1 \\ 2 \\ (1,2) \\ (2,1) \end{array} \begin{pmatrix} -4 & 1 & 3 & 0 & 0 \\ 2 & -5 & 0 & 3 & 0 \\ 4 & 0 & -5 & 0 & 1 \\ 0 & 0 & 2 & -2 & 0 \\ 0 & 4 & 0 & 0 & -4 \end{pmatrix} .$$

Then we solve the system of linear equations $\alpha Q = 0$ with $\alpha \mathbf{e} = 1$, which is easy to do with a computer and is not too hard by hand. Just as with DTMC's, one of the equations in $\alpha Q = 0$ is redundant, so that with the extra added equation $\alpha \mathbf{e} = 1$, there is a unique solution. Performing the calculation, we see that the limiting probability vector is

$$\alpha \equiv (\alpha_0, \alpha_1, \alpha_2, \alpha_{(1,2)}, \alpha_{(2,1)}) = \left(\frac{44}{129}, \frac{16}{129}, \frac{36}{129}, \frac{24}{129}, \frac{9}{129} \right) .$$

Thus, the long-run proportion of time that copier 1 is working is $\alpha_0 + \alpha_2 = 80/129 \approx 0.62$, while the long-run proportion of time that copier 2 is working is $\alpha_0 + \alpha_1 = 60/129 \approx 0.47$. The long-run proportion of time that the repairman is busy is $\alpha_1 + \alpha_2 + \alpha_{(1,2)} + \alpha_{(2,1)} = 1 - \alpha_0 = 85/129 \approx 0.659$,

---

(d) Now suppose, instead, that machine 1 is much more important than machine 2, so that the repairman will always service machine 1 if it is down, regardless of the state of machine 2. Formulate a CTMC for this modified problem and find the stationary distribution.

With this alternative repair strategy, we can revise the state space. Now it does suffice to use 4 states, letting the state correspond to the set of failed copiers, because now we know what the repairman will do when both copiers are down; he will always work on copier 1. Thus it suffices to use the single state $(1,2)$ to indicate that both machines have failed. There now is only one possible transition from state $(1,2)$: $Q_{(1,2),2} = \mu_1$. We display the revised rate diagram in Figure 2 below.

The associated rate matrix is now

$$Q = \begin{array}{c} 0 \\ 1 \\ 2 \\ (1,2) \end{array} \begin{pmatrix} -(\gamma_1 + \gamma_2) & \gamma_1 & \gamma_2 & 0 \\ \beta_1 & -(\gamma_2 + \beta_1) & 0 & \gamma_2 \\ \beta_2 & 0 & -(\gamma_1 + \beta_2) & \gamma_1 \\ 0 & 0 & \beta_1 & -\beta_1 \end{pmatrix}$$

or, with the numbers assigned to the parameters,

$$Q = \begin{array}{c} 0 \\ 1 \\ 2 \\ (1,2) \end{array} \begin{pmatrix} -4 & 1 & 3 & 0 \\ 2 & -5 & 0 & 3 \\ 4 & 0 & -5 & 1 \\ 0 & 0 & 2 & -2 \end{pmatrix} .$$

Just as before, we obtain the limiting probabilities by solving $\alpha Q = 0$ with $\alpha \mathbf{e} = 1$. Now we obtain

$$\alpha \equiv (\alpha_0, \alpha_1, \alpha_2, \alpha_{(1,2)}) = \left(\frac{20}{57}, \frac{4}{57}, \frac{18}{57}, \frac{15}{57} \right) .$$

Thus, the long-run proportion of time that copier 1 is working is $\alpha_0 + \alpha_2 = 38/57 = 2/3 \approx 0.67$, while the long-run proportion of time that copier 2 is working is $\alpha_0 + \alpha_1 = 24/57 \approx 0.42$. The

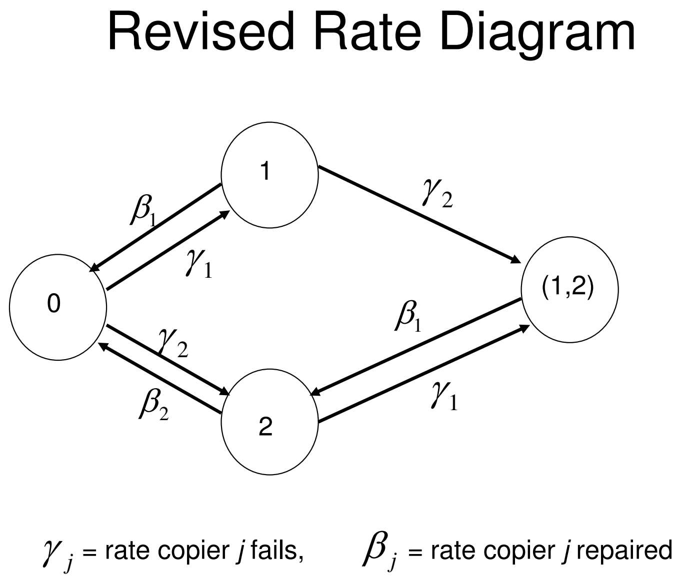

Figure 2: A revised rate diagram showing the transition rates among the 4 states in Problem 2, where the repairman always works on copier 1 first when both have failed.

new strategy has increased the long-run proportion of time copier 1 is working from 0.62 to 0.67, at the expense of decreasing the long-run proportion of time copier 2 is working from 0.47 to 0.42. The long-run proportion of time the repairman is busy is $1 - \alpha_0 = 37/57 \approx 0.649$, which is very slightly less than before.

We conclude by making some **further commentary**. We might think that the revised strategy is wasteful, because the repairman quits working on copier 2 when copier 1 fails after copier 2 previously failed. By shifting to work on copier 1, we might think that the repairman is being inefficient, "wasting" his expended effort working on copier 2, making it more likely that both copiers will remain failed. In practice, under other assumptions, that might indeed be true, but here because of the lack-of-memory property of the exponential distribution, the expended work on copier 2 has no influence on the remaining required repair times. From a pure efficiency perspective, it might be advantageous to give one of the two copiers priority at this point, but not because of the expended work on copier 2. On the other hand, we might prefer the original strategy from a "fairness" perspective. In any case, the CTMC model lets us analyze the consequences of alternative strategies. As always, the relevance of the conclusions depends on the validity of the model assumptions. But even when the model assumptions are not completely realistic or not strongly verified, the analysis can provide insight.