

IEOR 3106: Introduction to Operations Research: Stochastic Models

Fall 2012, Professor Whitt

Solutions to Homework Assignment 6: Due on Tuesday, October 23.

Chapter 6: Continuous-Time Markov Chains

Read Sections 6.1-6.5 in Ross. Do the following exercises at the end of Chapter 6.

As usual, you are only required to turn in the problems without answers in the back, but you should do all the problems.

2. The solution is in the back of the book, but we elaborate.

The object here is to construct a CTMC model. There are two key steps: (1) defining the state space and (2) specifying the local transition structure.

The state space in this example is the set of pairs (n, m) of nonnegative integers, where n is the number of organisms in state A and m is the number of organisms in state B . Thus, as indicated in the back of the book, the stochastic process that is the CTMC can be $\{(N_A(t), N_B(t)) : t \geq 0\}$, where $N_A(t)$ is the number of organisms in state A at time t , while $N_B(t)$ is the number of organisms in state B at time t .

The “local” transition structure of a CTMC can be defined in two ways. First, following Section 6.2 of Ross, we can specify the transition matrix P for the embedded discrete-time Markov chain (DTMC) operating at the transition epochs of the CTMC plus the rates of the exponential holding times in each state. As in Ross, we let ν_i denote the rate of the exponential holding time in state i ; i.e., the holding time in state i has an exponential distribution with mean $1/\nu_i$. The elements of the matrix P and the vector ν are given in the book.

An alternative way to specify a CTMC is to specify the (infinitesimal) transition rates. The (infinitesimal) transition rate of a transition from state i to state j is denoted by $Q_{i,j}$. The diagonal entries $Q_{i,i}$ of the matrix Q can be unspecified or can be defined as $Q_{i,i} \equiv -\sum_{j,j \neq i} Q_{i,j} = -\nu_i$. In terms of the model elements P and ν introduced above, we have

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

3. Recall that a birth-and-death (BD) process is a CTMC defined on a subset of the nonnegative integers that moves only to neighboring states; i.e., from state i it can only go next to one of $i - 1$ or $i + 1$. The transition rate from i to $i + 1$ is denoted by λ_i instead of $Q_{i,i+1}$, as it would be in the CTMC representation. Similarly, the transition rate from i to $i - 1$ is denoted by μ_i instead of by $Q_{i,i-1}$, as it would be in the CTMC representation.

This example cannot be represented directly as a BD process using the obvious state space, because the state space cannot simply be the number of machines working. We must keep track of which machine is working, if any. Moreover, the states cannot be relabeled so that this is a BD process. First we show that the system can be modeled as a CTMC.

We can analyze the system as a CTMC by letting the states be:

- b = both machines are working
- 1 = machine 1 is working, but 2 is not
- 2 = machine 2 is working, but 1 is not
- 0_1 = both machines are down, machine 1 is being serviced
- 0_2 = both machines are down, machine 2 is being serviced.

With these states we have a CTMC. We can define the CTMC by specifying the local transition structure in either of the two ways indicated for problem 2. The transition rates are

$$\begin{aligned}
 Q_{b,1} &= \mu_2 \\
 Q_{b,2} &= \mu_1 \\
 Q_{1,0_2} &= \mu_1 \\
 Q_{2,0_1} &= \mu_2 \\
 Q_{1,b} &= \mu \\
 Q_{2,b} &= \mu \\
 Q_{0_1,1} &= \mu \\
 Q_{0_2,2} &= \mu
 \end{aligned}$$

Alternatively, we can specify the model elements ν and P . For example,

$$\nu_b = Q_{b,1} + Q_{b,2} = \mu_2 + \mu_1 ,$$

while

$$P_{b,1} = \frac{Q_{b,1}}{Q_{b,1} + Q_{b,2}} = \frac{\mu_2}{\mu_2 + \mu_1} .$$

We might wonder if we could get a BD process from this CTMC by simply relabeling the states. For example, we may define new states by mapping the old states into the integers as shown below:

new state	old state	defining property
2	b	both machines are working
1	1	machine 1 is working, but 2 is not
3	2	machine 2 is working, but 1 is not
4	0_1	both machines are down, machine 1 is being serviced
0	0_2	both machines are down, machine 2 is being serviced.

From new state 2, we can go only to new states 1 and 3. From new state 3, we can go only to new states 2 and 4. From new state 1, we can go only to new states 0 and 2. Everything works so far, but there is a problem: From new state 0, we can only go to new state 3, and from new state 4, we can only go to new state 1. So it is not possible to get a BD process by relabeling the states.

We do not go on to “solve” this CTMC. For example, we could go on to solve for the steady-state distribution. If α is the steady-state probability vector, then $\alpha Q = 0$ and the sum

of the elements of the vector α is 1 (assuming that $Q_{i,i} = -\nu_i$). In other words, we solve the balance equations, setting the rate out of state i equal to the rate into state i :

$$\alpha_i \nu_i = \sum_{j:j \neq i} \alpha_j Q_{j,i} \quad \text{for all } i .$$

9. Since the death rate is constant, it follows that the number of deaths evolves as a Poisson process until the system becomes empty. Hence, $P_{i,j}(t)$ is the probability a Poisson distribution assumes the value $i - j$ for $i > j$; i.e.,

$$P_{i,j}(t) = \frac{e^{-\mu t} (\mu t)^{i-j}}{(i-j)!} \quad \text{for } 0 < j \leq i$$

and

$$P_{i,0}(t) = 1 - \sum_{j=1}^{j=i} P_{i,j}(t) = \sum_{k=i}^{\infty} \frac{e^{-\mu t} (\mu t)^k}{k!} .$$

13. This is an $M/M/1/1$ queueing model, having one server and one extra waiting space (a total capacity of 2). The two M 's refer to Markov, the first because the interarrival-time distribution is exponential (with the lack-of-memory property), and the second because the service-time distribution is also exponential; see Section 8.3.

Let $N(t)$ be the number of customers in the shop at time t . The process $\{N(t) : t \geq 0\}$ is a birth-and-death process with state space $\{0, 1, 2\}$.

The birth rates (per hour) are

$$\lambda_0 = 3 \quad \text{and} \quad \lambda_1 = 3 .$$

The death rates are

$$\mu_1 = 4 \quad \text{and} \quad \mu_2 = 4 .$$

The steady-state distribution can be found directly quite easily by solving local balance equations as indicated in Remark (iii) in Section 6.5. There is a standard form for any BD process (see (6.20), which we exploit below.

The steady-state vector α has the form

$$\alpha_i = \frac{r_i}{r_0 + r_1 + r_2} \quad \text{for } 0 \leq i \leq 2 ,$$

where

$$r_0 = 1, \quad r_1 = \frac{\lambda_0}{\mu_1} \quad \text{and} \quad r_2 = r_1 \times \frac{\lambda_1}{\mu_2} ,$$

i.e., $r_0 = 1$, $r_1 = 3/4$ and $r_2 = 9/16$, so that

$$\alpha_0 = \frac{16}{37}, \quad \alpha_1 = \frac{12}{37}, \quad \text{and} \quad \alpha_2 = \frac{9}{37} .$$

(a) Hence the average number in the shop is

$$E[N(\infty)] = (0 \times \alpha_0) + (1 \times \alpha_1) + (2 \times \alpha_2) = \frac{30}{37} .$$

By “average” number in the shop, we mean the expected value of the limiting steady-state distribution. Here we have $N(t)$ converge in distribution to $N(\infty)$ as $t \rightarrow \infty$.

(b) The proportion of customers that do not enter the shop is $\alpha_2 = 9/37$, so the proportion of customers that enter the shop is $28/37$.

(c) The service rate or death rate would increase from 4 to 8. The r vector would change to

$$r_0 = 1, \quad r_1 = 3/8 \quad \text{and} \quad r_2 = 9/64 .$$

Hence the steady-state vector changes to

$$\alpha_0 = \frac{64}{97}, \quad \alpha_1 = \frac{24}{97}, \quad \text{and} \quad \alpha_2 = \frac{9}{97} .$$

Thus the loss proportion decreases from $9/37$ to $9/97$. The served proportion thus increases from $28/37 = 0.7568$ to $88/97 = 0.9072$. The difference is 0.1505. Arrivals come at 3 per hour. Of these he serves 90.7% instead of 75.7%, a gain of 15.1%.

20. Let the state be the number of machines that are down. Then the state space is the set $\{0, 1, 2\}$. The stochastic process $\{X(t) : t \geq 0\}$, where $X(t)$ is the number of machines down at time t , is a birth-and-death process, because the process is a CTMC that can move only to a neighboring state at each transition.

The birth rates are $\lambda_i = \lambda$ for $i = 0, 1$; the death rates are $\mu_i = \mu$ for $i = 1, 2$. This model coincides with a $M/M/1/1$ queueing model, which has 1 server and 1 extra waiting space. Customer arrivals in the queue correspond to failures in the machines. Service times in the queue correspond to repair times for the machines.

(a) Referring to pages 356-359, we see that we want $E[T_0 + T_1]$, because T_0 is the time to go from state 0 to state 1, while T_1 is the time to go from state 1 to state 2. Then, following the top of page 357, we get

$$E[T_0 + T_1] = E[T_0] + E[T_1] = \left(\frac{1}{\lambda}\right) + \left(\frac{1}{\lambda} + \frac{\mu}{\lambda^2}\right) = \frac{2}{\lambda} + \frac{\mu}{\lambda^2} .$$

(b) To treat the variances, we look at the end of Section 6.3. In particular,

$$Var(T_0 + T_1) = Var(T_0) + Var(T_1) = \left(\frac{1}{\lambda^2}\right) + \left(\frac{1}{\lambda(\lambda + \mu)} + \frac{\mu}{\lambda^3} + \frac{\mu}{\mu + \lambda} \left(\frac{2}{\lambda} + \frac{\mu}{\lambda^2}\right)^2\right) .$$

The steady-state distribution has a simple form for a birth-and-death process. Let α_i be the steady-state probability of having i failed machines. Here, we want

$$\alpha_0 + \alpha_1 = 1 - \alpha_2 = \frac{1 + \lambda/\mu}{1 + \lambda/\mu + (\lambda/\mu)^2} .$$

23. Let the state be the number of machines that are down. The state space is thus $\{0, 1, 2, 3\}$. The number of machines that are down at time t is a BD process with birth rates $\lambda_0 = 3/10$, $\lambda_1 = 2/10$, $\lambda_2 = 1/10$ and death rates $\mu_1 = 1/8$, $\mu_2 = 2/8$, $\mu_3 = 2/8$.

The steady-state vector α has the form

$$\alpha_i = \frac{r_i}{r_0 + r_1 + r_2 + r_3} \quad \text{for} \quad 0 \leq i \leq 3 ,$$

where

$$r_0 = 1, \quad r_1 = \frac{\lambda_0}{\mu_1}, \quad r_2 = r_1 \times \frac{\lambda_1}{\mu_2}, \quad r_3 = r_2 \times \frac{\lambda_2}{\mu_3},$$

i.e., $r_0 = 1$, $r_1 = 24/10 = 12/5$, $r_2 = 48/25$ and $r_3 = 96/125$, so that

$$\alpha_0 = \frac{125}{761}, \quad \alpha_1 = \frac{300}{761}, \quad \alpha_2 = \frac{240}{761}, \quad \text{and} \quad \alpha_3 = \frac{96}{761}.$$

(a) The average number of machines not in use is

$$(0 \times \alpha_0) + (1 \times \alpha_1) + (2 \times \alpha_2) + (3 \times \alpha_3) = \frac{1068}{761}.$$

(b) The proportion of time both repairman are busy is

$$\alpha_3 + \alpha_2 = \frac{336}{761}.$$