CS 798: Homework Assignment 3 (Queueing Theory)

Assigned: October 6, 2009

1.0 Little's law

Patients arriving to the emergency room at the Grand River Hospital have a mean waiting time of three hours. It has been found that, averaged over the period of a day, that patients arrive at the rate of one every five minutes. (a) How many patients are awaiting treatment on average at any given point in time? (b) What should be the size of the waiting room so that it can accommodate everyone?

Solution: (a) The mean waiting time is 180 min, and the arrival rate is 0.2 patients/minute. Thus, the mean number of patients is their product = 180*0.2 = 36. (b) We do not have enough information to determine the maximum size of the waiting room! We know we need at least 36 spaces, but it's possible that a burst of a hundred patients may arrive, for example, due to an incident of mass food poisoning. But, as a rule of thumb, some small integer multiple of the mean, such as three or four times the mean, ought to be enough. In real life, we are forced to work with such 'fudge factors' because it is often too difficult or too expensive to determine the exact arrival process, which, in any case, may abruptly change over time.

2.0 A stochastic process

Consider that in Example 4, a person is on an infinite staircase on stair number 10 at time 0 and potentially moves once every clock tick. Suppose that he moves from stair i to stair $i+1$ with probability 0.2, and from stair i to stair $i-1$ with probability 0.2 (the probability of staying on stair i is 0.6). Compute the probability that the person is on each stair at time 1 (after the first move), time 2, and time 3.

Solution: At time 0, $P[X_0=10] = 1.0$.

At time 1, $P[X_1 = 9] = 0.2$; $P[X_1 = 10] = 0.6$; $P[X_1 = 11] = 0.2$.

At time 2, $P[X_2 = 8] = 0.2(0.2) = 0.04$; $P[X_2 = 9] = 0.2(0.6) + 0.6(0.2) = 0.24$; $P[X_2 = 10] = 0.2(0.2) + 0.6(0.6) + 0.2 (0.2) = 0.44$, and, by symmetry, $P[X_2 = 11] = 0.24$; $P[X_2 = 12] = 0.04$.

3.0 Discrete and continuous state and time processes

Come with your own examples for all four combinations of discrete state/discrete time/continuous state/continuous time processes.

4.0 Markov process

Is the process in Exercise 2 a Markov process? Why or why not?

Solution: The process is Markovian, because the probability of moving from stair i to stairs $i-1$, i, and $i+1$ do not depend on how the person reached stair i.

5.0 Homogeneity

Is the process in Exercise 2 homogeneous? Why or why not?

Solution: The transition probabilities are time-independent, and therefore the process is homogeneous.

6.0 Representation

(a) Represent the process in Exercise 2 using a transition matrix and a state transition diagram. (b) Do the rows in this matrix have to sum to 1? Do the columns in this matrix have to sum to 1? Why or why not? (c) Now, assume that the staircase has only 4 steps. Make appropriate assumptions (what are these?) to represent this finite process as a transition matrix and a state transition diagram.

Solution: (a)

$$\begin{bmatrix} \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0.2 & 0.6 & 0.2 & 0 & \dots & \dots \\ \dots & 0 & 0.2 & 0.6 & 0.2 & 0 & \dots \\ \dots & \dots & 0 & 0.2 & 0.6 & 0.2 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \end{bmatrix}$$

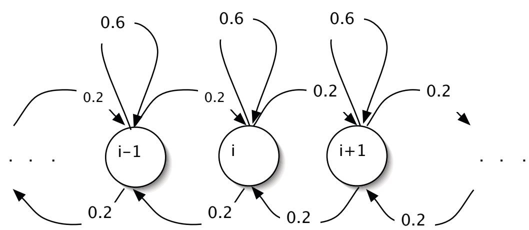

(b) The rows need to sum to 1, because at each time step, the process has to move to *some* state. The columns do not need to sum to 1 (think of a star-shaped state transition diagram with N states surrounding state 0, where state 0 has $1/N$ probability of going to any other state, and every state returns to state 0 with probability 1).

(c) We need to assume the boundary conditions. Suppose the at stair 1, the probability of staying at the same stair is 0.8, and at stair 4, the probability of staying at the same stair is also 0.8. Then, the transition matrix and state transition diagram are as shown below.

$$\begin{bmatrix} 0.8 & 0.2 & 0 & 0 \\ 0.2 & 0.6 & 0.2 & 0 \\ 0 & 0.2 & 0.6 & 0.2 \\ 0 & 0 & 0.2 & 0.8 \end{bmatrix}$$

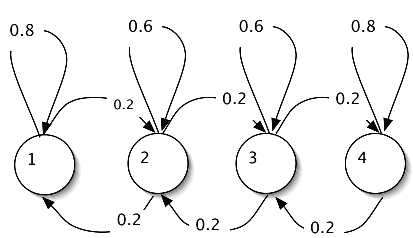

7.0 Reducibility

Is the chain in Exercise 2 reducible? Why or why not?

Solution: The chain is irreducible because every state can be reached from every other state.

8.0 Recurrence

Is state 1 in the chain in Exercise 6(c) recurrent? Compute f_1^1, f_1^2 and f_1^3.

Solution: State 1 is recurrent because the chain is finite and irreducible. f_1^1 is the probability that the process first returns to state 1 after one time step, and this is clearly 0.8. f_1^2 is the probability that the process first returns to state 1 after two time steps, and this is 0.2 * 0.2 = 0.04. f_1^3 is the probability that the process first returns to state 1 after three time steps. This can happen after a transition to state 2, a self loop in state 2, and then back. Thus, the value is 0.2*0.6*0.2 = 0.024.

9.0 Periodicity

Is the chain in Exercise 2 periodic? If not, give an example of a chain with period N for arbitrary $N > 1$.

Solution: The chain is not periodic because of the self-loop in every state. A trivial chain with period N is a ring with N states, with the transition probability of going from state i to state $(i+1) \bmod N = 1$.

10.0 Ergodicity

Is any state in the chain of Exercise 6(c) non-ergodic? Why or why not?

Solution: No state in the chain is non-ergodic because the chain is finite aperiodic and irreducible.

11.0 Stationary probability

Compute the stationary probability distribution of the chain in Exercise 6(c).

Solution:

From Theorem 2, because the chain is ergodic, we obtain:

$$\begin{aligned}
\pi_1 &= 0.8\pi_1 + 0.2\pi_2 \\
\pi_2 &= 0.2\pi_1 + 0.6\pi_2 + 0.2\pi_3 \\
\pi_3 &= 0.2\pi_2 + 0.6\pi_3 + 0.2\pi_4 \\
\pi_4 &= 0.2\pi_3 + 0.8\pi_4 \\
1 &= \pi_1 + \pi_2 + \pi_3 + \pi_4
\end{aligned}$$

This can be easily solved to obtain $\pi_1 = \pi_2 = \pi_3 = \pi_4 = 0.25$. (If you choose other assumptions for the boundary states, your computation will differ).

12.0 Residence times

Compute the residence time in each state of the Markov chain in Exercise 6(c).

Solution: $p_{11} = p_{44} = 0.8$, so the residence times in these states is 1/(1-0.8) = 1/0.2 = 5. $p_{22} = p_{33} = 0.6$, so the residence times in these states is 1/0.4 = 2.5.

13.0 Stationary probability of a birth-death-process

Consider the state-rate-transition diagram shown below.

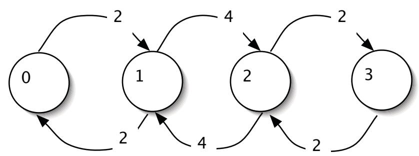

(a) Compare this with the state transition probability diagram in Exercise 6(c). What features are the same, and what differ?

(b) Write down the $\boldsymbol{Q}$ matrix for this system.

(c) Use the $\boldsymbol{Q}$ matrix to compute the stationary probability distribution of this chain.

Solution:

(a) Similarities: both are graphs with each node corresponding to a discrete state. Differences: the notation on an edge is the transition rate, not transition probability. The sum of rates leaving a node does not add up to 1, but total ingress rate matches total egress rate at each node.

(b) $$\begin{bmatrix} -2 & 2 & 0 & 0 \\ 2 & -6 & 4 & 0 \\ 0 & 4 & -6 & 2 \\ 0 & 0 & 2 & -2 \end{bmatrix}$$

(c) We have:

$-2P_0 + 2P_1 = 0$

$2P_0 - 6P_1 + 4P_2 = 0$

$4P_1 - 6P_2 + 2P_3 = 0$

$2P_2 - 2P_3 = 0$

$P_0 + P_1 + P_2 + P_3 = 1$

This yields: $P_0 = P_1 = P_2 = P_3 = 0.25$.

14.0 Poisson process

Prove that the inter-departure time of a pure-death process is exponentially distributed.

Solution: Consider a pure-death process, i.e. a birth-death process whose birth rates are zero. Clearly, the inter-departure times are nothing more than the residence times in each state. But we know that the residence times in a homogeneous continuous-time Markov chain are exponentially distributed (see Section A3.3.2 on page 16). QED.

15.0 Stationary probabilities of a birth-death process

Use Equation 30 to compute the stationary probability of the birth-death process in Exercise 13.

Solution: We see that in this chain, $\lambda_i = \mu_{i+1}$ so immediately we get $P_0 = P_1 = P_2 = P_3$. By summing them to 1, we can see that they are all 0.25.

16.0 M/M/1 queue

Is the birth-death-process in Exercise 13 M/M/1? Why or why not?

Solution: It is not M/M/1 because the state-transition rates are state-dependent.

17.0 M/M/1 queue

Consider a link to which packets arrive as a Poisson process at a rate of 450 packets/sec such that the time taken to service a packet is exponentially distributed. Suppose that the mean packet length is 250 bytes, and that the link capacity is 1 Mbps.

(a) What is the probability that the link's queue has 1, 2 and 10 packets respectively?

(b) What is the mean number of packets in the system? What is the mean number in the queue?

(c) What is the mean waiting time?

Solution:

(a) The packet length is 250 bytes = 2,000 bits, so that the link service rate of 1,000,000 bits/sec = 500 packets/sec. Therefore, the utilization is 450/500 = 0.9. When the link queue has 1 packet, it is in state j=2, because one packet is being served at that time. Thus, we need $P_2 = 0.9^2 * 0.1 = 0.081$. For the queue having two packets, we compute $P_3 = 0.9^3 * 0.1 = 0.0729$. For 10 packets in the queue, we compute $P_{11} = 0.9^{11} * 0.1 = 0.031$. (Compare these with values in Example 20 where the load is 0.8).

(b) The mean number of packets in the system is 0.9/1-0.9 = 9. Of these, 8 are expected to be in the queue.

(c) The mean waiting time is (1/500)/(1-0.9) = 0.002/0.1 = 0.02 s = 20 milliseconds.

18.0 Responsive (M/M/∞) server

Compute the ratio of P_j for a responsive server to the same value for an M/M/1 queue. How does this ratio behave as a function of j?

Solution: The ratio is:

$$\frac{e^{-\rho}\rho^j\frac{1}{j!}}{\rho^j(1-\rho)} = \frac{e^{-\rho}}{j!(1-\rho)} = \frac{1}{j!(1-\rho)e^{\rho}} = \frac{C}{j!}$$

, where C is a constant with respect to j. Therefore, for an M/M/∞ queue the probability of being in state j diminishes proportional to $j!$ compared to being in state j for an M/M/1 queue. Clearly, this favors much lower queue lengths for the M/M/∞ queue.

19.0 M/M/1/K server

Assume that the queueing system in Exercise 17 has 10 buffers. Compute an upper bound on the probability of packet loss.

Solution: Packet losses happen when there is an arrival and the system is in state j=11. This is upper bounded by P_{11}, which is given by

$$P_{11} = \frac{1-\rho}{1-\rho^{K+1}}\rho^j = \frac{0.1}{1-0.9^{12}}0.9^{11} = 0.0437 .$$

20.0 M/D/1 queue

Compute the mean number of customers in an M/D/1 system that has a utilization of 0.9. (a) How does this compare with a similarly loaded M/M/1 system? (b) Compute the ratio of the mean number of customers as a function of ρ. (c) Use this to compare the behavior of an M/D/1 queue with that of an M/M/1 queue under heavy load.

Solution: (a) The mean number of customers in the system for such a queue is given by

$\rho + \frac{\rho^2}{2(1-\rho)} = 0.9 + \frac{0.9^2}{2(0.1)} = 4.95$, which is roughly half the size of an equivalently loaded M/M/1 queue (from Exercise 17).

(b) The ratio is $\frac{\rho + \frac{\rho^2}{2(1-\rho)}}{\frac{\rho}{1-\rho}} = 1 - \frac{\rho}{2}$. This tends to 0.5 as the utilization tends to 1.

(c) Under heavy loads, the mean waiting time for an M/D/1 queue is half that of a similarly loaded M/M/1 queue.