

IEOR 6711: Stochastic Models I

SOLUTIONS to Second Midterm Exam, Chs. 3-4, November 20, 2007

Justify your answers; show your work.

1. Selling Flour: The Grainery. (30 points)

The Grainery is a store that sells flour to bakers. Suppose that bakers come to the Grainery according to a Poisson process with a rate of λ per week. Suppose that each baker asks to buy a random amount of flour, which is distributed according to a random variable X having cumulative distribution function (cdf) G with density g and a mean of m pounds. Let the successive quantities requested by the bakers be independent and identically distributed, independent of the arrival process.

The store uses an (s, S) inventory policy: Whenever the store's inventory level of flour drops below s pounds, the store immediately places an order from a centralized warehouse to bring its inventory level up to S pounds, where $S > s$. Assume that the delivery time from the warehouse to the Grainery is negligible; i.e., assume that these orders are delivered to the Grainery immediately after the last baker whose request takes the inventory level below s leaves. (The baker whose request takes the inventory level below s cannot receive any of the new order. That baker receives only part of his request if his request exceeds the current supply. This is the case of "lost sales;" i.e., there is no commitment to fill the rest of the order later. After that baker leaves, the new inventory is S .)

This is an elaboration of Example 3.4A on page 118. The key observation is that the counting process of the successive order times (by the Grainery from the warehouse) is a renewal counting process.

SCORING: (a) and (b) each 7, (c) and (e) each 6, and (d) 4.

(a) Give an expression for the long-run rate that the Grainery places orders to replace its stock of flour. Indicate how the numerical value can be obtained.

The rate is the reciprocal of the mean time between successive orders. To find it, let Y_n be the interarrival time between the $(n - 1)^{\text{st}}$ and n^{th} baker; and let X_n be the amount of flour requested by the n^{th} baker. Then $\{Y_n : n \geq 1\}$ is a sequence of IID exponential random variables with mean $1/\lambda$, while $\{X_n : n \geq 1\}$ is a sequence of IID random variables with cdf G and mean m . Let $N_G(t)$ be the renewal counting process associated with the sequence $\{X_n : n \geq 1\}$. Then the times between successive orders by the Grainery are IID, each distributed as

$$T = \sum_{k=1}^{N_G(S-s)+1} Y_k,$$

where the summands Y_k are IID and independent of the random variables $N_G(S - s) + 1$. Hence we can simply condition and uncondition (we do not need Wald) to get the mean

$$E[T] = E \left[\sum_{k=1}^{N_G(S-s)+1} Y_k \right] = E[(N_G(S - s) + 1)]E[Y_1] = \frac{m_G(S - s) + 1}{\lambda},$$

where m_G is the renewal function associated with N_G , i.e.,

$$m_G(t) \equiv E[N_G(t)], \quad t \geq 0.$$

Thus, the rate of orders by the Grainery is

$$\text{rate} = \frac{1}{E[T]} = \frac{\lambda}{m_G(S-s) + 1}.$$

(b) Let $X(t)$ be the inventory of flour at the Grainery at time t . Explain why the limit of $P(X(t) \geq x)$ as $t \rightarrow \infty$ exists and give an expression for that limit.

Let the cycles be the successive periods between orders by the Grainery. As noted above, these are IID. For $s \leq x \leq S$, the process $X(t)$ is first above x and then eventually below x . These successive periods above and below x produce an alternating renewal process. We can thus apply Theorem 3.4.4, which in turn can be proved by the key renewal theorem. We need to observe that the distribution of T is non-lattice, because it is a random sum of exponential random variables. Indeed, the sum of two independent random variables has a density if one of the two has a density; see Theorem 4 on p. 146 of Feller (1971), vol. II. See Asmussen (2003) for more on this issue. I was not expecting a detailed proof of this condition. Here we can obtain that representation by writing

$$T = Y_1 + \sum_{k=2}^{N_G(S-s)+1} Y_k.$$

Moreover, we need $E[T] < \infty$, which we have. Thus, we have

$$\lim_{t \rightarrow \infty} P(X(t) \geq x) = \frac{E[\text{reward per cycle}]}{E[\text{length of cycle}]} = \frac{(m_G(S-x) + 1)/\lambda}{(m_G(S-s) + 1)/\lambda} = \frac{m_G(S-x) + 1}{m_G(S-s) + 1},$$

as given in the last formula of the example in Ross on page 119.

(c) How does the answer in part (b) simplify when the cdf G is exponential?

When the cdf G is exponential, $m_G(t) = t/m$ for all $t \geq 0$. Hence, we obtain

$$\lim_{t \rightarrow \infty} P(X(t) \geq x) = \frac{E[\text{reward per cycle}]}{E[\text{length of cycle}]} = \frac{(m_G(S-x) + 1)/\lambda}{(m_G(S-s) + 1)/\lambda} = \frac{(S-x) + m}{(S-s) + m}.$$

(d) Is the cumulative distribution function of the limiting inventory level in part (c) continuous? Explain.

No. It is continuous except for an atom at S . Indeed, the distribution is uniform except for this atom. Note that

$$\lim_{t \rightarrow \infty} P(X(t) \geq S) = \frac{(S-S) + m}{(S-s) + m} = \frac{m}{(S-s) + m},$$

but clearly

$$\lim_{t \rightarrow \infty} P(X(t) \geq S) = \lim_{t \rightarrow \infty} P(X(t) = S).$$

(e) How does the answer (b) simplify when the mean m becomes relatively small compared to $S - s$? (To make this precise, suppose that the mean- m random variable $X \equiv X(m)$ introduced above is constructed from a mean-1 random variable Z by letting $X(m) = mZ$. Then let $m \downarrow 0$.)

Apply the elementary renewal theorem, stating that $m_G(t)/t \rightarrow 1/m$ as $t \rightarrow \infty$. When we scale down the mean, it is equivalent to increasing t in the renewal function. We want to look at

$$\lim_{m \downarrow 0} \lim_{t \rightarrow \infty} P(X(t) \geq x).$$

But, under our scaling relation, the limit is the uniform distribution on $[s, S]$. That should be consistent with intuition.

To proceed carefully, we can first apply part (b) to the mean-1 random variable Z . We use a subscript to indicate the mean of the X variable. When we start with Z , we have mean 1. We get

$$\lim_{t \rightarrow \infty} P(X_1(t) \geq x) = \frac{E[\text{reward per cycle}]}{E[\text{length of cycle}]} = \frac{(m_Z(S - x) + 1)/\lambda}{(m_Z(S - s) + 1)/\lambda} = \frac{m_Z(S - x) + 1}{m_Z(S - s) + 1},$$

Now, if we introduce the scaling by m , that is tantamount to re-scaling time by $1/m$ and keeping the random variable Z , which still have mean 1. Hence we get

$$\begin{aligned} \lim_{t \rightarrow \infty} P(X_m(t) \geq x) &= \lim_{t \rightarrow \infty} P(X_1(t/m) \geq x) \\ &= \frac{m_Z((S - x)/m) + 1}{m_Z((S - s)/m) + 1} \\ &\rightarrow \frac{S - x}{S - s}, \end{aligned} \tag{1}$$

because, by the elementary renewal theorem,

$$\frac{m_Z(t/m)}{t/m} \rightarrow 1 \quad \text{as } m \downarrow 0,$$

since $E[Z] = 1$.

2. Random Clockwise Walk Around the Circle. (35 points)

Consider a random walk on the circle, where each step is a clockwise random motion. At each step, the angle is at one of the values $k\pi/2$, $1 \leq k \leq 4$. That is, there is a sequence of independent clockwise motions on the circle among the four angles $k\pi/2$, $1 \leq k \leq 4$. Let transitions take place at positive integer times. In each step, the walk moves in a clockwise motion $j\pi/2$ with probability $(j + 1)/10$ for $0 \leq j \leq 3$. Henceforth, let state k represent $k\pi/2$,

then we have the following transition matrix for transitions among the four states 1, 2, 3 and 4:

$$P = \begin{matrix} & \begin{matrix} 1 & 2 & 3 & 4 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{pmatrix} 0.1 & 0.2 & 0.3 & 0.4 \\ 0.4 & 0.1 & 0.2 & 0.3 \\ 0.3 & 0.4 & 0.1 & 0.2 \\ 0.2 & 0.3 & 0.4 & 0.1 \end{pmatrix} \end{matrix} .$$

SCORING: Each part counted 7, so the total was 35 points. But part (b) was hard, so that it should be regarded as a "bonus" question.

(a) Show that there exists a distance (metric) d on the space of probability vectors of length 4 and a constant c with $0 < c < 1$ such that, for any two probability vectors $u \equiv (u_1, u_2, u_3, u_4)$ and $v \equiv (v_1, v_2, v_3, v_4)$,

$$d(uP, vP) \leq cd(u, v) .$$

This is an example of the contraction proof of the steady-state limit for irreducible aperiodic finite-state DTMC's, discussed in the lecture on November 8, and the posted notes for that day. See those notes. We use the metric associated with the l_1 norm:

$$d(u, v) \equiv \|u - v\|_1 \equiv \sum_{k=1}^m |u_k - v_k| .$$

As shown in the notes, the contraction constant c can be taken to be one minus the sum of the minimal elements of the columns. Hence here $c = 0.6$.

(b) Find the smallest such constant c in part (a), such that the inequality is valid for all u and v , and prove that it is smallest.

The contraction constant $c = 0.6$ specified above is *not* best possible. It is possible to find the answer by writing out the terms in detail of $d(uP, vP)$, as shown by Jaehyun Cho. The following is his method:

$$\begin{aligned} d(uP, vP) &= |0.0(u_1 - v_1) + 0.1(u_2 - v_2) + 0.2(u_3 - v_3) + 0.3(u_4 - v_4)| \\ &+ |0.3(u_1 - v_1) + 0.0(u_2 - v_2) + 0.1(u_3 - v_3) + 0.2(u_4 - v_4)| \\ &+ |0.2(u_1 - v_1) + 0.3(u_2 - v_2) + 0.0(u_3 - v_3) + 0.1(u_4 - v_4)| \\ &+ |0.1(u_1 - v_1) + 0.2(u_2 - v_2) + 0.3(u_3 - v_3) + 0.0(u_4 - v_4)| \\ &= | - 0.1(u_1 - v_1) + 0.0(u_2 - v_2) + 0.1(u_3 - v_3) + 0.2(u_4 - v_4)| \\ &+ |0.2(u_1 - v_1) - 0.1(u_2 - v_2) + 0.0(u_3 - v_3) + 0.1(u_4 - v_4)| \\ &+ |0.1(u_1 - v_1) + 0.2(u_2 - v_2) - 0.1(u_3 - v_3) + 0.0(u_4 - v_4)| \\ &+ |0.0(u_1 - v_1) + 0.1(u_2 - v_2) + 0.2(u_3 - v_3) - 0.1(u_4 - v_4)| \\ &\leq | - 0.1(u_1 - v_1)| + |0.0(u_2 - v_2)| + |0.1(u_3 - v_3)| + |0.2(u_4 - v_4)| \\ &+ |0.2(u_1 - v_1)| + | - 0.1(u_2 - v_2)| + |0.0(u_3 - v_3)| + |0.1(u_4 - v_4)| \\ &+ |0.1(u_1 - v_1)| + |0.2(u_2 - v_2)| + | - 0.1(u_3 - v_3)| + |0.0(u_4 - v_4)| \end{aligned}$$

$$\begin{aligned}
& + |0.0(u_1 - v_1)| + |0.1(u_2 - v_2)| + |0.2(u_3 - v_3)| + |-0.1(u_4 - v_4)| \\
= & 0.4(|u_1 - v_1| + |u_2 - v_2| + |u_3 - v_3| + |u_4 - v_4|)
\end{aligned}$$

The second equality is obtained by subtracting $0.1(u_1 - v_1) + 0.1(u_2 - v_2) + 0.1(u_3 - v_3) + 0.1(u_4 - v_4)$ from each term, which changes nothing because it is 0, since u and v are probability vectors.

It is easy to show that $c = 0.4$ is best possible by an example: $u = (1, 0, 0, 0)$ and $v = (0, 0, 1, 0)$. Clearly, $d(u, v) = 2$ and $d(uP, vP) = 0.8 = 0.4d(u, v)$.

(c) Use part (a) to show that there exists a unique stationary probability vector for P , i.e., a probability vector $\pi \equiv (\pi_1, \pi_2, \pi_3, \pi_4)$ satisfying

$$\pi = \pi P.$$

See the notes.

(d) Use part (a) to show that

$$d(uP^n, \pi) \leq Kc^n, \quad n \geq 1,$$

where c and K are constants, independent of u , and identify the best possible constant K (the smallest constant that is valid for all u).

See the notes. The notes give $d(u/\pi)$ on the right instead of K . Note that $d(u, v) \leq 2$ for all probability vectors u and v . However, for $\pi = (1/4, 1/4, 1/4, 1/4)$, which is the limit here (see part (e) below), we can do better. We can have K at most $2 \times (3/4) = 3/2$. That is attained by $u = (1, 0, 0, 0)$.

(e) Find the stationary probability vector π .

Since the transition matrix P is doubly stochastic, we know right away that the stationary probability vector is $\pi = (1/4, 1/4, 1/4, 1/4)$.

3. Customers at an ATM. (50 points)

Suppose that customers arrive at a single automatic teller machine (ATM) according to a Poisson process with rate λ per minute. Customers use the ATM one at a time. There is unlimited waiting space. Assume that all potential customers join the queue and wait their turn. (There is no customer abandonment.)

Let the successive service times at the ATM be mutually independent, and independent of the arrival process, with a cumulative distribution function G having density g and mean $1/\mu$. Let $Q(t)$ be the number of customers at the ATM at time t , including the one in service, if any.

This is the $M/G/1$ queue, as discussed in class on November 15. See the posted notes and the cited sections in Ross.

SCORING: All parts counted 5 points, except part (d), which counted 10 points. The total possible score was 50. Parts (h) and (i) can be considered bonus questions.

(a) Prove that $Q(t) \rightarrow \infty$ as $t \rightarrow \infty$ with probability 1 if $\rho \equiv \lambda/\mu > 1$.

Let $A(t)$ count the number of arrivals in $[0, t]$, let $D(t)$ count the number of departures in $[0, t]$; let $S(t)$ count the number of service completions in $[0, t]$ if the server were to work continuously without interruption, using the given sequence of service times. Clearly, $D(t) \leq S(t)$ w.p.1. Note that $Q(t) = A(t) - D(t)$ for each t . By the SLLN for a Poisson process (a special case of Proposition 3.3.1), $A(t)/t \rightarrow \lambda$. By the SLLN for the renewal process, $S(t)/t \rightarrow \mu$. Thus,

$$\liminf_{t \rightarrow \infty} \frac{A(t) - D(t)}{t} \geq \lim_{t \rightarrow \infty} \frac{A(t) - S(t)}{t} = \lambda - \mu > 0,$$

which implies that

$$\liminf_{t \rightarrow \infty} \frac{Q(t)}{t} \geq \lambda - \mu > 0,$$

so that $Q(t) \rightarrow \infty$ w.p.1.

(b) Is the stochastic process $\{Q(t) : t \geq 0\}$ a Markov process? Explain.

No. The future of $Q(t)$ depends on the elapsed service time in progress at time t .

(c) Identify random times T_n , $n \geq 1$, such that the stochastic process $\{X_n : n \geq 1\}$ is an irreducible infinite-state discrete-time Markov chain (DTMC), when $X_n = Q(T_n)$ for $n \geq 1$.

The random times can be the departure times. See the notes and the book.

(d) Find conditions under which the state X_n at time n of the DTMC $\{X_n : n \geq 1\}$ in part (c) converges in distribution to a proper steady-state limit as $n \rightarrow \infty$, and determine that limiting steady-state distribution.

We follow Example 4.3 A on page 177. We need to have $\lambda < \mu$ or, equivalently, $\rho \equiv \lambda/\mu < 1$. We apply Theorem 4.3.3 stating that it suffices to solve $\pi = \pi P$, under the conditions. We determine the steady-state probability vector π via its generating function. The generating function of π is what was wanted.

(e) In the setting of part (d), what is the steady-state probability that the system is empty (at these embedded random times)?

We get $\pi_0 = 1 - \rho$ as part of the analysis in the previous part.

(f) How does the steady-state distribution determined in part (d) simplify when the service-time cdf is

$$G(x) = 1 - e^{-\mu x}, \quad x \geq 0 ?$$

The distribution becomes geometric; see the notes.

(g) Find

$$\lim_{n \rightarrow \infty} P(X_n = j)$$

in the setting of part (f).

$$\lim_{n \rightarrow \infty} P(X_n = j) = (1 - \rho)\rho^j.$$

(h) What is the heavy-traffic approximation for the steady-state distribution found in part (d)?

A heavy-traffic approximation is the exponential distribution with the exact mean. The exact mean is

$$E[X_{\rho, \infty}] = \rho + \left(\frac{\rho^2}{1 - \rho} \right) \left(\frac{c_s^2 + 1}{2} \right),$$

where c_s^2 is the SCV of the service time. The limit in the next part does not directly yield exactly that approximation, because as $\rho \uparrow 1$, some of the ρ terms get replaced by 1. A direct application of the limit would yield first

$$(1 - \rho)X_{\rho, \infty} \approx \left(\frac{c_s^2 + 1}{2} \right) Z,$$

where Z is a mean-1 exponential random variable, and then

$$X_{\rho, \infty} \approx \left(\frac{c_s^2 + 1}{2(1 - \rho)} \right) Z,$$

but, given the exact mean, we could refine the approximation. The main thing we learn for $M/G/1$ is that the complicated distribution tends to be approximately exponential for large ρ .

(i) State and prove a limit theorem justifying the approximation in part (h).

See the notes. Paralleling the approximation above, the limit is

$$(1 - \rho)X_{\rho, \infty} \Rightarrow L,$$

where L is a random variable with an exponential distribution having mean $(c_s^2 + 1)/2$.

The detailed proof is given in the posted pages from Gnedenko and Kovalenko. The idea is to exploit Taylor's series, which becomes valid because of the scaling.
