IEOR 6711, Stochastic Models, I: Final Exam

Fall 2013, SOLUTIONS

There are three questions, each with several parts.

1. Customers Coming to a Group of Automatic Teller Machines (35 points: 3 points each for first 5 parts; 4 points each for last 5 parts)

Customers arrive one at a time to a **group of 4 ATM's** (automatic teller machines) to withdraw money. Customers **arrive** according to a Poisson process with rate $\lambda \equiv 1$ per minute. The **service times** of successive customers at the ATM can be regarded as i.i.d. random variables, each distributed as the random variable S, with mean $E[S] = 2$. However, the customers are **highly impatient** and are unwilling to wait if the 4 ATM's are all busy, so that they leave immediately if all ATM's are busy.

First Part: Exponential Service Times

For this first part assume that the service times are exponentially distributed.

(a) Identify an appropriate Markov process that can be used to analyze the steady-state and transient behavior of this system.

The number of busy ATM's at time t, denoted by $X(t)$, is a birth-and-death process if the service times are exponentially distributed. The model is the Erlang loss model $M/M/4/0$. The birth rates are $\lambda_k = \lambda = 1$, $0 \le k \le 3$, while the death rates are $\mu_k = k\mu = k/2, 1 \le k \le 4$.

(b) What is the long-run proportion of time that all 4 ATM's are simultaneously busy?

The probability that k ATM's are busy at an arbitrary time in steady state is α_k, where

$$\alpha_k = \frac{r_k}{\sum_{j=0}^{4} r_j} = \frac{Ga^k}{k!}, \quad 0 \le k \le 4, \tag{1}$$

where $a \equiv \lambda/\mu = 1 \times 2 = 2$, because $\lambda = 1$, $1/\mu = E[S] = 2$, and G is the normalization constant, i.e.,

$$G \equiv G(4) \equiv G(4,a) = \sum_{j=0}^{4} \frac{a^j}{j!}.$$

By the PASTA (Poisson Arrivals See Time Averages) property, $\alpha_4 = B(s,a) = B(4,2)$, the Erlang blocking probability for an arrival in steady state.

(c) True or False: The stationary departure process of served customers is a Poisson process. Justify your answer.

FALSE. All birth-and-death processes are reversible, and so this process $X(t)$ is reversible, but unlike the case of an unlimited waiting room, the stationary departure process is NOT a Poisson process. Notice that a proportion of arrivals are blocked (lost), while only the served customers depart. The overflow traffic tends to be more highly variable than a Poisson process, with overflows occurring in bunches, only when all servers are busy. In contrast the departure process of served customers, the so-called "carried traffic," tends to be smoother than Poisson.

Providing a proof is somewhat complicated, however. It is easily verified by simulation. Here is a rough line of reasoning: If the stationary departure process of served customers were Poisson, then it would have stationary independent increments. However, conditional on having NO departures over a long interval $[0,T]$ in steady state, we can show that the conditional expected number of busy servers is higher than the steady-state mean, so that the conditional departure rate after time T is higher than its long-run average, which violates the stationary and independent increments property.

Second Part: Uniform Service Times

For this second part, assume that the service times are uniformly distributed on the interval $[0,4]$.

(d) What is the long-run proportion of time that all 4 ATM's are simultaneously busy? ATM's?

This is the $M/GI/4/0$ Erlang loss model, which has the insensitivity property, as discussed in §9 of the CTMC notes and §5.7.2 of the textbook. The steady-state number of busy servers has the same distribution as in the associated $M/M/4/0$ birth-and-death process, given in part (b) above.

(e) Identify an appropriate Markov process that can be used to justify your answer in part (d) and study the transient behavior of this system?

The main thing is to define an appropriate state of the Markov process. Let the state be denoted by $(n;x)$ where n is the number of busy ATM"s and $x \equiv (x_1, x_2, \ldots, x_n)$, depending on n, is a vector of length n giving the times (ages) that each of the n ATM's have been occupied by their current customer, arranged in ascending order, so that $x_1 < \cdots < x_n$ for each n. When $n = 0$, the state is simply 0. The stochastic process can be represents as $\{(X(t), Y(t)) : t \geq 0\}$, where $X(t) = n$ and $Y(t) = (Y_1, \ldots, Y_n(t) = (x_1, \ldots, x_n)$ if the state is $(n;x)$ at time t.

(f) Exhibit the steady-state distribution of the Markov process in part (e). (Be as explicit as possible.)

$$lim_{t\to\infty} P(X(t) = n, Y_j(t) \leq x_j, 1 \leq j \leq n) = \alpha_n \prod_{j=1}^{n} n! F_e(x_j),$$

where α_n is given in part (b) and F_e is the cdf of the stationary-excess distribution associated with the service-time cdf F. Since the pdf f is uniform on $[0,4]$, the cdf of the service time is

$$F(t) = \frac{t}{4}, \quad 0 \leq t \leq 4, \quad \text{and} \quad F(t) = 1, \quad t \geq 4,$$

and the pdf of the stationary-excess variable is

$$f_e(t) \equiv \frac{1 - F(t)}{E[S]} = \frac{1}{2} - \frac{t}{8}, \quad 0 \le t \le 4,$$

so that the cdf is

$$F_e(t) \equiv \frac{t}{2} - \frac{t^2}{16}, \quad 0 \le t \le 4, \quad \text{and} \quad F_e(t) = 1, \quad t \ge 4.$$

(g) Prove that the steady-state distribution in part (f) is valid, stating all theorems used, possibly including theorems not actually proved in the book or in class.

The proof desired is in §9.3 of the CTMC notes or §5.7.2 of the book. The theorem to apply and state is the continuous-state analog of Kelly's lemma, Theorem 6.8, which was not actually proved in class or in the book. Let $(n; x)$ be a possible state and $\alpha(n; x)$ be its steady-state probability (a density with respect to x). Let $q((n; x), (n'; x'))$ represent the transition intensity from state $(n; x)$ to another state $(n'; x')$ and let $\overleftarrow{q}((n; x), (n'; x'))$ represent the transition intensity from state $(n; x)$ to another state $(n'; x')$ for a candidate reverse-time process. Here is the theorem that we want to apply

Theorem 0.1 *Consider a Markov process with states $(n; x)$ as specified above and evolution according to the infinitesimal rate (generator) $q((n; x), (n'; x'))$. If we can find candidate steady-state distributions specified by $\alpha(n; x)$ and a reverse-time Markov process specified by $\overleftarrow{q}((n; x), (n'; x'))$ such that*

$$\alpha(n; x)q((n; x), (n'; x')) = \alpha(n'; x')\overleftarrow{q}((n'; x'), (n; x)) \quad \textit{for all} \quad \textit{states}(n; x) \quad \textit{and} \quad (n'; x'),$$

then $\alpha(n; x)$ characterizes the unique steady-state distribution of the Markov process and $\overleftarrow{q}((n; x), (n'; x'))$ characterizes a bonafide reverse-time process.

Given this theorem, the desired proof is the proof of Theorem 9.3 in the notes. We guess that the reverse time process is also a loss model with state $(n; x)$, but in reverse time x_j represents the remaining service time instead of the age. We see that there are two transitions to consider, due to an arrival and a service completion. We see that the balance equation above is satisfied in each case. (See the notes and the book.)

(h) Let T be the time after an arrival finds the system full in steady state until an ATM first becomes free. What is $P(T > 2 \ \text{minutes})$? (explicit numerical answer desired, recall that $E[S] = 2$)

Let $H \equiv P(T \le x)$. Then T is distributed as the minimum of 4 i.i.d. random variables, each distributed with cdf F_e in part (c). Hence,

$$P(T > x) \equiv H^c(x) = F_e^c(x)^4 = \left(1 - \frac{x}{2} + \frac{x^2}{16}\right)^4$$

so that

$$P(T > 2) = \left(1 - \frac{2}{2} + \frac{2^2}{16}\right)^4 = \left(\frac{1}{4}\right)^4 = \frac{1}{256}$$

Note that $P(T > 2$ is much smaller than $P(S > 2) = 1/2$.

(i) Show that the stochastic process representing the number of busy ATM's at time t is a regenerative process by identifying the regeneration times.

Observe that the instants that arrivals come to find an empty system, with all 4 ATM's idle, are the epochs of a renewal process. Moreover, the future evolution has the same distribution from each of those epochs, because the Poisson arrival process has stationary and independent increments, and there are no service times in process. The next part shows that these times between renewals have finite mean. The distribution can be shown to be nonlattice as well. This provides another way to prove that the stochastic process $\{X(t) : t \geq 0\}$ and the stochastic process $\{(X(t), Y(t)) : t \geq 0\}$ have proper limiting distributions. Both are regenerative stochastic processes.

(j) What is the expected time between successive regeneration times in part (i)?

Observe that the steady-state probability that all ATM's idle has been determined as $\alpha_0 = 1/G(4)$ in parts (b) and (d) above. On the other hand, by the renewal reward theorem, that steady-state probability must coincide with the long-run proportion of time that all 4 ATM's are idle, which also can be represented as the expected reward per cycle divided by the expected length of a cycle, where a cycle X is the time between renewals (whose mean EX we want to determine). To get that, we can assign reward at rate 1 whenever all 4 ATM's are idle. Thus the expected reward per cycle is the expected time after all 4 ATM's become idle until the next arrival. Since the arrival process is Poisson, the expected reward is the mean time between arrivals. Thus we have

$$\alpha_0 = \frac{1/\lambda}{E[X]},$$

so that

$$E[X] = \frac{1}{\lambda \alpha_0} = \frac{1}{\alpha_0},$$

where α_k is given in parts (b) and (d) above, because $\lambda = 1$.

---

2. The Movement of a Taxi
(35 points: 4 points each, except 3 for (a))

A continuously operating taxi serves three locations: A, B and C.

idle times:
The taxi sits idle at each location an exponential length of time before departing to make a trip to one of the other two locations. The mean idle times are 2 minutes at A, 1 minute at B and 2 minutes at C. The idle times and travel times are mutually independent.

transition probabilities:
From A, the taxi next goes to B with probability $1/3$ and to C with probability $2/3$.

From B, the taxi next goes to A with probability $1/2$ and to C with probability $1/2$.
From C, the taxi next goes to B with probability $1/3$ and to A with probability $2/3$.

travel times:

The travel times between A and B in either direction are uniformly distributed in the interval $[5, 15]$ minutes.
The travel times between A and C in either direction are uniformly distributed in the interval $[20, 60]$ minutes.
The travel times between B and C in either direction are uniformly distributed in the interval $[20, 40]$ minutes.

Please provide explicit numerical answers to all questions below, showing your method, unless indicated otherwise.

(a) What is the long-run proportion of all taxi trips starting from location A?

First, we observe that this problem is a variation on Exercise 4.50 in Ross, which relates to §4.8 on semi-Markov processes.

This first part only involves a finite-state DTMC. The transition matrix among locations, based on the transition probabilities given directly above, is

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

By symmetry, it is evident that the stationary vector must satisfy $\pi_A = \pi_C$. We thus easily solve $\pi = \pi P$ to get $\pi = (3/8, 2/8, 3/8)$. The long-run proportion of all trips starting from A is $\pi_A = 3/8$.

(b) What is the long-run proportion of time that the taxi's most recent stop was at location A?

We can solve this problem in (at least) two ways. First, we can apply the standard formula for the steady-state distribution of an SMP. Thus, in the context of the three states defined in part (a), we can compute the mean time spent in each state. We add the mean idle time to the mean travel time. We get the mean vector

$$m \equiv (m_A, m_B, m_C) = (32, 21, 116/3).$$

We then apply the steady-state formula for a semi-Markov process (SMP) in Theorem 4.8.3 of Ross to get the limiting probability. Let α_A be the long-run proportion of time that the last stop was in location A. (The taxi may be idle at A or traveling away from A, toward either B or C.) We then have by Theorem 4.8.3 of Ross that

$$\alpha_A = \frac{\pi_A m_A}{\pi_A m_A + \pi_B m_B + \pi_C m_C} = \frac{(96/8)}{(254/8)} = \frac{96}{254} = \frac{48}{127} \approx 0.378$$

The second way is to first increase the number of states to distinguish between being idle and traveling. That is done in part (c) below. After doing that, the answer becomes the long-run proportion of time the taxi is in one of the three states A, (A,B) and (A,C), where A means idle at A and (A,B) means traveling from A to B. From the details given below in part (c), we see that this long run proportion is

$$\frac{(96/16)}{(254/16)} = \frac{96}{254} = \frac{48}{127} \approx 0.378$$

Fortunately, the answers agree.

(c) What is the long-run proportion of time that the taxi is idle at location A?

Again, we can solve this in several ways. For this question, it is convenient to add six more states: (A,B), (A,C), (B,A), (B,C), (C,A) and (C,B), with (A,B) meaning that the taxi is traveling from A to B. We then have the 9-state transition matrix

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

From one of the three locations, we go to traveling; from traveling, we go next to one of the three locations. From the answer to part (a), it is easy to determine the stationary probability vector for this 9×9 chain. Since we are at the locations on half the transitions, we divide the previous probabilities by 2. We then get $\pi_{(A,B)} = \pi_A P_{A,B}$ using the 3×3 transition matrix on the right, and so forth. Hence, it is easy to see that the stationary probability vector for this 9×9 discrete-time Markov chain must be

$$\pi = (3/16, 2/16, 3/16, 1/16, 2/16, 1/16, 1/16, 2/16, 1/16).$$

This is also easily verified by direct calculation: We see that indeed $\pi = \pi P$ for this 9-state DTMC. (This representation plays a role in part (i) below, so I am anticipating that in my approach, but nevertheless this approach is natural.)

We are now ready to answer the question. Let α_A be the long-run proportion of time spent by the taxi idle at location A. We then have by Theorem 4.8.3 of Ross in this more detailed context that

$$\alpha_A = \frac{\pi_A m_A}{\sum_i \pi_i m_i},$$

where m_i is the mean time spent in state i. The mean vector is

$$m = (m_A, m_B, m_C, m_{(A,B)}, m_{(A,C)}, m_{(B,A)}, m_{(B,CB)}, m_{(C,A)}, m_{(C,B)}) = (2, 1, 2, 10, 40, 10, 30, 40, 30).$$

Hence,

$$\alpha_A = \frac{(3/16)2}{\sum_i \pi_i m_i} = \frac{(6/16)}{254/16} = \frac{6}{254} = \frac{3}{127} \approx 0.0236$$

On the other hand, we can exploit renewal theory directly. We can start by identifying an embedded renewal process. Let the times of successive arrivals to A constitute renewals. We can write

$$\alpha_A = \frac{m_A}{m_{A,A}} \quad \text{so that} \quad m_{A,A} = \frac{m_A}{\alpha_A} = \frac{32}{96/254} = \frac{254}{3},$$

applying part (b). We then write

$$\frac{E[\text{reward per cycle}]}{E[\text{length of cycle}]} = \frac{2}{254/3} = \frac{6}{254} = \frac{3}{127} \approx 0.0236$$

(d) Let $P_t(A)$ be the probability that the taxi is idle at location A at time t. Does $P_t(A)$ converge to a proper limit as $t \to \infty$? Why or why not? If so, what is that limit?

The answer is YES by Proposition 4.8.1, which in turn is implied by the limit theorem for alternating renewal processes, Theorem 3.4.4, which in turn is implied by the key renewal theorem. We first observe that the SMP is irreducible. We next observe that the return times to each state have non-lattice distributions. Since all the time random variables have densities, so there is no problem with non-lattice. The limit is the same as the answer in part (c). The alternating renewal process step is useful, because the function $h(t)$ in the renewal equation that must be d.R.i. (directly Riemann integrable) is nondecreasing and bounded, which is a convenient condition for d.R.i. Specifically, the renewal equation is

$$g(t) = h(t) + \int_0^t g(t-s)\, dF(s),$$

and its solution is

$$g(t) = h(t) + \int_0^t h(t-s)\, dm(s),$$

where here $g(t) = P_t(A)$ and $h(t) = P(U > t)$, where U is the length of the "up" or "on" interval in the alternating renewal process with cycle cdf F and $m(t)$ is the renewal function associated with F.

(e) What is the rate (per unit of time) at which the taxi makes trips departing from location A heading toward location B? Start by defining what is meant by "rate" here.

Let $N_{A,B}(t)$ be the counting process recording the number of taxi trips starting from A headed toward B in the time interval $[0,t]$. By "rate," we mean the limit of $N_{A,B}(t)/t$ as $t \to \infty$. To be precise, we count at the instant the taxi leaves A after being idle, but the exact time will not matter in the limit, because we have divided by t. We see that this is a renewal counting process.

The long-run rate that the taxi makes trips from A is $1/m_{A,A}$, where $m_{A,A}$ is the mean time between arrivals to A. The instants between arrivals to A can serve as an embedded renewal process. A proportion $P_{A,B}$ of these visits are followed by a trip to B. So the long-run rate of trips from A to B is

$$P_{A,B}/m_{A,A}.$$

On the other hand, we know that $\alpha_A = m_A/m_{A,A}$ by Proposition 4.8.1. Hence,

$$\frac{P_{A,B}}{m_{A,A}} = \frac{\alpha_A P_{A,B}}{m_A} = \frac{(3/127)(1/3)}{2} = \frac{1}{254}.$$

(f) What is the long-run conditional probability that the taxi will come next to location B, given that the taxi is now traveling away from location A?

$$\lim_{t\to\infty} P(X(t) = (A,B)|X(t) \in \{(A,B) \cup (A,C)\}) = \frac{\lim_{t\to\infty} P(X(t) = (A,B))}{\lim_{t\to\infty} P(X(t) \in \{(A,B) \cup (A,C)\})}$$
$$= \frac{\pi_{(A,B)} m_{(A,B)}}{\pi_{(A,B)} m_{(A,B)} + \pi_{(A,C)} m_{(A,C)}} = \frac{(1/16)10}{(1/16)10 + (2/16)40} = \frac{1}{9}.$$

Note that this conditional probability does not simply equal $P_{A,B} = 1/3$. We need to take time into account.

(g) What is the long-run proportion of time that the taxi is traveling from A to C and the remaining time before getting to C is at least 30 minutes?

This is a variant of Theorem 4.8.4 in Ross. Let $m_{(A,C),(A,C)}$ be the mean time between beginning a trip from A to C. Let $T_{A,C}$ be the uniformly distributed travel time on a trip from A to C. let $m_{A,C} \equiv E[T_{A,C}]$. We already have two expressions for $\alpha_{(A,C)}$, one being $m_{A,C}/m_{(A,C),(A,C)}$ and the other from the reasoning in part (c), involving Theorem 4.8.3 of Ross.

$$\begin{aligned}
\lim_{t\to\infty} P(X(t) = (A,C), Y(t) > 30) &= \frac{E[(T_{A,C} - 30)^+]}{m_{(A,C),(A,C)}} \\
&= \frac{\int_0^\infty P((T_{A,C} - 30)^+ > y)\, dy}{m_{(A,C),(A,C)}} \\
&= \frac{\int_0^{30} P(U(20,60) - 30 > y\, dy}{m_{(A,C),(A,C)}} \\
&= \frac{\int_0^{30} (30 - y)/40\, dy}{m_{(A,C),(A,C)}} \\
&= \frac{(900)/80}{m_{(A,C),(A,C)}} = \frac{(900)/80)\alpha_{(A,C)}}{m_{(A,C)}} \\
&= \left(\frac{900}{80}\right)\left(\frac{(80/254)}{40}\right) = \left(\frac{45}{4}\right)\left(\frac{1}{127}\right) = \frac{45}{508} \approx 0.089
\end{aligned}$$

(h) Which of the previous answers would change if the travel times were changed from uniform to exponential with the same mean? (You need not do any new computations?)

Only the previous part, part (g), would have a different answer. All the others had formulas that depend only on the mean travel time.

(i) Suppose that the travel times are indeed changed from uniform to exponential with the same mean. Let $X(t)$ be the state of the taxi at time t, e.g., idle at A or traveling from A to B. Give an explicit formula (not numerical value) for the conditional probability

$$P(X(2) = \text{idle at } B \quad \text{and} \quad X(7) = \text{idle at } C | X(0) = \text{idle at } A).$$

Under the new exponential assumption, the SMP using the 9 states becomes a CTMC. With the state notation in the solution to part (c), we have

$$\begin{aligned} P(X(2) = \text{idle at } B \quad \text{and} \quad X(7) = \text{idle at } C | X(0) = \text{idle at } A) \\ \equiv P(X(2) = B, X(7) = C | X(0) = A) = P_{A,B}(2) P_{B,C}(5), \end{aligned}$$

where $P_{i,j}(t) \equiv P(X(t+s) = j | X(s) = i)$ is the transition probability for the CTMC. (Note that the A and B that appear as subscripts in $P_{A,B}(2)$ above are separate states indicating "idle at A" and "idle at B," respectively. If we had meant a transition from "traveling from A to B" to "traveling from B to C," in time $t = 2$, then we would have written $P_{(A,B),(B,C)}(2)$.)

The transition probability $P_{i,j}(t)$ in turn is an element of the transition matrix $P(t)$, which can be obtained by solving one of the ODE's $\dot{P}(t) = QP(t)$ or $\dot{P}(t) = P(t)Q$. More directly, the solution of these ODE's can be represented explicitly as the matrix exponential

$$P(t) = e^{Qt} \equiv \sum_{k=0}^{\infty} \frac{Q^k t^k}{k!},$$

where Q is the 9×9 rate matrix for the CTMC. Here the rate matrix is

$$Q = \begin{array}{c} A \\ B \\ C \\ (A,B) \\ (A,C) \\ (B,A) \\ (B,C) \\ (C,A) \\ (C,B) \end{array} \begin{pmatrix} -1/2 & 0 & 0 & 1/6 & 2/6 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 1/2 & 1/2 & 0 & 0 \\ 0 & 0 & -1/2 & 0 & 0 & 0 & 0 & 2/6 & 1/6 \\ 0 & 1/10 & 0 & -1/10 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/40 & 0 & -1/40 & 0 & 0 & 0 & 0 \\ 1/10 & 0 & 0 & 0 & 0 & -1/10 & 0 & 0 & 0 \\ 0 & 0 & 1/30 & 0 & 0 & 0 & -1/30 & 0 & 0 \\ 1/40 & 0 & 0 & 0 & 0 & 0 & 0 & -1/40 & 0 \\ 0 & 1/30 & 0 & 0 & 0 & 0 & 0 & 0 & -1/30 \end{pmatrix}$$

We get these entries as follows. For the transition rate from A to (A,B), we multiply the reciprocal of the mean idle time by the transition probability, getting $1/2 \times 1/3 = 1/6$. For the transition rate from (A,B) to B, we use the reciprocal of the mean travel time. The diagonal entries are minus the off-diagonal row sum.

Alternatively in this framework you could use a uniformization expression in the 9-state framework above. It is also possible to give a more complicated SMP expression, without introducing the extra states, but that is rather cumbersome.

---

3. Birth-and-Death (BD) Processes (30 points: 6 points each)

(a) Prove or disprove: For every irreducible finite-state BD process, we can express the transition probability matrix function $P(t) = (P_{i,j}(t)$ as

$$P(t) = \tilde{P}_{N(t)}, \quad t \geq 0.$$

where $\tilde{P}$ is the transition matrix of a finite-state discrete-time Markov chain (DTMC) and $\{N(t) : t \geq 0\}$ is a Poisson process that is independent of the DTMC.

This statement is VALID for any finite-state CTMC. It holds by the uniformization construction in §3.4 of the CTMC notes. A good answer would give the construction in (3.26)-(3.30) and then state and prove Theorem 3.4.

(b) Prove or disprove: For every irreducible finite-state BD process, we can express the transition probabilities as

$$P_{i,j}(t) = \sum_{n=0}^{\infty} \frac{a_{i,j}(n) t^n}{n!}, \quad t \geq 0,$$

where $a_{i,j}(n)$ are real numbers depending on n, the birth rates λ_k and the death rates μ_k (but are independent of t).

This statement is VALID for any finite-state CTMC. This is a minor re-wording of the matrix exponential representation

$$P(t) = e^{Qt}$$

which follows from the Kolmogorov ODE, as stated in Theorem 3.2 in the notes. Two proofs are given there: on pages 10 and 16. A good answer would give one.

(c) Prove or disprove: For every irreducible finite-state BD process, we can express the transition probabilities as

$$P_{i,j}(t) = b_j + \sum_{k=1}^{m-1} a_k e^{-r_k t}, \quad t \geq 0,$$

where b_j, a_k and r_k are real numbers with $b_j > 0$ for all j and $r_k > 0$ for all k, with b_j depending on j, the birth rates and the death rates (but is independent of i and t), while a_k depends on i, j, the birth rates and the death rates (but is independent of t).

This statement is VALID for any BD process, but NOT every finite-state CTMC. It is valid for all REVERSIBLE CTMC's. This is by the spectral representation discussed in §11.1. The stated formula here is (11.8) on p. 55. A good answer would state and prove Theorem 11.1.

(d) Prove or disprove: For every irreducible infinite-state BD process and initial state i, the expected time until the process makes k transitions is finite for all k.

This statement is VALID for any BD process. It is valid for any pure jump process (without any instantaneous transitions. For an irreducible BD process, the proof is especially easy. Starting in state i, the BD process can only visit the $2k+1$ states between $i-k$ and $i+k$ in the first k transitions. The time of each transition in state i is exponentially distributed with rate γ_i, and thus mean $1/\gamma_i$, where

$$\gamma_i \equiv \lambda_i + \mu_i,$$

where λ_i is the birth rate and μ_i is the death rate in state i. Let γ be the minimum of all these rates γ_i in all these $2k+1$ states. The expected time of each of the first k transitions is bounded above by $1/\gamma$. Thus the expected time for the process to make k transitions from state i is bounded above by k/γ.

(e) Prove or disprove: For every irreducible infinite-state BD process and initial state i, the expected time until the process makes infinitely many transitions is infinite.

This statement is FALSE. It is possible that there could be infinitely many transitions in finite time; i.e., there could be an explosion, as discussed in §10 of the CTMC notes. A pure birth process turns out to have an explosion if and only if the expected time until the process has infinitely many births is finite, i.e., if and only if

$$\sum_{k=0}^{\infty} \frac{1}{\lambda_k} < \infty,$$

as stated in Theorem 10.1. However, the pure-birth process is not irreducible. But we can see that the same problem occurs if the birth rates are $\lambda_n = n^2$, while the death rates are $\mu_n = 1$.

It is easy to see that the DTMC is transient, and will diverge to infinity. Then we can show that regularity does not hold for this example by applying applying Theorem 10.5.
