IEOR 6711: Stochastic Models, I

Fall 2012, Professor Whitt, Final Exam

SOLUTIONS

There are four questions, each with several parts.

1. Customers Coming to an Automatic Teller Machine (ATM) (30 points)

Customers arrive one at a time at a single automatic teller machine (ATM), with unlimited waiting space, to withdraw money. Customers **arrive** according to a Poisson process with rate $\lambda \equiv 48$ per hour. The **service times** of successive customers at the ATM can be regarded as i.i.d. random variables, distributed as the random variable S having a gamma distribution with mean $ES \equiv 1$ minute and standard deviation $\sqrt{Var(S)} \equiv 0.5$, and thus probability density function (pdf)

$$g(t) \equiv g_S(t) \equiv \frac{128t^3e^{-4t}}{3}, \quad t \geq 0,$$

and Laplace transform

$$\hat{g}(s) \equiv E[^{-sS}] \equiv \int_0^\infty e^{-st} g(t)\, dt = \left(\frac{4}{4+s}\right)^4.$$

The successive **withdrawals** can be regarded as i.i.d. random variables distributed as W having a gamma distribution with mean $EW \equiv 100$ dollars and standard deviation $SD(W) \equiv 110$ and thus Laplace transform

$$E[e^{-sW}] \equiv \left(\frac{1}{1+121s}\right)^{(1/1.21)}.$$

[first part: arrivals (numerical answers plus explanation desired)] (4 points)

(a) Suppose that 20 arrivals come during a given hour. What are the mean and variance of the number of these arrivals that come during the first 15 minutes of that hour?

The key fact is that, conditional on the number of arrivals in a Poisson process being n in a given interval, those n arrival times are distributed as n i.i.d. random variables, each uniformly distributed over the interval. Thus each arrival occurs in the first 15 minutes with probability $1/4$ and the total number that occur in the first 15 minutes has a binomial distribution with parameters $n = 20$ and $p = 1/4$. Thus,

$$\text{mean } = np = 20 \times (1/4) = 5 \quad \text{and} \quad \text{variance } = np(1-p) = 15/4.$$

(b) What is the probability that the first arrival completes service before the second customer arrives? (Assume that the system is initially empty.)

Let T_k be the interarrival time between the $(k-1)st$ arrival and the k^{th} arrival. Let S_k be the k^{th} service time. It suffices to look at the system starting at the moment the first customer arrives and immediately enters service. Measure time in minutes, so that $\lambda = 48/60 = 0.8$ Then

$$\begin{aligned} P(T_2 > S_1) &= \int_0^\infty e^{-\lambda t} g(t)\, dt = \hat{g}(\lambda) = \left(\frac{4}{4+\lambda}\right)^4 \\ &= \left(\frac{4}{4+.8}\right)^4 = \left(\frac{5}{6}\right)^4 = \frac{625}{1296}. \end{aligned}$$

[second part: total money withdrawn (numerical answers plus explanation desired)] (6 points)

(c) What are the mean and variance of the total amount of money withdrawn by all customers during one hour?

The key fact is that the total amount of money withdrawn in the interval $[0, t]$ as a function of t is a compound Poisson process with

$$\text{mean } = \lambda t E[W] = 48 \times 100 = 4,800$$

and

$$\begin{aligned} \text{variance } &= \lambda t E[W^2] = 48 \times (10,000 + 12,100) \\ &= 48 \times 22,100 = 1,060,800 \end{aligned}$$

Note that the second moment of W appears in the variance formula.

(d) What is the expected conditional total amount withdrawn in a given hour, given that the amount withdrawn in the previous hour is exactly two times the mean?

A compound Poisson process has stationary and independent increments so that conditioning event does not alter the mean. The mean is the same as in part (a): $4,800$.

(e) What is the approximate probability (to within 0.1) that the total amount of money withdrawn by all customers during one hour exceeds \$10,000?

We use a normal approximation, which is justified by the central limit theorem for a compound Poisson process. From part (c), the variance is approximately $1,000,000$, so that the standard deviation is approximately $1,000$. Thus, $10,000$ is more than 5 standard deviations above the mean. Hence,

$$\text{approximate probability is} \quad 0.$$

[third part: number of customers at the ATM] (20 **points)**

Let $X(t)$ be the number of customers at the ATM at time t.

(f) Is $\{X(t) : t \geq 0\}$ an irreducible aperiodic continuous-time Markov chain? Explain. (2 points)

No, because the Markov property fails. Since the service-time distribution is gamma, and not exponential (the special case of gamma with variance equal to the square of the mean), it does not have the lack-of-memory property. When $X(t) > 0$, the evolution of $\{X(s) : s \geq t\}$ after t depends on more than the state at time t. It also depends on when the customer in service at time t entered service.

(g) Identify random times T_n, $n \geq 0$, such that $\{X(T_n) : n \geq 0\}$ is an irreducible aperiodic discrete-time Markov chain with state space $\{0, 1, 2, \ldots\}$. (2 points)

We are now beginning to analyze this system as an $M/G/1$ queue, as discussed in the lecture notes of November 1. It suffices to let T_n be the departure time of the n^{th} arrival. If we look at these times, the system is either empty or a new service time is just beginning. Since the arrival process is Poisson, the times between arrivals are exponential and have the lack-of-memory property. It is easy to see that the process can always change by any amount greater than or equal to -1 from a positive value or by any amount greater than or equal to 0 after being 0. It is thus irreducible and aperiodic.

(h) For the DTMC identified in part (g), exhibit the transition probabilities. (4 points)

These transition probabilities are given on page 164 of the book, as part of Example 4.1 (A).

(i) For the DTMC identified in parts (g) and (h), prove that $X(T_n) \Rightarrow L$ as $n \to \infty$ for some random variable L with $P(L < \infty) = 1$ and characterize the probability distribution of L. (6 points)

Now we apply the analysis of Example 4.3 (A) in the book, starting on p. 177, as indicated in the lecture notes of November 1. Since the traffic intensity is $\rho \equiv \lambda E[S] = (48/60) \times 1 = 0.8 < 1$, the system is stable and a proper limit exists. By Theorems 4.3.1 and 4.3.3, it suffices to find a solution to the equation $\pi = \pi P$ which is nonnegative and sums to 1. The unique such π can be found by writing down the transition probabilities and using generating functions. That produces a single equation for the generating function, which can be solved That yields the Pollaczek-Khintchine transform (generating function), which is expressed in terms of the Laplace transform of g. In particular,

$$\hat{\pi}(z) \equiv E[z^L] \equiv \frac{(1-\rho)(z-1)\hat{g}(\lambda(1-z))}{z - \hat{g}(\lambda(1-z))} \ ,$$

where $\rho = 0.8$. Details are wanted here. See the book and the lecture notes for them.

(j) State and prove a limit theorem describing the distribution of $L(\lambda)$ if we let the arrival rate λ increase from 48 per hour to 60 per hour (while holding the service-time distribution unchanged). (6 points)

This is the heavy-traffic limit at the end of the notes on November 1. The limit of the scaled random variables $(1-\rho)L(\rho)$ as $\rho \uparrow 1$ is exponential with mean $(1+c_s^2)/2 = (1.25)/2 = 0.625$ As $\lambda \uparrow 60$, we have $\rho \uparrow 1$. We use Taylor series approximation for the Laplace transform $\hat{g}(s)$ and the characteristic function of $(1-\rho)L(\rho)$. See the lecture notes for the details.

---

2. The Department of Motor Vehicles (DMV) (26 points)

You can get a new automobile driver's license at the New York Department of Motor Vehicles (DMV) on 34th Street. The standard process for getting a new license involves passing through three stages of service. First, you get in a single line to wait for your turn to be served by a single clerk to get the correct form; second, you get in a single line to wait for a photographer to take your picture; and, third, you wait in a single line (you actually sit in a waiting room with order maintained by being assigned a number) for one of several clerks to complete the processing, prepare your license and collect your money. Your goal is to analyze the performance in this service system.

Make the following assumptions: Suppose that customers arrive according to a Poisson process with a rate of 2 per minute. Suppose that the service times at each step have exponential distributions. Suppose that the service times of different customers and of the same customers at different steps are all mutually independent. Suppose that all arriving customers go to the first clerk. Suppose the mean service time there is 15 seconds. Suppose that there is unlimited waiting space and that all customers are willing to wait until they can be served (Nobody gets impatient and abandons.) Suppose that only 1/4 of all arriving customers seek new licenses and must complete the second and third steps. The other 3/4 of the arriving customers go elsewhere after finishing the first stage. Suppose that the mean time required for the photographer to take a picture is 1 minute. Suppose that the mean time for each of the final clerks to complete the processing, prepare your license and collect your money is 10 minutes.

[first part: basic performance analysis] (10 points)

(a) How many clerks are need at the third stage (to complete the processing, prepare your license and collect your money) in order for the total customer arrival rate at the third stage to be strictly less than the maximum possible service rate (assuming all servers are working)?

Because only 1/4 of the customers go on to the second stage to have their pictures taken, the arrival rate at the second and third stages is $2/4 = 0.5$ per minute. Each clerk at the third stage serves at rate 0.1 per minute. If there are s clerks, then the maximum service rate at the third stage (which prevails whenever all the clerks there are busy) is $0.1s$. We require that

$$0.5 < 0.1s \quad \text{or} \quad s > 5.$$

That is, we must have at least 6 clerks at the third stage in order for the total customer arrival rate at the third stage to be strictly less than the maximum possible service rate (assuming all servers are working). That condition is needed to ensure the existence of a proper steady state. Otherwise the number in the system would explode (without customer abandonment).

Henceforth assume that there are 8 clerks at the third stage.

(b) Suppose that the first customer of the day (who finds a completely empty system upon arrival) wants to get a new driver's license. What are the mean and variance of the total time that this initial customer must spend at the DMV in order to get the license?

The first customer never has to wait. This customer only experiences his three service times. This customer's expected times at the three stations are, respectively, 0.25, 1.0 and 10.0 minutes, so the customers expected total time in the system is 11.25 minutes.

Since the service times are independent exponential random variables, the variance of the sum is the sum of the variances, while the variance of each is the square of the mean. Hence the variance of the times spent are $(0.25)^2 = 0.0625$, $1^2 = 1$ and $(10)^2 = 100$. Finally, the variance of the total time is 101.0625.

(c) Suppose that the second customer of the day also wants to get a new driver's license, and suppose that this second customer finds the first customer still being served by the first clerk when he arrives. What is the expected total time that this second customer must spend at the DMV in order to get the license?

We use properties of exponential distributions. First, by the lack of memory property the remaining time of the first customer at the first clerk is exponential with mean 0.25 minutes, just as it was at the beginning. After that, we should consider what happens first: service of the first customer by the photographer or service of the second customer by the first clerk. (We now use properties of the minimum of two independent exponential random variables.) The expected time required for the first of these two events is $1/(1 + 4.0) = 0.2$ minutes. The second customer finishes first with probability $4/(4+1) = 4/5$, while the first customer finishes first with probability $1/5$. If the first customer finishes first, then the first customer no longer can be in the way of the second customer. So the total remaining expected wait for the second customer is 11.25, by part (b) above. On the other hand, if the second customer finishes first, then the remaining expected time for the second customer is $1 + 1 + 10 = 12$ minutes, because both customers have a full service time at the photographer, after which the second customer will have his own clerk at the third stage. Thus the expected total time required for the second customer is

$$0.25 + 0.20 + (1/5)(11.25) + (4/5)(12.00) = 0.45 + 2.25 + 9.60 = 12.30$$

The expected waiting time for the second customer is 1.05 minutes longer than the expected waiting time of the first customer under these circumstances.

(d) Give an expression for the steady-state probability that there are 3 customers either waiting or being served at the first clerk, and 5 customers either waiting or being served at the photographer.

This is an open queueing network as discussed in §of the CTMC notes. Recall from the discussion of reversibility, that the steady-state numbers at the three queues are independent random variables; see Theorem 6.7 of the CTMC notes. Let X_j be the number of customers either waiting or being served at queue j. Let ρ_j be the traffic intensity at queue j. Hence,

$$P(X_1 = 3 \quad \text{and} \quad X_2 = 5) = P(X_1 = 3)P(X_2 = 5) = (1-\rho_1)\rho_1^3(1-\rho_2)\rho_2^5 = (1/2)^{10} = 1/1024$$

(e) Give an expression for the steady-state probability that exactly 4 customers complete service from the photographer during one specified minute in steady state.

Now we use the fact that in steady state the departure process from the first queue is also a Poisson process with the same rate as the arrival rate, which is $2 \times (1/4) = 0.5$ per minute; see Theorem 6.5 of the CTMC notes. Let $C(t)$ be the number of customers that complete service from the first clerk during an interval of length t. Then

$$P(C(t) = 4) = \frac{e^{-0.5t}(0.5t)^4}{4!}$$

Since $t = 1$, we get

$$P(C(1) = 4) = \frac{e^{-0.5}(0.5)^4}{4!}$$

[second part: supporting theory] (10 points)

Let $X_j(t)$ be the number of customers at station j, either waiting or being served.

(f) Prove or disprove: With an appropriate stationary distribution, the stochastic process $\{(X_1(t), X_2(t), X_3(t)) : t \geq 0\}$ is a time reversible irreducible CTMC. (3 points)

This statement is false. As discussed on top of p. 34 of the CTMC lecture notes, reversibility fails. It is possible to go from state(i,j,k) to state $(i-1, j+1, k)$ by a service completion at queue 1, but it is not possible to go in the reverse direction.

(g) Prove or disprove: The stochastic process $\{(X_1(t), X_2(t), X_3(t)) : t \geq 0\}$ has a proper limiting distribution, i.e.,

$$\lim_{t\to\infty} P(X_1(t) = j_1, X_2(t) = j_2, X_3(t) = j_3)|X_1(0) = i_1, X_2(0) = i_2, X_3(0) = i_3) = \alpha_{(j_1,j_2,j_3)}$$

for all vectors of nonnegative integers (j_1, j_2, j_3) and (i_1, i_2, i_3), where

$$\sum_{j_1=0}^{\infty}\sum_{j_2=0}^{\infty}\sum_{j_3=0}^{\infty} \alpha_{(j_1,j_2,j_3)} = 1.$$

(7 points)

This statement is true. This is a minor variant of Theorem 6.7. We need a slight extension because some of the departures from the first queue leave the system. Thus we apply Theorem 6.9. To do so, we verify the conditions of Theorem 6.8.

[third part: starting over] (6 points, 2 points each)

Suppose that, with probability 1/5, independent of the history, each customer after completing all three stages of service has to return immediately to the end of the first queue and start the process over, with new independent service times.

(h) What is the new answer to part (a) above under this new condition?

We now have more general traffic rate equations to solve. The net arrival rate to the first queue is λ. We have the equation

$$2 + \frac{\lambda}{20} = \lambda$$

because 1/4 of the input to station 1 goes on to station s 2 and 3, and then 1/5 of the final flow comes back to station 1. Hence,

$$\lambda = 40/19$$

and the net arrival rate at station 3 is $(40/19) \times (1/4) = 10/19$. As before, we see that 5 severs at station 3 is not enough, but 6 is enough. With $s = 6$ servers at station 3, the traffic intensity there is

$$\rho_3 = \frac{(10/19) \times 10}{6} = \frac{100}{114} = 0.877 < 1.$$

(i) How do the answers to parts (f) and (g) above change under this new condition?

No change. We still have a Markovian queueing network.

(j) Prove or disprove: Under this new condition, the arrival process at station 2 when the system is in steady state (with a stationary initial distribution) is a Poisson process.

The arrival process is **not** a Poisson process. Even though there is a well defined steady-state distribution, which has the same form as if the arrival process to each queue were a Poisson process, the arrival processes to the queues are not Poisson processes because of the feedback. Here is a counterexample: We show that the arrival process does not have independent increments. We can see that the arrival process in the future after time t depends on more than the current state. If we condition on having 10^{100} arrivals in the interval $[t - 100, 100]$, then there are likely to be more arrival after time t than if we did not have that unlikely condition, because many of these recent arrivals will return.

---

3. Random Walk on a Graph (24 **points**)

Consider the graph shown in the figure on the top of the next page. There are 7 nodes, labeled with capital letters and 8 arcs connecting some of the nodes. On each arc is a numerical weight. Six of the arcs have weight 1, while two of the arcs have weight m. Consider a random walk on this graph, which moves randomly from node to node, always going to a neighbor, via a connecting arc. Let each move be to one of the current node's neighbors, with a probability proportional to the weight on the connecting arc, independent of the history prior to reaching the current node. Thus the probability of moving from node A to node C in one step is $1/(1+m)$, while the probability of moving from node C to node A in one step is $1/(1+1+1)$ $= 1/3$. Let X_n be the node occupied after the n^{th} step of the random walk. Suppose that $X_0 = A$.

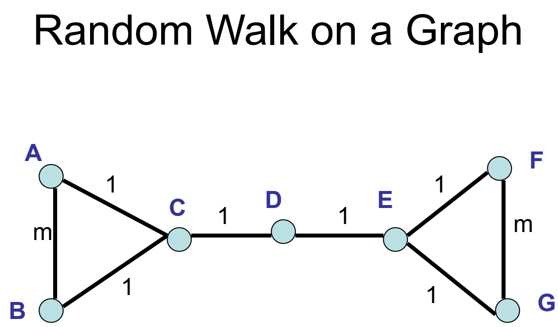

(a) (4 points) For any nodes A and B, let $T_{A,B}$ be the first passage time (number of steps) required to go from A to B, with $T_{A,A} \geq 1$. Calculate $E[T_{A,A}]$. Briefly explain.

Here we recognize that this is the random walk on a weighted graph, involving reversibility in DTMC's. Let π_A be the long-run proportion of moves ending in the node A. We know that π_A is the sum of the weights out of A divided by the sum over all nodes of the sum of weights out of that node. Thus

$$\pi_A = \frac{m+1}{4(m+1)+8} = \frac{m+1}{4m+12}.$$

Then $E[T_{A,A}]$ is the reciprocal of π_A,

$$\frac{1}{\pi_A} = \frac{4m+12}{m+1}.$$

That follows from renewal theory.

(b) (4 points) Let $Z_{A,B,C}$ be the number of visits to B starting in A before hitting C. Give an expression for $E[Z_{A,B,F}]$ and justify your answer.

This part is answered by applying the absorbing theory of DTMC's. We make node F an absorbing state. Then the remaining 7 nodes become transient states. We then put the transition matrix in canonical form, which we can denote by

$$P = \begin{pmatrix} I & 0 \\ R & Q \end{pmatrix},$$

where the first row is for the single absorbing state, I is a 1×1 identity matrix, giving the transition probabilities among the absorbing states, which is only the single state F, 0 is a 1×7 matrix of 0's giving the transition probabilities from the absorbing state to the transient states, R is a 7×1 transition matrix giving the one-step probabilities of being absorbed in the absorbing state starting in each of the 7 transient states, and Q is the 7×7 square transition matrix among the 7 transient states (all states except F). The expected value desired is $E[Z_{A,B,F}] = N_{A,B}$, which is the (A, B) entry of the square matrix

$$N \equiv (I - Q)^{-1},$$

corresponding to row A and column B. This formula is justified in §2.3.1 of the notes of October 23.

(c) (3 points) Calculate numerical values for $E[Z_{A,B,C}]$ and $E[Z_{F,G,E}]$.

We can apply the previous part, after changing the absorbing state from F to C. However, if we make C absorbing and start from A, then we never can reach the states D, E, F and G. It suffices to consider only the states A, B and C. It suffices to look at

$$P = \begin{pmatrix} I & 0 \\ R & Q \end{pmatrix},$$

where the first rows is the single absorbing state, I is a 1×1 identity matrix, giving the transition probabilities among these absorbing states, 0 is a 1×2 matrix of 0's giving the transition probabilities from the absorbing state to the two transient states A and B, R is a 2×1 transition matrix giving the one-step probabilities of being absorbed in the absorbing state starting in each of the 2 transient states, and Q is the 2×2 square transition matrix among the 2 transient states. The expected value desired is $N_{A,B}$, which is the (A, B) entry of the square matrix

$$N \equiv (I - Q)^{-1},$$

corresponding to row A and column B.

Let $N_{A,B} \equiv E[Z_{A,B,C}]$ be the expected total number of visits to B starting in A. To have positive expected value, we must get to B in the first transition, which occurs with probability $m/(m + 1)$. Then to get more contribution, we must return to A, which again occurs with probability $m/(m + 1)$, after which we start over. (Recall Example 3.13 on p. 110.) Thus, we can write down the recursion:

$$N_{A,B} = \left(\frac{m}{m+1}\right)\left(1 + \left(\frac{m}{m+1}\right) N_{A,B}\right).$$

We can then solve for $N_{A,B}$, getting

$$E[Z_{A,B,C}] \equiv N_{A,B} = \frac{m/(m+1)}{1 - (m^2/(m+1)^2} = \frac{m(m+1)}{2m+1} \approx \frac{m}{2} \quad \text{for} \quad \text{large} \quad m.$$

We have

$$E[Z_{F,G,E}] = E[Z_{A,B,C}]$$

by **symmetry**. The model is the same if we reflect the model about a vertical line through node D. That is, we map the state vector (A, B, C, D, E, F, G) into the new state vector

(F, G, E, D, C, A, B). There are other symmetries here as well. By similar reasoning, we also have $E[T_{A,D}] = E[T_{B,D}]$. In this case we can reflect the graph about the horizontal axis, through nodes C, D and E. We then change the state vector (A, B, C, D, E, F, G) into the new states (B, A, C, D, E, G, F).

(d) (4 points) Calculate $E[z^{Z_{A,B,C}}]$.

We can apply a modification of the previous recursion to calculate the probability generating function $x(z) \equiv E[z^{Z_{A,B,C}}]$ of the random variable $Z_{A,B,C}$:

$$\begin{aligned} x(z) &= (1/(m+1))z^0 + (m/(m+1))z + (m/(m+1))^2 E[z^{1+Z_{A,B,C}}] \\ &= (1/(m+1)) + (m/(m+1))z + (m/(m+1))^2 (zx(z)), \end{aligned}$$

so that

$$x(z) = \frac{1/(m+1) + (m/(m+1))z}{1 - z[m/(m+1)]^2}.$$

Alternatively, we can directly show by mathematical induction that

$$P(Z_{A,B,C} = k) = \left(\frac{m}{m+1}\right)^{2k-1} \left(\frac{2m+1}{(m+1)^2}\right), \quad k \geq 1.$$

Initially, we get

$$P(Z_{A,B,C} = 0) = \frac{1}{m+1},$$

$$P(Z_{A,B,C} = 1) = \left(\frac{m}{m+1}\right)\left(1 - \left(\frac{m}{m+1}\right)^2\right) = \left(\frac{m}{m+1}\right)\left(\frac{2m+1}{(m+1)^2}\right)$$

and

$$P(Z_{A,B,C} = 2) = \left(\frac{m}{m+1}\right)^3 \left(\frac{2m+1}{(m+1)^2}\right).$$

(e) (3 points) Prove or disprove: If a discrete-time stochastic process $\{X_n : n \geq 0\}$ has a unique proper limiting distribution as $n \to \infty$, then that limiting distribution is the unique stationary distribution.

This statement is false. A counterexample is given in Example 4.1 (b) of the CTMC notes. The continuous-time example there can easily be converted to discrete time.

(f) (3 points) Prove or disprove: Every finite-state DTMC $\{X_n : n \geq 0\}$ has a stationary distribution.

This is TRUE. This follows because there has to be at least one closed communication class. It is impossible for all states in a finite DTMC to be transient. That follows from Theorems 4.1 and 4.2 in the book. Each closed communication class has a unique stationary distribution concentrating on that class. The class of *all* stationary distributions for the entire DTMC is thus the set of all positive convex combinations of the unique stationary distributions for the closed classes.

(g) (3 points) For the random walk on the graph, prove that there exists a DTMC with transition matrix denoted by $P(\infty)$ such that

$$P(m) \to P(\infty) \quad \text{as} \quad m \to \infty,$$

where by convergence of matrices we mean that all elements converge. Exhibit all stationary probability vectors for the DTMC with transition matrix $P(\infty)$.

Note that convergence takes place. It suffices to consider the transitions out of A, B, F and G, because they are the only ones that depend upon m. But these transitions all have the same form, so we actually only need to consider one of these nodes. We see that the limiting probability transition matrix is

$$P(\infty) = \begin{array}{c} A \\ B \\ C \\ D \\ E \\ F \\ G \end{array} \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1/3 & 1/3 & 0 & 1/3 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 0 & 1/2 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 0 & 1/3 & 1/3 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

It is a bona fide transition matrix for a DTMC because the elements are nonnegative and the row sums are all 1. However, unlike $P(m)$, there are two closed communication class and one open communication class. The two closed communication classes are $\{A, B\}$ and $\{F, G\}$. The open class contains the remaining states. If we put it in canonical form, then we should group the two closed sets at the top.

$$P(\infty) = \begin{array}{c} A \\ B \\ F \\ G \\ C \\ E \\ D \end{array} \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1/3 & 1/3 & 0 & 0 & 0 & 0 & 1/3 \\ 0 & 0 & 1/3 & 1/3 & 0 & 0 & 1/3 \\ 0 & 0 & 0 & 0 & 1/2 & 1/2 & 0 \end{pmatrix}.$$

As pointed out in the solution to the previous part, the set of all stationary probability vectors in the set of all positive convex combinations of the stationary vectors on the two closed classes $\{A, B\}$ and $\{F, G\}$. The unique stationary vector on each class is $(1/2, 1/2)$. This produces a one-dimensional infinite family In the canonical form, the set of all stationary probability vectors on the state vector (A, B, F, G, C, E, D) is

$$\{(a/2, a/2, (1-a)/2, (1-a)/2, 0, 0, 0) : 0 \leq a \leq 1\}.$$

Recall that the DTMC's $P(m)$ are all irreducible. But as m increases, this DTMC becomes **nearly completely decomposable (NCD)**. The limit is $P(\infty)$ reducible. There is a substantial theory of NCD NC's and many applications.

---

4. Sums of i.i.d Positive Random Variables (20 points, 5 points each)

Let $\{X_n : n \geq 1\}$ be a sequence of i.i.d. positive random variables with pdf f and mean $E[X_n] = 2$. Let $S_n \equiv X_1 + \cdots + X_n$, $n \geq 1$, and $S_0 \equiv 0$. For any $t > 0$, let

$$w(t) \equiv \sum_{n=1}^{\infty} P(S_n \leq t)$$

and

$$V(t) \equiv S_n - t,$$

for the n such that $S_{n-1} \leq t < S_n$.

(a) Prove or disprove: $w(t) < \infty$ for all t, $0 \leq t < \infty$.

This question is asking standard questions about renewal theory, but in a slightly disguised form. This statement is true. Notice that $w(t)$ is the renewal function, usually expressed as $m(t) = E[N(t)]$; i.e., apply Proposition 3.2.1. Then this is Proposition 3.2.2. Its proof is desired.

(b) Prove or disprove: $\lim_{t\to\infty} (w(t)/t) = c$ for some constant c, $0 < c < \infty$.

After the interpretation, this is the elementary renewal theorem, Theorem 3.3.4. Its proof is desired.

(c) Determine $E[e^{-sV(t)}]$.

Notice that $V(t)$ is just the familiar excess, usually expressed as $Y(t) \equiv S_{N(t)+1} - t$. What is wanted, therefore, is the renewal equation for it and then its solution obtained via Laplace transforms. See (20) in §6 and §2 of the notes of Oct 9 and §3.1 of the notes of Oct. 11. We see that the answer is $\hat{h}(s)/(1-\hat{f}(s))$, where $h(t) = F(t+x) - F(t)$, so that $\hat{h}(s) = (e^{sx}-1)\hat{f}(s)/s$.

(d) Prove or disprove: $V(t)$ converges in distribution to a proper limit as $t \to \infty$.

The limit exists by the key renewal theorem, as discussed in the notes of October 11. The limit is distributed as F_e, where F_e is the stationary-excess cdf, i.e.,

$$F_e(x) \equiv \frac{1}{E[X]} \int_0^x P(X > s)\, ds.$$
