Assignment #2 (Due February 13)

1. Consider the distribution function

$$F(x)=\begin{cases}0 & \text{for } x<-2\\ (x/4)+(1/2) & \text{for } -2\le x<0\\ (x^2+1)/2 & \text{for } 0\le x<1/2\\ 7/8 & \text{for } 1/2\le x<3\\ 1 & \text{for } x\ge 3\end{cases}$$

a) Sketch the function $F(x)$. Assuming that a uniform random number generator is available, use the inversion method to provide an algorithm for generating a random variable X having distribution function F.

b) Compute $E[X]$ and $Var[X]$ analytically. (One way of verifying your random number generator would be to compute the sample mean and variance of a large number of generated samples, and see whether what you get is close to the analytical values.)

c) Prove the following converse of the inversion method: if X is a random variable having a continuous, strictly increasing distribution function F, then $U=F(X)$ is uniformly distributed on $[0,1]$. As discussed later, this result can be used to generate pairs of random values with specified correlation. [Hint: the assumptions on F ensure that $F^{-1}\big(F(x)\big)=x$ for any real number x, and $F\big(F^{-1}(y)\big)=y$ for any $y\in[0,1]$. Also note that F^{-1} is an increasing function.]

2. (**Computing problem**: Hospital unit) A hospital unit consists of a standard care unit (SCU) containing b_1 beds (numbered $1,2,\ldots,b_1$) and an intensive care unit (ICU) containing b_2 beds (numbered $1,2,\ldots,b_2$). Patients arrive according to a renewal process with the successive inter-arrival times i.i.d. as a random variable A having probability density function

$$f(x;a)=\begin{cases}2x/a^2 & 0\le x\le a\\ 0 & \text{otherwise}\end{cases},$$

where a is a parameter of the distribution. (Throughout, all times are measured in units of days.) With probability p, an arrival is an SCU patient, and with probability $1-p$, an arrival is an ICU patient. An arrival to the SCU is assigned to the lowest-numbered available bed, and similarly for the ICU. If there is no bed available for an arriving patient, the patient is immediately rerouted to another hospital and is thus lost from the system. The length of stay for an SCU patient (in the absence of a "critical event" as described below) and the length of stay in the ICU are distributed as random variables L_{SCU} and L_{ICU}, respectively, each having a 2-Erlang(λ) distribution. Such a random variable is distributed as the sum of two independent exponential random variables, each with intensity λ. An SCU patient may experience a "critical event" at some point after arrival, at which point the patient becomes an

ICU patient, and stays in the ICU for an amount of time distributed as L_{ICU} (i.e., the same as for a freshly arrived ICU patient). The time until a critical event for an SCU patient is distributed as a random variable Q having a symmetric triangular distribution on the interval [1/2, 3/2]; cf Problem 3 of Assignment #1. If, for a given patient, the critical-event time Q exceeds the hospitalization time L_{SCU}, then the patient is released from the hospital without ever experiencing a critical event. If an SCU patient experiences a critical event and there are no available beds in the ICU, the patient is immediately transferred to another hospital. Define the state of the system at time t as

$$X(t) = \Big(M_1(t), M_2(t), \dots, M_{b_1}(t), N_1(t), N_2(t), \dots, N_{b_2}(t)\Big),$$

where $M_i(t) = 1$ if the ith SCU bed is occupied at time t and $M_i(t) = 0$ otherwise, and $N_i(t) = 1$ if the ith ICU bed is occupied at time t, and $N_i(t) = 0$ otherwise. Assume that the hospital unit is initially empty (set the clock for the initial patent arrival as a fresh sample from the distribution of A.)

a) Specify $\{X(t) : t \ge 0\}$ as a GSMP with event set $E = \{e_0\} \cup \{e_{1,1}, \dots, e_{1,b_1}\} \cup \{e_{2,1}, \dots, e_{2,b_2}\}$ $\cup \{e_{3,1}, \dots, e_{3,b_1}\}$, where e_0 = "arrival of patient", $e_{1,i}$ = "departure of patient from SCU bed i", $e_{2,i}$ = "departure of patient from ICU bed i", and $e_{3,i}$ = "critical event for patient in SCU bed i". Give your specification generically in terms of $A, L_{\text{SCU}}, L_{\text{ICU}}, Q$. [Hint: you may want to define auxiliary quantities such as $i_{\text{SCU}}(s)$ = the lowest-numbered available bed in the SCU when the state is $s = (m_1, \dots m_{b_1}, n_1, \dots, n_{b_2})$, where $i_{\text{SCU}}(s) = -1$ if there are no empty beds.]

b) In terms of the process $\{X(t) : t \ge 0\}$, give a precise specification of (i) the expected time until the hospital unit (SCU + ICU) first fills up all of its beds, (ii) the expected fraction of time during the first thirty days of operation during which there are 7 or more empty beds, (iii) the probability that the time until the hospital unit first fills up all of its beds is less than or equal to 55 days, and (iv) the expected fraction of patient arrivals in [0, 30] that are immediately transferred to another hospital. [Define the fourth performance measure directly in terms of the GSSMC $\{(S_n, C_n) : n \ge 0\}$, using notions defined in class, such as the nth transition time ζ_n, the trigger-event function $E^*(s, c)$, and so forth. You may also want to use indicator functions.]

c) Now assume that $b = (5, 4)$, $a = 0.08$, $E[L_{\text{SCU}}] = E[L_{\text{ICU}}] = 0.25$, and $p = 0.68$. Using the general algorithm for simulating GSMPs given in class, estimate the four performance measures from part (b) to within $\pm 1\%$ with probability 99%.

3. (Extra Credit: Markov chain Monte Carlo) In this problem we develop a Monte Carlo method for estimating an infinite sum of the form $S = \sum_{i=-\infty}^{\infty} g(i)\pi(i)$, where g is a real-valued function and π is a probability distribution on the integers. For example, π might be of the form $\pi(0) = \theta$ and $\pi(i) = \theta i^{-2} e^{\cos(i)}$ for $i \ne 0$, where θ is a normalizing constant chosen so that $\sum_{i=-\infty}^{\infty} \pi(i) = 1$. In applications (especially Bayesian data analysis), θ is often unknown, so that the distribution π is

known only up to a normalizing constant. (Also, the distribution π is usually multivariate, typically with very high dimension, in which case the MCMC method developed here is even more effective.) Consider a state-transition matrix Q for a DTMC that is "irreducible" in that for each i and j there exists $n = n(i,j)$ such that $Q^n(i,j) > 0$. (Intuitively, an irreducible DTMC can, with positive probability, go from one state to any other state in a finite number of transitions.) Set

$$\alpha(i,j) = \min\left(1, \frac{\pi(j)Q(j,i)}{\pi(i)Q(i,j)}\right)$$

for all i and j, and consider a real-valued DTMC $\{X_n : n \geq 0\}$ that is generated according to the following algorithm (which can be viewed as a sophisticated version of acceptance-rejection sampling):

(1) Initialize X_0 and set $m = 0$.
(2) Generate a random variable Y according to $Q(X_m, \cdot)$.
(3) Generate a Uniform[0,1] random variable U.
(4) If $U \leq \alpha(X_m, Y)$ then set $X_{m+1} = Y$, else set $X_{m+1} = X_m$.
(5) Set $m \leftarrow m+1$ and go to Step (2).

a) Show that $\pi(i)Q(i,j)\alpha(i,j) = \pi(j)Q(j,i)\alpha(j,i)$ for all i and j. [Hint: evaluate $\alpha(i,j)$ and $\alpha(j,i)$, both when $\pi(j)Q(j,i) > \pi(i)Q(i,j)$ and when $\pi(j)Q(j,i) \leq \pi(i)Q(i,j)$.]

b) Let P be the state-transition matrix for the DTMC $\{X_m : m \geq 0\}$. For i and j with $i \neq j$, show that $P(i,j) = Q(i,j)\alpha(i,j)$, and then show that π is a stationary distribution for $\{X_m : m \geq 0\}$. [Hint: recall from Lecture #2, slide 10, that π is a stationary distribution if $\pi(j) = \sum_i P(i,j)\pi(i)$. To establish the desired property for π, first use part (a) to show that $\pi(i)P(i,j) = \pi(j)P(j,i)$ for all i and j.]

c) It can be shown that the DTMC $\{X_m : m \geq 0\}$ obeys a strong law of large numbers in that $\lim_{n\to\infty}(1/n)\sum_{i=1}^{n} g(X_i) = E[g(X)]$ a.s., where X is distributed according to the stationary distribution π. Using this fact, write down a complete and detailed algorithm that will generate $X_0, X_1, \ldots, X_n$ (where n is a given, large value) and produce a strongly consistent point estimate of the infinite sum S defined above, for the specific distribution π given above. (Don't worry about how to select the value of n or about providing a confidence interval; in general, we can use the regenerative or batch means methods—discussed later in the course—to do this.) Your algorithm should incorporate and refine the general algorithm given above. Take $X_0 = 0$ and, given X_m, have your algorithm generate Y as $X_m + W$, where W is uniformly distributed on the set $\{-k, -k+1, \ldots, -1, 0, 1, \ldots, k-1, k\}$ for a fixed integer k. Your algorithm should *not* require that you know or calculate the value of the normalizing constant θ in the definition of π. Show that, for our specific method of generating a candidate Y, the general definition of $\alpha(i,j)$ takes on a somewhat simpler form, so that step (4) of the general algorithm simplifies.

4. (Extra Credit: Analytical solution of gambling game) We can analytically compute the expected length, $E[L]$, and hence the expected gain, $E[8.99 - L] = 8.99 - E[L]$, for the gambling game, using Markov chain theory. Let X_n be the absolute difference between the number of heads and tails after n flips of the coin.

   a) Prove that $\{X_n : n \geq 0\}$ is a Markov chain with state space $S = \{0,1,2,3\}$. [Hint: Write down a recurrence relation.]

   b) Write down the initial distribution μ and transition matrix P for the chain. (Assume that when the absolute difference of heads and tails becomes equal to 3, the chain just stays in that state forever.)

   c) Denote by $E_i[L]$ the expected length of the game when the initial state is $X_0 = i$. Using the Markov property, it can be shown that the following "first-step decomposition" holds:

   $$E_i[L] = 1 + \sum_{j \in S} P(i,j) \cdot E_j[L] \text{ for } i = 0,1,2$$

   and $E_3[L] = 0$. Intuitively, if the current state is $i \in \{0,1,2\}$, then the expected length is 1 (for the current flip) plus the expected number of remaining flips. If the current flip results in a transition to state j, then, by the Markov property, the expected number of remaining flips is the same as if you *started* the game in state j, that is, $E_j[L]$. Solve the set of equations defined by the first-step decomposition to compute the actual expected gain, which is $8.99 - E_0[L]$.