OR221 MIDTERM SOLUTIONS

Due: 9:30 am, Thursday, March 9, 2006

P1. (10 pts.) **Uniformization.** Consider a two state CTMC on state space {0,1} with transition rates $q_{0,1} = \lambda$ and $q_{1,0} = \mu$. Compute the transition probability matrix $P(t)$ for this CTMC by using the method of uniformization.
Hint: The uniformization method works with any q bounded below by the one given in Eq. (6.66). Try using $q = \lambda + \mu$ in Equation (6.72).

Ans. We have

$$Q = \begin{bmatrix} -\lambda & \lambda \\ \mu & -\mu \end{bmatrix}.$$

Use $q = \lambda + \mu$ as the uniformization constant. From Equation (6.69) we get

$$\hat{Q} = \begin{bmatrix} \frac{\mu}{\lambda+\mu} & \frac{\lambda}{\lambda+\mu} \\ \frac{\mu}{\lambda+\mu} & \frac{\lambda}{\lambda+\mu} \end{bmatrix}.$$

Hence we get $\hat{Q}^0 = I$ and $\hat{Q}^n = \hat{Q}, \;\; n \geq 1$. Substituting in Equation (6.72) and carrying out the summation, we get the expression for $P(t)$ given in Equation (6.38).

P2. (10 pts.) **Gasoline Prices.** Let $X(t)$ be the price of gasoline (in \$ per gallon) at time t (time is measured in weeks). Suppose that $X(t) = 2.00 + .1Y(t)$, where $\{Y(t), t \geq 0\}$ is a birth and death process on $S = \{0, 1, ..., 9\}$ with the following birth and death rates:

$$\lambda_i = .5, \;\; i = 0, 1, ..., 8,$$

$$\mu_i = .5, \;\; i = 1, 2, ..., 9.$$

Consider two customers Al and Bob. Al buys 10 gallons of gasoline at 8 am on every Monday. Bob also buys 10 gallons of gasoline every time he purchases gasoline, but the times between two of Bob's consecutive purchases are i.i.d. random variables that are exponentially distributed with mean 1 week.

(a). Suppose both Bob and Al buy gas at 8am on the first Monday in February and pay \$2.10 per gallon . What is the expected cost of the next purchase for Al and Bob?

(b). In the long run how much money do Al and Bob spend on gasoline per week?

Ans.

(a). Let Y_n be the state of the birth and death process at 8am on the n th Monday. Then $\{Y_n, n \geq 0\}$ is a DTMC with transition matrix $P(1) = expm(Q)$, where Q is the generator matrix of the Y process. Hence, using Matlab, the distribution of Y_1 given $Y_0 = 1$ is seen to be [0.0011 0.0082 0.0499 0.2079 0.4658 0.2079 0.0499 0.0082 0.0010 0.0001]. Hence the expected cost of Al's purchase on the next Monday is $\$10 * (2 + .1 * E(Y_1|Y_0 = 1)) = \$21.08$.
Now let Z_n be the state of the Y process at the time of the n th purchase by Bob. Then $\{Z_n, n \geq 0\}$ is a DTMC with transition matrix $(I - Q)^{-1}$, since Bob observes the Y process at the events of a PP(1). Hence the expected cost of Bob's next purchase is $\$10 * (2 + .1 * E(Z_1|Z_0 = 1)) = \$21.0981 = \$21.10$.

(b). For the birth and death process $\{Y(t), t \geq 0\}$ we get

$$\rho_j = 1, \quad j = 0, 1, 2, ..., 9.$$

Hence the limiting distribution of $Y(t)$ is given by

$$p_j = \lim_{t \to \infty} P(Y(t) = j) = .1, \quad 0 \leq j \leq 9.$$

Hence

$$p_j = \lim_{t \to \infty} P(X(t) = 2 + .1j) = .1, \quad 0 \leq j \leq 9.$$

Now, since Bob visits the gas pump according to PP, he sees the time averages, according to PASTA. Hence he pays $ 2 + .1j per gallon with probability .1 for $0 \leq j \leq 9$. Thus his weekly cost of gasoline is 10*(2 + .1(0+1+...+9)*.1) = $ 24.50. To compute Al's cost, we need to know the distribution of the price on each Monday. Let Y_n be the state of the birth and death process at 8am on the n th Monday. Then $\{Y_n, n \geq 0\}$ is a DTMC with transition matrix $P(1) = expm(Q)$, where Q is the generator matrix of the Y process. Then the n-step transition probability matrix of the Y_n process is given by $expm(Q*n)$. Hence the limiting distribution of the DTMC Y_n is the same as that of the CTMC Y. Hence Al's weekly cost is also given by $24.50.

P3. (10pts.) **Queue with Catastrophes.** Customers arrive at a single server queue with finite waiting room of size K according to $PP(\lambda)$ and request i.i.d. $Exp(\mu)$ service times. Any customer who arrives at a full system is lost. The system is subject to catastrophes that occur according a $PP(\theta)$. When a catastrophe occurs, all the customers in the system are lost. It takes an $Exp(\alpha)$ amount of time to recover from a catastrophe. During this recovery phase no customers can be admitted to the system, and any occurrences of catastrophes have no effect. Model this as CTMC as follows:
(1). Show the state space.
(2). Show the rate diagram.
(3). Write down the balance equations.
(4). Solve them. When is this system stable?
(5). Compute the expected number of customers lost when a catastrophe occurs in the long run.

Ans.

(1). $S = \{c, 0, 1, 2, ..., K\}$, where c means the system is under recovery from a catastrophe, and i $(0 \le i \le K)$ means that the server is up and there are i customers in the system.

(2). The positive transition rates are:

$$q_{c,0} = \alpha,$$

$$q_{0,c} = \theta, \quad q_{0,1} = \lambda,$$

$$q_{i,i-1} = \mu, \quad q_{i,i+1} = \lambda, q_{i,c} = \theta, \quad 1 \le i \le K-1,$$

$$q_{K,K-1} = \mu, \quad q_{K,c} = \theta.$$

(3). The balance equations are:

$$\alpha p_c = \theta(p_0 + p_1 + ...p_K),$$

$$(\theta + \lambda)p_0 = \alpha p_c + \mu p_1,$$

$$(\theta + \lambda + \mu)p_i = \lambda p_{i-1} + \mu p_{i+1}, \quad 1 \le i \le K-1,$$

$$(\theta + \mu)p_K = \lambda p_{K-1}.$$

The normalizing equation:

$$p_c + p_0 + ... + p_K = 1.$$

(4). Using the normalizing equation in the first equation for p_c, we get

$$p_c = \theta/(\alpha+\theta).$$

Now try $p_i = r^i, 0 \le i \le K$ in the third equation to get

$$(\theta+\lambda+\mu)r^i = \lambda r^{i-1} + \mu r^{i+1}$$

which yields

$$\mu r^2 - (\theta+\lambda+\mu)r + \lambda = 0.$$

This has two solutions

$$r_1 = \frac{(\theta+\lambda+\mu) - \sqrt{(\theta+\lambda+\mu)^2 - 4\lambda\mu}}{2\mu},$$

$$r_2 = \frac{(\theta+\lambda+\mu) + \sqrt{(\theta+\lambda+\mu)^2 - 4\lambda\mu}}{2\mu}.$$

Hence the general solution is

$$p_i = Ar_1^i + Br_2^i, \quad 0 \le i \le K.$$

The unknowns A and B are obtained by using the equations for p_0 and p_K.

(5). Since the catastrophes arrive according to a Poisson process, they see the system in steady state (PASTA). Hence the the number of customers lost is given by

$$M = 0 * p_c + \sum_{i=0}^{K} i p_i.$$

If we restrict ourselves to catastrophes that arrive when the server is functional, the customers lost at such catastrophes is given by

$$M/(1-p_c) = (\alpha+\theta)M/\alpha.$$

P4. (10pts.) **Optimal Routing.** a queueing system consists of K servers, each with its own queue. Customers arrive at the system according to a PP(λ). A system controller routes an incoming customer to server k with probability α_k, where $\alpha_1 + \alpha_2 + \cdots + \alpha_K = 1$. Customers assigned to server k receive i.i.d. exp(μ_k) service times. Assume that $\mu_1+\mu_2+\cdots+\mu_K > \lambda$. It costs h_k dollars to hold a customer for one unit of time in queue k.

(a). What are the feasible values of α_k's so that the resulting system is stable?

(b). Compute the the expected holding cost per unit time as a function of the routing probabilities α_k $(1 \le k \le K)$ in the stable region.

(c). Compute the optimal routing probabilities α_k that minimize the holding cost per unit time for the entire system.

Ans.

(a). We must have $\lambda\alpha_k < \mu_k$. Hence the feasible region is

$$0 \le \alpha_k < \mu_k/\lambda, \;\; 1 \le k \le K,$$

$$\alpha_1 + ... + \alpha_K = 1.$$

(b). The long run cost per unit time for the entire system is given by

$$c(\alpha) = \sum_{k=1}^{K} h_k \frac{\lambda\alpha_k}{\mu_k - \lambda\alpha_k}.$$

(c). We need to minimize $c(\alpha)$ subject to $\alpha_k < \mu_k/\lambda$ for $1 \le k \le K$, and $\alpha_1 + ...\alpha_K = 1$. We ignore the inequality constraints, and use Lagrangian multipliers to solve the constrained optimization. We get the KK conditions as

$$\frac{\partial c(\alpha)}{\alpha_k} = \frac{\lambda h_k \mu_k}{(\mu_k - \lambda\alpha_k)^2} = a(constant), \;\; 1 \le k \le K.$$

This yields

$$\alpha_k = \sqrt{\frac{\mu_k}{\lambda}}(\sqrt{\frac{\mu_k}{\lambda}} - \sqrt{\frac{h_k}{a}}).$$

Thus the constant a is chosen to satisfy

$$\sum_{k=1}^{K} \sqrt{\frac{\mu_k}{\lambda}}(\sqrt{\frac{\mu_k}{\lambda}} - \sqrt{\frac{h_k}{a}}) = 1.$$

The solution is given by

$$a = \frac{\lambda}{\mu^2}(\sum_{k=1}^{K} \sqrt{h_k\mu_k})^2.$$

It can be seen that with this a the resulting α_k will be automatically feasible.