CS 286a | # HW 4a: Beyond M/M/* | **Guru: Adam Wierman**

Assigned: 11/10/08 | Due: 11/19/08, my office, 3pm

I encourage you to discuss these problems with others, but you need to write up the actual solutions alone. At the top of your homework sheet, please list all the people with whom you discussed. Crediting help from other classmates will not take away any credit from you. Also, please limit your use of the web as it is difficult to develop good problems, and so enough search will probably lead you to find the solutions online.

1 Warmup: z-transforms [10 points]

We'll need to use z-transforms in Wednesday's lecture, so I want you to practice with them a bit beforehand. The z-transform of a discrete function on the non-negative integers is defined as

$$G_p(z) = \sum_{i=1}^{\infty} p(i) z^i$$

When we write the z-transform of a random variable X having p.m.f. $p(i)$ I will often write $\widehat{X}(z) := G_p(z)$. (Similarly, for the Laplace transform of X I will often write $\widetilde{X}(s) := E[e^{-sX}]$.)

(a) Let A_t denote the number of arrivals by time t when arrivals occur according to a Poisson process with rate λ. Derive $\widehat{A_t}(z)$.

(b) Let A_X denote the number of arrivals during one service time X under a Poisson process with rate λ. Prove that $\widehat{A_X}(z) = \widetilde{X}(\lambda(1-z))$.

2 Moment matching with PH distributions [20 points]

As we discussed in class, one technique for analyzing systems with non-exponential job size or interarrival distributions is to use phase-type (PH) distributions. However, to do this one is faced with the task of fitting a PH distribution to the desired (often measured) distributions. One way in which this is done is to choose a PH distribution that matches the first few moments of the desired distribution. In this problem, you'll figure out how to match the first two moments using PH distributions. Throughout, suppose the desired mean is μ and variance is σ^2.

(a) Recall that an Erlang(n, μ) distribution has $C^2 = 1/n \leq 1$. Characterize the set of achievable C^2 can be attained by an MixedErlang(n, p, μ), where a MixedErlang distribution is with probability p_k the sum of k i.i.d. Exponential(μ) random variables for $k = 1, \ldots, n$.

(b) Prove that all hyperexponential distributions have $C^2 \geq 1$.

(c) Give a moment matching algorithm that uses hyperexponential (when $C^2 > 1$) and MixedErlang (when $C^2 \leq 1$) distributions to specify a PH distribution that has mean μ and variance σ^2 for arbitrary μ and σ^2. The algorithm should use only closed form calculations.

(d) Use your algorithm to fit a distribution to the cases of $\mu = 1$ and $C^2 = \pi, 1/\pi$.

3 More fun with PH distributions [20 points]

For this problem, we will consider discrete time PH distributions, i.e., distributions determined by the absorbtion time in a DTMC. This is simply to make the proofs a little bit easier, but if you would prefer to do the proofs for continuous time PH distributions feel free!

(a) Prove that if X and Y are independent PH distributions, with n and m phases respectively, then $Z = X + Y$ can be represented as an $n + m$ state PH distribution.

(b) Prove that if X and Y are independent PH distributions, with n and m phases respectively, then

$$Z = \begin{cases} X, & \text{with probability p;} \\ Y, & \text{with probability 1-p.} \end{cases}$$

can be represented as an $n + m$ state PH distribution.

(c) Prove that if X and Y are independent PH distributions, with n and m phases respectively, then $Z = \min(X, Y)$ can be represented as an nm state PH distribution. Note: you may find the Kronecker product $\otimes$ useful in this part. Please look up the definition of it if you haven't seen it before.

(d) Prove that any distribution with finite support on the natural numbers is a discrete PH distribution.

4 Beyond First Come First Served [20 points]

We have focused almost entirely on FCFS scheduling so far in the course. We will move to other scheduling policies soon, so here is an intro to some other ones.

(a) Under Processor Sharing (PS) the server is divided evenly among the jobs in the system. If there are 3 jobs in the system, each gets 1/3 of the total service rate. Derive $E[T]$ and $E[N]$ for the M/M/1/PS queue.

(b) Under Preemptive Last Come First Served (PLCFS) the server is always devoted entirely to the most recent arrival present. So, whenever a job arrives, the job at the server is interrupted and returned to the queue. (We will assume there is no overhead associated with this preemption.) Note that this is exactly like pushing and popping from a stack. Derive $E[T]$ and $E[N]$ for the M/M/1/PS queue.

(c) Comment on the similarity of the prior two analyses. How far can this similarity be pushed? Does it hold beyond the M/M/1?

(d) Now, let us consider a priority queue. Suppose there are two classes of jobs and that class 1 customers have absolute priority over class 2 customers. That is, a class 2 customer cannot receive service while a class 1 customer is in the system. (Again, assume there is no overhead associated with preemptions.) Class 1 and 2 customers both have Exponential(μ) job sizes and arrive according to Poisson processes with rates λ_1 and λ_2 respectively.

Derive the $E[T]$. Also derive and contrast $E[T_1]$ and $E[T_2]$, the mean response time for class 1 and class 2 customers respectively.

5 Non-poisson arrivals [25 points]

In class we illustrated how matrix analytic techniques allow generalization of the job size distribution. In this problem, you will apply the same ideas to the interarrival time distribution.

Consider a single server FCFS queue with Exponential(μ) job sizes where the arrival process is time varying. It is either a Poisson process with rate λ_1 or λ_2 and the rate switches after a period of time that is Exponential(α).

(a) Define the state space and draw the CTMC.

(b) Write out the generator matrix and define the $F, L, B, \ldots$ submatricies.

(c) Write the balance equations.

(d) Assume $\lambda_1 = 1$, $\lambda_2 = 2$, $\alpha = 0.5$, and $\mu = 3$ for the remainder of the problem. Solve for the R matrix iteratively using Matlab (or some other program). Please hand in your (commented) Matlab code.

(e) Solve for the limiting probabilities.

(f) Compute $E[T]$ and $E[N]$.

6 When can R be found without Matlab? [15 points]

Though solving for the R matrix is often a numerical task, sometimes it can be done analytically. In this problem, you'll determine one example when this is the case.

Suppose we have the same QBD structure that we studied in class, but that $B = v \cdot \alpha$ where

$$v = \begin{pmatrix} v_0 \\ \vdots \\ v_m \end{pmatrix}, \qquad \alpha = (\alpha_0, \ldots, \alpha_m), and \qquad \alpha e = 1$$

where e is a vector of ones. What this all means is that the rows of B are equivalent up to scaling. Prove that in this case

$$R = -F(L + Fe\alpha)^{-1}.$$

Use this result to verify your answer to problem 4.

Extra Credit: Prove that the R matrix can be found explicitly when the equivalent condition holds for F, i.e., $F = w \cdot \beta$ where

$$w = \begin{pmatrix} w_0 \\ \vdots \\ w_m \end{pmatrix}, \qquad \beta = (\beta_0, \ldots, \beta_m), and \qquad \beta e = 1$$

Note – this is harder!