Introduction to Queuing Networks

Solutions to Problem Sheet 4

---

1. Following the hint, we define the state $X(t)$ to be the number of waiting customors minus the number of waiting taxis at time t. Since customers and taxis arrive according to independent Poisson processes, therefore $X(t)$ is a continuous-time Markov chain. Its state space is $\{-K, -K+1, \dots, 0, 1, 2 \dots\}$ as the maximum number of waiting taxis is K. The transition rates of $X(t)$ are given by $q_{n,n+1} = \lambda$ and $q_{n,n-1} = \mu 1(n > -K)$, with all other transition rates being zero.

   The Markov chain $X(t)$ is a birth-death chain and hence is reversible. The invariant distribution π can be found by solving the detailed balance equations,

$$\mu \pi_{n+1} = \lambda \pi_n, \quad n \geq -K,$$

   which imply that $\pi_n = \pi_{-K} \rho^{n+K}$, where $\rho = \lambda/\mu$. Combining this with the normalisation condition $\sum_{n=-K}^{\infty} \pi_n = 1$, we get

$$\pi_n = (1-\rho)\rho^{n+K}, \quad n \geq -K. \tag{1}$$

   Note that the existence of an invariant distribution relies on the stability assumption that $\lambda < \mu$.

   (a) The mean number of waiting customers is given by

$$E[N_C] = \sum_{n=0}^{\infty} n\pi_n = (1-\rho)\rho^K \sum_{n=0}^{\infty} n\rho^n = \frac{\rho^{K+1}}{1-\rho}.$$

   We have used the hint to obtain the last inequality.

   Now, by Little's law, since the arrival rate of customers is λ, the mean waiting time for a customer is

$$E[W_C] = \frac{E[N_C]}{\lambda} = \frac{\rho^K}{\mu - \lambda}.$$

   (b) The mean number of waiting taxis is given by

$$E[N_T] = \sum_{n=0}^{K} n\pi_{-n} = (1-\rho)\rho^K \sum_{n=0}^{K} n\rho^{-n} = \frac{K - (K+1)\rho^K + \rho^{K+1}}{1-\rho}.$$

   We have used the hint to obtain the last inequality.

   Now, by Little's law, since the arrival rate of customers is μ, the mean waiting time for a customer is

$$E[W_T] = \frac{E[N_T]}{\mu} = \frac{K - (K+1)\rho^K + \rho^{K+1}}{\mu - \lambda}.$$

   (c) As K tends to infinity, the mean waiting time of a customer decreases to zero, while the mean waiting time of a taxi increases to infinity. In the limit, the queue for taxis is unstable (since $\mu > \lambda$) and builds up to infinity. Hence, eventually there are always taxis waiting and customers never have to wait, whereas taxis have to wait forever (on average).

2. (a) The states are $\{0, 1, 2, \dots\}$. The possible transitions from any state $n > 0$ are to $n+1$ and $n-1$, while the only possible transition from 0 is to 1. The transition rates are $q_{n,n+1} = \lambda$ and $q_{n,n-1} = \frac{\mu(n+1)}{n} 1(n \geq 1)$.

(b) Let π denote the invariant distribution, if one exists. The Markov process N_t is a birth-death process, so it is reversible if it has an invariant distribution, which can then be computed by solving the detailed (local) balance equations. These equations say

$$\lambda \pi_n = \frac{(n+2)\mu}{n+1} \pi_{n+1} \text{ i.e., } \pi_{n+1} = \frac{n+1}{n+2} \rho \pi_n, \quad , n = 0, 1, 2, \dots,$$

where $\rho = \lambda/\mu$. Iterating this yields a telescoping product, so that

$$\pi_n = \frac{\rho^n}{n+1} \pi_0, \quad n = 0, 1, 2, \dots$$

We need to solve for π_0, which we do by noting that a probability distribution has to sum to 1. The sum of π_n converges only if $\rho < 1$, i.e., $\lambda < \mu$, so it is only in this case that the Markov process has an invariant distribution. Otherwise, the queue is unstable and blows up over time. If $\rho < 1$, then using the hint, we get

$$\frac{-\log(1-\rho)}{\rho} \pi_0 = 1.$$

Substituting this in the expression for π_n, the invariant distribution is given by

$$\pi_n = \frac{1}{-\log(1-\rho)} \frac{\rho^{n+1}}{n+1}, \quad n = 0, 1, 2, \dots$$

(c) The generating function for the invariant queue length distribution is given by

$$G(z) = \sum_{n=0}^{\infty} \pi_n z^n = \frac{1}{-z \log(1-\rho)} \sum_{n=0}^{\infty} \frac{(\rho z)^{n+1}}{n+1} = \frac{\log(1-\rho z)}{z \log(1-\rho)},$$

for all $|z| < 1/\rho$. In particular, $z = 1$ is inside this radius of convergence, so we can differentiate at this point to obtain the mean queue length. We get

$$E[Q] = G'(1) = \frac{1}{\log(1-\rho)} \Big(\frac{-\rho z}{(1-\rho z) z^2} - \frac{\log(1-\rho z)}{z^2} \Big) \Big|_{z=1} = \frac{-\rho}{(1-\rho)\log(1-\rho)} - 1.$$

By Little's law, the mean sojourn time is given by $E[W] = E[Q]/\lambda$.

3. (a) Let $X(t)$ denote the number of customers in the restaurant at time t. Then $X(t)$ is a birth-death Markov chain with birth rate $\frac{n+1}{n+2}\lambda$ and death rate $n\mu$ in state n. (It is a modification of an $M/M/\infty$ queue, with state-dependent arrival rate.)

The invariant distribution π, if there is one, satisfies the detailed balance equations

$$\frac{(n+1)\lambda}{n+2} \pi_n = (n+1)\mu \pi_{n+1}, \quad n \geq 0,$$

which imply that $\pi_n = \frac{1}{(n+1)!} \rho^n \pi_0$ for $n \geq 0$, where $\rho = \lambda/\mu$. Since π_n is a probability distribution, it must sum to 1. This yields that

$$\frac{1}{\pi_0} = \sum_{n=0}^{\infty} \frac{\rho^n}{(n+1)!} = \frac{1}{\rho} \sum_{m=1}^{\infty} \frac{\rho^m}{m!} = \rho^{-1}(e^\rho - 1).$$

Thus, the invariant distribution is given by

$$\pi_n = \frac{\rho}{e^\rho - 1}\frac{\rho^n}{(n+1)!} = \frac{1}{e^\rho - 1}\frac{\rho^{n+1}}{(n+1)!}, \quad n \geq 0.$$

This holds for all λ and μ and doesn't require a stability condition.

(b) Let X_A denote the (random) number of customers seen by an arrival at time t who decides to enter the restaurant. We assume that the system is in equilibrium. By Bayes' theorem, we have

$$\begin{aligned} P(X_A = n) &= P(X(t) = n | \text{arrival at time } t \text{ decides to enter}) \\ &= \frac{P(X(t) = n \text{ and arrival at } t \text{ decides to enter})}{P(\text{arrival at time } t \text{ decides to enter})} \\ &= \frac{\pi_n P(\text{arrival at } t \text{ enters given there are } n \text{ customers})}{\sum_{k=0}^\infty \pi_k P(\text{arrival at } t \text{ enters given there are } k \text{ customers})} \\ &= \frac{\rho^{n+1}/(n+2)!}{\sum_{k=0}^\infty \rho^{k+1}/(k+2)!} \\ &= \frac{\rho^{n+2}}{(n+2)!}\frac{1}{\sum_{k=0}^\infty \rho^{k+2}/(k+2)!} \\ &= \frac{\rho^{n+2}}{(n+2)!}\frac{1}{e^\rho - 1 - \rho}. \end{aligned}$$

Note that this isn't exactly the same as π_n.

4. (a) The mean service time is obtained as

$$E[S] = \int_0^\infty x f_S(x)dx = \int_0^\infty \frac{3x}{(1+x)^4}dx = \int_0^\infty \Big[\frac{3}{(1+x)^3} - \frac{3}{(1+x)^4}\Big]dx = \frac{1}{2}.$$

Hence, $\lambda = 2$ is the largest arrival rate for which the system can be stable, no matter what service discipline is used.

(b) First, we compute the second moment of the service time distribution. We have

$$\begin{aligned} E[S^2] &= \int_0^\infty x^2 f_S(x)dx = \int_0^\infty \frac{3x^2}{(1+x)^4}dx \\ &= \int_0^\infty \Big[\frac{3}{(1+x)^4} - \frac{6}{(1+x)^3} + \frac{3}{(1+x)^2}\Big]dx = 1. \end{aligned}$$

Hence, by the Pollaczek-Khinchin formula, the mean sojourn time is given by

$$E[W] = E[S] + \frac{\lambda E[S^2]}{2(1-\rho)} = 0.5 + \frac{0.5}{2(1-0.25)} = \frac{5}{6}.$$

Hence, by Little's law, the mean number in system is given by

$$E[N] = \lambda E[W] = \frac{5}{12}.$$

(c) If the server employs the PS (processor-sharing) discipline, then the invariant distribution is the same as for the $M/M/1$ queue, namely

$$\pi_n = (1-\rho)\rho^n, \quad n = 0, 1, 2, \ldots,$$

where $\rho = \lambda E[S] = 0.25$. Hence, the mean number in system is given by

$$E[N] = \frac{\rho}{1-\rho} = \frac{1}{3},$$

and, by Little's law, the mean sojourn time is $E[W] = E[N]/\lambda = \frac{2}{3}$.

(d) Comparing the answers to parts (b) and (c), it is clear that the PS policy minimises mean delay (sojourn time) and mean queue length. Hence, that is the policy I would recommend.

If the service time is deterministically equal to the same mean, 0.5, then $E[S^2] = 0.5^2 = 0.25$. Hence, by the Pollaczek-Khinchin formula, the mean sojourn time is $E[W] = 7/12$ and the mean number in system is $E[N] = 7/24$. The corresponding numbers for the PS queue are the same as before, namely $2/3$ and $1/3$. Hence, FCFS is now the better policy.

The intuition for why the mean delay under FCFS increases with increasing variability of the service time is that, if there is a large job, there are proportionately more small jobs (in proportion to its size) waiting behind it, all of whose sojourn times are correspondingly increased. If the variance gets too large (specifically, larger than that of an exponential with the same mean), then it becomes beneficial to use processor sharing, under which jobs don't interfere with each other.

5. Let us model the cafetaria as an $M/G/\infty$ queue, as the arrivals are Poisson, and all customers entering the system are "served" in parallel. (Here, service corresponds to the time the customer spends eating lunch.) The approximation involved is that the number of servers is not infinite, but is only equal to the number of dining spaces. If this number is K, then the exact model is $M/G/K/K$ (with K servers and total waiting room of size K) rather than $M/G/\infty$. However, while we can analyse an $M/M/K/K$ queue, we have no results for an $M/G/K/K$ queue, which prompts the approximation.

For the $M/G/\infty$ queue, we know that the invariant distribution is the same as for an $M/M/\infty$ queue with the same arrival rate and mean service time. In our example, the arrival rate is $\lambda = 1$ per minute, and the mean service time is 25 minutes, so the service rate is $\mu = 1/25$ per minute. Hence, $\rho = \lambda/\mu = 25$. Hence, the invariant distribution is Poisson(25).

Since the arrival process is Poisson, then by the **PASTA property**, the distribution of the number in system seen by an arrival is the same as the invariant distribution, Poisson(25). Suppose there are K places in the cafetaria. We will approximate the probability that aa arriving customer sees the cafetaria full (in the actual $M/G/K/K$ system) by the probability that an arrival sees K or more customers in system, in the $M/G/\infty$ approximating system. Letting N denote the random number of customers in the system as seen by an arrival, we want to compute $P(N \geq K)$.

We know that N has a Poisson(25) distribution, which has mean and variance equal to 25. Following the hint, we will approximate the Poisson distribution by a normal distribution with the same mean and variance, i.e., $\mu = 25$ and $\sigma^2 = 25$. Then, $(N-\mu)/\sigma$ has a standard normal distribution. Hence, by the hint,

$$P\Big(\frac{N-\mu}{\sigma} > 2.23\Big) < 0.01 < P\Big(\frac{N-\mu}{\sigma} > 2.22\Big).$$

This suggests that we should take $N = 37$ in order to ensure that fewer than 1% of customers are turned away.

There are several approximations involved in the analysis above. This is common in applications, where the system rarely fits a standard model exactly, and you have to exercise judgement as to how you want to approximate it by some known model.