Introduction to Queuing Networks

Problem Sheet 6

---

1. Customers come into a bookshop according to a Poisson process with a rate of 1 per minute. (All customers enter individually, never as part of a group.) Each customer browses for a period uniformly distributed between 5 and 25 minutes, independent of other customers and of the arrival process. There are enough copies of each book that all customers can browse in parallel. At the end of the browsing period, each customer has a 0.3 chance of deciding to make a purchase, independent of the length of the browsing period, other customer decisions and the arrival process. The other 70% of customers leave without making a purchase.

   The check-out is manned by a single person, who takes an exponentially distributed time with a mean of 2 minutes to serve each customer who is making a purchase. Each customer who makes a purchase leaves the shop with probability 0.9, but decides to continue browsing with probability 0.1. If they decide to continue to browse, their chance of making a further purchase and all their future decisions are identical to those of a new customer.

   (a) Draw a figure modelling the system described above as a network of queues. For each node in the network, specify what type of queue it is, the external arrival rate γ_j into that node and the service rate μ_j at that node. Specify all the routing probabilities in the network, both for pairs of nodes i, j, and also from each node i to the outside. Solve the traffic equations to find the total arrival rate λ_j at each node.

   (b) What is the mean length of time spent in the bookshop by a customer on any one visit?

   (c) The shop provides 20 armchairs. Every customer who is browsing will use a chair if one becomes available. What is the probability that a new customer entering the shop finds no free chairs? Justify your answer carefully.
   (Hint: If X is a Poisson(θ) random variable, then, for $n > \theta$, you may use the approximation,
   $$P(X \geq n) \approx \frac{n+1}{n+1-\theta} \frac{1}{\sqrt{2\pi n}} \exp\Big(n - \theta - n \log_e \frac{n}{\theta}\Big),$$
   which can be obtained using Stirling's formula.)

2. The handing out of first-time prison sentences (i.e. to individuals never previously sentenced to prison) in the UK follows a Poisson process of rate 4000 per month. The average length of a prison sentence is 9 months but its distribution is unknown. The recidivism rate for released prisoners, defined as the probability that they commit new crimes for which they are convicted and sentenced to prison, is 60%. The time between release and re-sentencing, for recidivists, is exponentially distributed with a mean of 2 years.

We make the simplifying assumption that the probability of recidivism is a fixed number, and that the time before recidivism is a random variable with fixed distribution. Specifically, their values don't depend on the past history of the prisoner (the number and duration of previous prison sentences). We also assume that the length of sentences handed out to recidivists has the same distribution as for first-time offenders, irrespective of the number of previous convictions.

(a) Draw a figure modelling the system described above as a network of queues. For each node in the network, specify what type of queue it is, the external arrival rate γ_j into that node and the service rate μ_j at that node. Specify all the routing probabilities in the network, both for pairs of nodes i, j, and also from each node i to the outside.

(b) Solve the traffic equations to find the total arrival rate λ_j at each node.

(c) What is the mean number of people in prison in equilibrium?

(d) The use of a normal approximation suggests that enough prison places should be provided for the mean plus three standard deviations of the random variable describing the prison population. How many prison places does the UK need?

3. Requests arrive at a Web server according to a Poisson process of rate 5 per second, and are always for a single item. If the requested item is already in the cache, it is served from there. Otherwise, it is first read from the disk to the cache and then served from the cache. There is a 20% chance that any given request can be served from cache, independent of all previous requests.

The time it takes to read an item from disk to cache is dominated by the disk access time. So we will ignore any dependence on the size of the item, and model the time to read each item as an exponential random variable with a mean of 200 milliseconds. The disk schedules requests in order of arrival, i.e., it is a FIFO (first-in-first-out) single server queue.

The service time for a single item from cache is given by the ratio of the item file size to the link bandwidth. We assume that the service time requirements $T_1, T_2, \ldots$ of distinct items are i.i.d. (independent and identically distributed) with a Pareto distribution whose density is given by

$$f_T(t) = \frac{0.02}{\left(1 + \frac{t}{100}\right)^3} 1(t \geq 0).$$

The unit of time used for specifying the density is milliseconds.

(a) Suppose that the cache uses the PS (processor sharing) service discipline. Model the progress of a request through the web server, including both disk and cache, by a queueing network. To do this, draw a figure specifying all the network nodes, the external arrival rates to each node, the service rate, number of servers and service discipline at each node, and the routing probabilities between every pair of nodes as well as from a node to outside the network.

(b) Compute the mean time between the arrival of a request at the web server and its service completion.

(c) It is suggested that the service discipline at the cache be changed from PS to FIFO. Compute the mean sojourn time of customers under the FIFO discipline and say whether it is preferable to PS.

4. Items brought into a dry-cleaners have to undergo pre-processing for a random time with mean 5 minutes and standard deviation 3 minutes. The pre-processing is done by a single server on a last-in first-out (LIFO) basis. After pre-processing, there is a choice of one of two treatments. Treatment 1 is carried out manually by a single server on a FIFO basis and takes an exponentially distributed time with mean 20 minutes. At the end of this treatment, the item is always clean and ready to leave the system. Treatment 2 is carried out by a machine with effectively unlimited capacity, and takes a deterministic period of 60 minutes. At the end of this treatment, the item is ready to leave with probability 0.9; with probability 0.1, it is not satisfactorily cleaned, and has to restart the process.

   Assume that items are brought into the dry-cleaners singly, at times which constitute a Poisson process of rate γ. Service times are mutually independent, and independent of the arrival process. Suppose that, after pre-processing, each item is assigned to Treatment 1 with probability α and to Treatment 2 with probability $1 - \alpha$, independent of its history, and of all other items and their histories.

   (a) Find the stable region of this queueing system, i.e., for each α, find the largest value of γ for which the system is stable. If we denote this by $\gamma(\alpha)$, then for what value of α is this largest?

   (b) For this value of α, and for $\gamma = \gamma(\alpha)/2$, find the mean length of time needed to complete the processing of a single item.

5. The application process for a Vogon driving licence involves N processing stages, to be completed in sequence. Each stage of processing is carried out by a large department, sufficiently well staffed to process all ongoing applications in parallel; it takes a week on average to deal with each application.

   At the completion of any processing stage between 1 and $N-1$, but not N, an error is found in the application with probability $1/2$, independent of how many stages the application has already gone through. If this happens, it is necessary to restart the application process from the first stage. Faced with this prospect, an applicant gives up with probability $1/2$, independent of past choices or the amount of time already spent in the process, and quits the system. With residual probability $1/2$, the applicant returns to the start of the process.

   We assume that new applicants for driving licences enter the system according to a Poisson process of rate η (in units of per week). Applicants who are fortunate enough to go through all N stages without an error being discovered get a licence and leave the system.

   (a) Draw a diagram describing the application process as a queueing network, marking the external arrival rates into all nodes, and all the non-zero routing probabilities between nodes, or for leaving the system. Also state the type of each queueing node (i.e., number of servers, service time distribution and service discipline).

   (b) Solve the traffic equations to determine the total arrival rate at each of the nodes in terms of the external arrival rate η into the first node.

   (c) What is the mean length of time spent in the process by an applicant, whether successful or not? If there is a condition on η required for this time to be finite, then state that condition.

   (d) What fraction of applicants succeed in getting a licence?