

Math 365 Queues

Submit your R script to tleise@amherst.edu by the Wednesday after spring break, and turn your handwritten part in by start of class on that day.

The objective of this lab is to further study queueing as a type of birth and death process using a combination of simulations and theoretical analysis.

1 Simulation of M/M/1 Queue

In general for birth and death queueing processes, we let λ_n denote the arrival rate when there are currently n customers, and μ_n denote the departure rate, where both arrivals and departures are treated as exponential processes. That is, if there are currently n customers, the time until the next event is an exponential random variable with rate $\lambda_n + \mu_n$; the probability that the next event is an arrival is $\frac{\lambda_n}{\lambda_n + \mu_n}$ and the probability that the next event is a departure is $\frac{\mu_n}{\lambda_n + \mu_n}$.

The M/M/1 queue has constant rates $\lambda_n = \lambda$ for $n \geq 0$ and $\mu_n = \mu$ for $n \geq 1$ ($\mu_0 = 0$).

We'll start by running some simulations of an M/M/1 queue to see some examples of what such processes look like for the positive recurrent case ($\lambda < \mu$) in which we expect to repeatedly return to an empty queue. Feel free to vary the values of λ and μ in the simulation. If it seems to be running too long, halt it and try a different pair of values (make sure $\lambda < \mu$ so you don't get exploding queues).

First initialize some variables:

```
# M/M/1 queue simulator
lambda <- 1 # arrival rate
mu <- 4 # service rate
duration <- 10000 # total duration of the simulation
t <- 0 # current time in the simulation
queue <- 0 # start with empty queue
s <- 0 # running sum for computing average queue length
```

We assume the queue is initially empty, so generate a first arrival to the queue with rate λ :

```
# first arrival to start process
T1 <- rexp(1,rate=lambda)
currentqueue <- 1
eventsTime <- T1
t <- T1
nEvents <- 1 # total number of events that have occurred
```

The next event can be either an arrival or departure. The time $T1$ until that next event is an exponential random variable with rate $\lambda + \mu$. To determine whether that event is an arrival or a departure, generate a random number p between 0 and 1; if $p < \frac{\lambda}{\lambda + \mu}$, then the event is an arrival, otherwise it's a departure.

```

while (t<duration) {

  nEvents <- nEvents+1
  if(currentqueue>0) { # nonempty queue
    T1 <- rexp(1,rate=lambda+mu) # time until next event
    # is event an arrival or departure?
    p <- runif(1,0,1)
    queue[nEvents] <- currentqueue # how many have been in queue before this new event
    currentqueue <- ifelse(p<lambda/(lambda+mu),
      currentqueue+1, # arrival
      currentqueue-1) # departure

  } else { # empty queue
    T1 <- rexp(1,rate=lambda)
    queue[nEvents] <- currentqueue
    currentqueue <- 1
  }

  t <- t+T1 # time of next arrival
  eventsTime[nEvents] <- T1 # inter-event time
  s <- s+T1*queue[nEvents] # spent T1 at nth queue length
}

```

Note that the R code above also has to check at each step whether the queue is empty, in which case only an arrival is possible. Once the simulation is complete, you can plot the queue length over time:

```

plot(cumsum(eventsTime),queue,type="s", xlab="Time",ylab="Queue length",
main=paste("M/M/1 Simulation, mean queue length=",s/t))

```

2 Analysis of M/M/1 Queue

What do we expect the average queue length to be? What proportion of the time do we expect the queue to be empty over the long run? What is the average time customers will spend in the queue? We can do some straightforward analysis to answer these questions. Begin by calculating the limiting probability distribution $\pi(n)$ for $n \geq 0$.

Exercise 1 For the M/M/1 queue, $\pi(n)$ satisfies the equilibrium conditions

$$(\lambda + \mu)\pi(n) = \lambda\pi(n - 1) + \mu\pi(n + 1) \text{ for } n \geq 1 \text{ and } \lambda\pi(0) = \mu\pi(1).$$

First solve for $\pi(1)$, then for $\pi(2)$, and then recognize the pattern that is occurring to find the general solution in terms of $\pi(0)$. Enforce $\sum_{n=0}^{\infty} \pi(n) = 1$ to find $\pi(0)$ (use the geometric series formula to simplify). State the resulting formula for $\pi(n)$, which gives the long-run proportion of time that the queue length equals n (equivalently, the probability at any given time t that the queue length is n).

Exercise 2 Run a simulation and check that the observed proportion of time in which the queue

length is 0 agrees well with the theoretical value for $\pi(0)$ found in Exercise 1. An easy way to compute this value in R is the following:

```
sum(eventsTime[which(queue==0)])/t
```

Exercise 3 Compare the observed proportion of time in which the queue length is 1 with the theoretical value for $\pi(1)$.

Exercise 4 Let L be the expected value of the queue length (average number of customers):

$$L = \sum_{n=0}^{\infty} n\pi(n).$$

Find a simple formula for L by substituting in your solution for $\pi(n)$ from Exercise 1 and simplifying using the formula $\sum_{n=1}^{\infty} nr^n = \frac{r}{(1-r)^2}$ (which can be obtained by taking a derivative of the geometric series formula).

Exercise 5 Compare the observed average queue length (`s/t` in R) with the theoretical value for L .

Exercise 6 Finally, the average time W each customer spends in line is the average queue length divided by the rate of arrivals: $W = L/\lambda$. This relation is known as “Little’s law” (http://en.wikipedia.org/wiki/Little's_law). State a formula for W using the result of Exercise 4.