Queueing models, branching processes, strong Markov property

Monday, November 10, 2008
11:57 AM

HW 3 posted, due noon on Thursday, November 20.
Office hours cancelled for Wednesday, November 12.

We will show how to use branching processes to answer two questions in discrete-time queuing theory:

1. During a consecutively busy period, how many demands are serviced until the server is free?
2. How many epochs must one wait before the server is free?

References: Feller, An Introduction to Probability Theory and its Applications I, Ch. XII.5

Karlin and Taylor, Sec. 3.5

- address similar questions but not explicitly in terms of branching processes; somewhat opaque

For the first question, all we really need to know is the probability distribution for the number of demands that arrive during a service period.

$$p_j = P(j \text{ arrivals during a service period})$$

This is similar to when we set up a discrete-time Markov chain model for a queue with epochs corresponding to the end of a service period, but here the probability distribution for the number of arrivals is slightly different because we do allow here for there to be no arrivals during a service period. Notice by setting up things the way that we are, we don't have to keep track separately of the details of the length of a service period and/or the rate at which demand arrives; just the single statistic described above will be the crucial one. Actually this calculation doesn't even require the assumption that time is discrete.

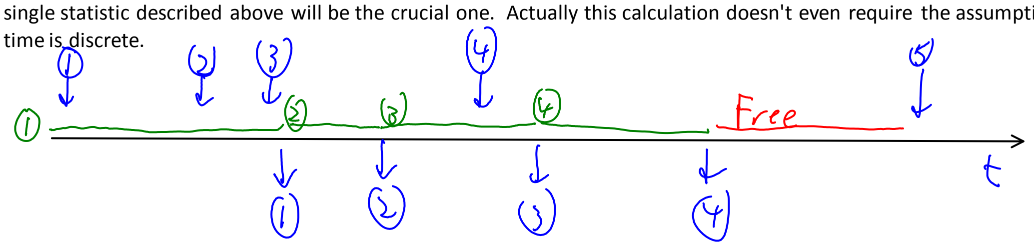

We can, without regard to whether we want to treat the arrivals and service times as continuous or discretely distributed, represent the queue as a branching process. For this question, we will consider each of the demands as an agent in the branching process, and the offspring associated to an agent will be the new demands which arrive while that agent is being served.

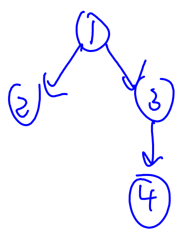

The branching process terminates once the server is free.

Therefore the question about how many demands must be satisfied before the queue becomes free is equivalent to the question of what is the total progeny of a branching process, meaning the total number of nodes generated by the branching process before it terminates. Let's set up the means to calculate this.

Branching process X_n = # agents at epoch n

$$X_{n+1} = \sum_{j=1}^{X_n} X_{n,j}$$

$\{X_{n,j}\}$ = # offspring of agent #j at epoch n

iid rvs

$$P(\xi = j) = p_j$$

Total progeny $\quad Z = \sum_{n=0}^{\infty} X_n \quad$ (could be ∞)

To set up the calculation for the probability distribution for the total progeny, we first set up a simpler calculation which is the total progeny generated up until epoch m.

$$Z_m = \sum_{n=0}^{m} X_n$$

This is a sum of dependent random variables, but they have a recursive structure because of the branching process and probability generating functions work very well under recursion.

$$\mathcal{P}_{Z_m}(s) = \mathbb{E}\left(s^{Z_m} \mid X_0 = 1\right)$$

$$\text{*)} \quad \mathbb{E}\left(s^{Z_m} \mid X_0 = x_0\right) = \left(\mathcal{P}_{Z_m}(s)\right)^{x_0}$$

because each ancestral node generates an independent number of descendants, so the probability generating function for their sum is just the product of the generating function associated to each single ancestor.

As is typical in stochastic analysis, especially with recursive structure, we'll derive an equation for these probability generating functions by first step analysis

$$\mathcal{P}_{Z_m}(s) = \mathbb{E}\left[\mathbb{E}\left(s^{Z_m} \mid X_1, X_0 = 1\right) \Big| X_0 = 1\right]$$

Law of total expectation

(alternatively just sum over the possible values of X_1 as in the more concrete representation of law of total expectation)

$$= \mathbb{E}\left[\mathbb{E}\left(s^{X_0 + \sum_{n=1}^{m} X_n} \Big| X_1\right) \Big| X_0 = 1\right]$$

$$= \mathbb{E}\left(\mathbb{E}\left(s^{1 + \sum_{n=1}^{m} X_n} \Big| X_1\right) \Big| X_0 = 1\right]$$

$$= s\,\mathbb{E}\left(\mathbb{E}\left(s^{\sum_{n=1}^{m} X_n} \Big| X_1, X_0 = 1\right) \Big| X_0 = 1\right]$$

$\downarrow$ Markov property

$$= s\,\mathbb{E}\left(\mathbb{E}\left(s^{\sum_{n=1}^{m} X_n} \Big| X_1\right) \Big| X_0 = 1\right)$$

$$= s\, \mathbb{E}\left(\mathbb{E}\left(s^{\sum_{n=1}^{m} X_n} \,\middle|\, X_1 \right) \,\middle|\, X_0 = 1 \right)$$

$$= s\, \mathbb{E}\left(\left(P_{Z_{m-1}}(s) \right)^{X_1} \,\middle|\, X_0 = 1 \right)$$

time shift $(1, m) \to (0, m-1)$

and used

$$= s\, \mathbb{E}\left(\left(P_{Z_{m-1}}(s) \right)^{X_{0,1}} \,\middle|\, X_0 = 1 \right)$$

$$= s\, \mathbb{E}\left(\left(P_{Z_{m-1}}(s) \right)^{X_{0,1}} \right)$$

$$P_{Z_m}(s) = s\, P_X\left(P_{Z_{m-1}}(s) \right)$$

$P_{Z_m}(s)$ ↓ p.g.f for # total descendants through m generations

P_X ↓ p.g.f. for # offspring of one agent

$$P_X(s) = \sum_{j=0}^{\infty} p_j s^j$$

$P_X(s)$ ↓ specified by model

$$Z_0 = 1 \Rightarrow P_{Z_0}(s) = 1$$

and then use this relationship to recursively calculated probability generating function (and therefore probability distribution) for any $Z_m,\ m \geq 0$.

What if we want to extend indefinitely to all descendants of a given ancestor $m \to \infty$?

$$\lim_{m \to \infty} Z_m = Z$$

Then $P_Z(s) = \mathbb{E}\, s^Z = \lim_{m \to \infty} P_{Z_m}(s)$

Here we are using the Continuity Theorem which is discussed in Resnick Sec. 1.5.

We just have to be careful to note that Z can be a defective random variable, meaning that

$$\sum_{i=0}^{\infty} P(Z = i) < 1$$

We just have to be careful to note that Z can be a defective random variable, meaning that

$$\sum_{j=0}^{\infty} P(Z=j) < 1$$

sum over all finite values

Reason is $P(Z=\infty)$ could be positive.

In the probability generating function, defective random variables are no problem because we just treat

$$s^{\infty} = 0 \quad \text{since} \quad 0 \le s < 1$$

Taking limits in our recursion relation, we get

$$\underbrace{P_Z(s)}_{\text{solve for}} = s \underbrace{P_Y}_{\text{given}} \left(P_Z(s) \right)$$

This is like a family of equations for each value of s. See Feller for proof that this has unique solution. We solve this nonlinear equation to obtain the probability generating function $P_Z(s)$ for the total progeny and we can calculate statistics from it in the usual way.

Some simple statistics can actually be computed directly from the recursion equation.

$$P_Z(1) = \sum_{j=0}^{\infty} 1^j \, P(Z=j) = \sum_{j=0}^{\infty} P(Z=j) = P(Z<\infty)$$

$$= P(\text{extinction}) = a$$

From a previous lecture, we showed that this extinction probability satisfies the equation

$$a = P_Y(a)$$

And this agrees with the new recursion equation we derived.

What is the mean number of total progeny?

$$\langle Z \rangle = s \frac{d}{ds} P_Z(s) \Big|_{s=1} = P_Z'(1)$$

Differentiate recursion relation:

$$P_Z'(s) = P_Y(P_Z(s)) + s P_Y'(P_Z(s)) P_Z'(s)$$

Evaluate at $s = 1$:

$$\underbrace{P_Z'(1)}_{\langle Z \rangle} = P_Y(P_Z(1)) + P_Y'(P_Z(1))\underbrace{P_Z'(1)}_{\langle Z \rangle}$$

Clearly if the probability for the number of progeny to be infinite is nonzero, then the mean number of total progeny will be infinite. So let's just focus on the case when extinction is certain

$$a = 1 \iff \langle Y \rangle \leq 1$$

$$\Rightarrow q = P_Z(1) = 1$$

$$P_Y'(1) = \langle Y \rangle$$

$$P_Y(1) = 1$$

$$\langle Z \rangle = 1 + \langle Y \rangle \langle Z \rangle$$

$$\langle Z \rangle = \frac{1}{1 - \langle Y \rangle} \quad \text{for } \langle Y \rangle \leq 1$$

mean # total progeny

mean # offspring per agent

Now we turn to the second queuing question which is, how much time does a busy period for the server last?

It's tempting to proceed as follows:

$W =$ total busy period

$T_n =$ service time for nth demand

$$W = \sum_{n=1}^{Z} T_n$$

independent, identically distributed r.v.'s

This looks like a random sum of independent random variables, but Z is not independent of the service times since shorter service times are correlated with higher probability to terminate the branching process. We'll see once we discuss martingales that we can deal with some random sums of this type even though the number of terms in the sum is not independent of the summands.

So we'll go about this a different way. For this purpose, let's set up discrete time epochs

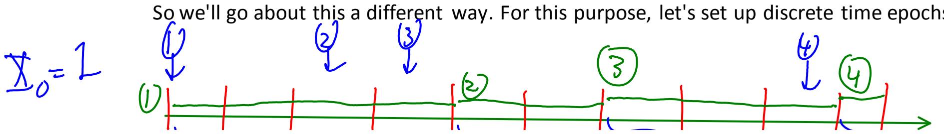

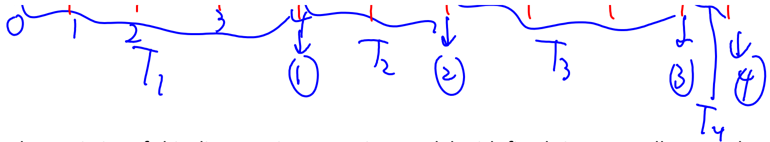

The statistics of this discrete-time queuing model with fixed-time, equally spaced epochs will be modeled as follows:

- one arrival occurs during any epoch with probability p
- the length of a service time is given by probability distribution $P(T=j) = p_{T,j}$

For a typical Markov chain model, we would have a probability q for a service to be completed at any given epoch, which would give a geometric distribution for the service time:

$$p_{T,j} = q(1-q)^{j-1} \qquad j = 1, 2, \ldots$$

but we don't need this assumption for the branching process calculation. That is, the queueing model itself need not have Markovian behavior with regard to its departures!

The way this gets mapped into a branching process model is by considering time epochs of the discrete-time queue model to be the agents of the branching process. An epoch has no descendants if no demands arrives during it. However, if a demand does arrive during an epoch, then the epochs during which that demand is serviced are considered the offpsring of the epoch during which the demand arrived.

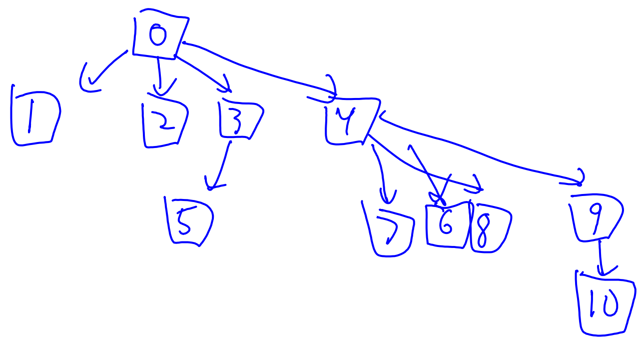

The total waiting time until the server is free is just the total progeny of this branching process. To calculate its statistics, we just need to use the branching process recursion formulas together with the probability distribution for the number of offspring of an epoch.

ξ = # of offspring of 1 agent

$$P(\xi = 0) = 1 - p$$

$$P(\xi = j) = P(\text{demand arrives, requiring } j \text{ units of service time})$$

$$= p \, p_{T,j} \quad \text{for } j \geq 1$$

From this compute $\mathcal{P}_\xi(s) \Rightarrow \mathcal{P}_Z(s)$

which gives the probability generating function for the total number of progeny in the branching process, which is the number of epochs over which the server remains busy until the first moment it is free.

We return now to thinking about continuous-time Markov chains. Last time, we discussed the numerical simulation of the Poisson counting process, and we want to move to discussing the numerical simulation of general continuous-time Markov chains. We need a preparatory concept:

Markov times and the Strong Markov Property

References: Resnick 1.8 (discrete time)
Karlin and Taylor Sec. 6.7 (continuous time)

Given a stochastic process X(t), we define $\mathcal{U}_t$ to be the σ-algebra of events that are decided by observations of the stochastic process up to and including time t.

For separable (nice, esp. continuous) stochastic processes with associated probability space Ω

$\mathcal{U}_t$ is generated by sets of the form:

$$A = \{ \omega \in \Omega : X(s_i, \omega) \in B_i \text{ for } i=1,\ldots,n \}$$

for Borel (nice) sets $B_i \subseteq S$

and times $0 \leq s_i \leq t$,

arbitrary $n \geq 1$.

We typically view these sigma-algebras as a family parameterized by time t and they satisfy the reasonable condition that $\mathcal{U}_t \subseteq \mathcal{U}_{t'}$ for $t \leq t'$

Such a family of sigma-algebras is called a filtration, and when we define the σ-algebras as above, we say they are the filtration generated by the stochastic process X(t). The filtration generated by a stochastic process is simply a way of describing which events are determined by observing the stochastic process up to particular moments of time.

Markov property in term of filtrations:

all info about $X(t)$ up to time s

$$P(X(t) = j \mid \mathcal{U}_s)$$

$$= P(X(t) = j \mid X(s))$$

provided $0 \leq s < t$.