CS 547 Lecture 20: Closed Queueing Networks

Daniel Myers

At this point, we've established all of the basic techniques required to solve open queueing models with Poisson arrivals. We now turn our attention to closed networks.

Closed Queueing Models

Closed networks are extremely important for computer systems modeling because they capture the notion of *interactivity* in the system. In an open network, customers simply arrive, receive service, and depart. In a closed network, we can model a set of users submitting requests to a system, waiting for results, then submitting more requests.

Many computer systems exhibit this type of interactive behavior – human users interacting with a system, threads contending for a lock, processes blocking for I/O, networked servers waiting for a response message. Closed models are arguably more common in computer systems modeling than open models.

Product Forum Queueing Networks

We can compute exact solutions for a special class of closed queueing models called *product form queueing networks* (PFQNs).

A PFQN consists of a collection of queueing and delay centers. It satisfies the following conditions.

- All queueing centers are either FCFS, PS, or LCFSPR
- Any number of delay centers are allowed
- If a center is FCFS, its service times must be exponentially distributed
- If a FCFS center has multiple service classes, they must all have the same average service time
- External arrivals, if any, are Poisson (we won't consider any mixed models of this type)
- Routing is state-independent (the path a customer takes cannot depend on the current state of the system)

When these conditions are satisfied, we can use the mean-value analysis (MVA) algorithm to compute exact residence times at each center and in the entire network.[1] When these conditions are not precisely satisfied, we can "bend" the MVA algorithm to produce acceptable approximate results.

The Arrival Theorem

Consider a FCFS queue with exponential arrivals in a closed network. We can analyze this queue using the tagged customer method, just as we did for the M/M/1 system.

[1] There is a small set of networks that do not satisfy these conditions but are still product form.

We are interested in deriving an expression for $Rbar_k(n)$, the average residence time at queueing center k when there are n total customers in the closed system. The number of customers, n, is a key parameter in a closed system, so we'll derive our results as functions of n.

When a new customer arrives to the queue, there will be some number of other customers already present and waiting, including a customer that is in service. Let $\overline{A}_k(n)$ represent the average number of customers present at an arrival to queueing center k when there are n total customers in the closed system.

Exponential service times are memoryless, so each of these customers will require an average service time of $\overline{s}$. The tagged customer also needs $\overline{s}$ for its own service time. The average residence time is thus

$$\overline{R}_k(n) = \overline{s}_k + \overline{A}_k(n)\overline{s}_k$$

We derived a similar equation for the M/M/1 queue. In that case, we were able to use the PASTA property to argue that $\overline{A}_k = \overline{Q}_k$. In a closed system, however, there is no Poisson arrival process, so PASTA does not apply. We need a new result to simplify this equation and solve for $\overline{R}_k(n)$.

The key result for closed PFQNs is the *Arrival Theorem*, which states that the state of the queue at an arrival instant is equal to the average state of the queue with one customer removed from the system. That is,

$$\overline{A}_k(n) = \overline{Q}_k(n-1)$$

This result is not obvious, but it makes some amount of sense. If there are n customers in the entire system, and one of those customers is arriving to queueing center k, then there are $n-1$ *other* customers distributed throughout the system. Therefore, the state of queue k at the instant one customer arrives depends on the average behavior of those other $n-1$ customers.

A proof of the Arrival Theorem is beyond the scope of this course, as it relies on some detailed analysis of closed PFQNs using Markov chains. We will examine some related results towards the end of the course when we talk about using Markov models to solve queueing systems.

We now have the first part of our closed model solution.

$$\overline{R}_k(n) = \overline{s}_k + \overline{Q}_k(n-1)\overline{s}_k$$

In the next lecture, we'll use this equation to develop the mean-value analysis algorithm, which we can use to actually calculate the residence times and throughputs in a closed system.