IEOR 3106: Fall 2012, Professor Whitt

Time Reversibility and Queueing Networks, Thursday, November 1

In this class we will go over Section 6 in the long CTMC notes. See Sections 6.6 and 8.4 in Ross for additional discussion.

I. A Cafeteria Example

An example of an **open network of queues** might be a cafeteria with 7 separate stations: salad bar, sandwich bar, pizza counter, vegetarian section, hot food area, drink area and cashiers. To get a nice simple probabilistic description of the state of the system in the "long run," we would assume that arrivals enter according to a Poisson process. People might then choose their first station independently according to specified probabilities. We also need to specify probabilities for where customers go after completing service at each station. Since different customers will have different preferences, we need probabilities. Then each customer might have a service time at each station that is exponentially distributed with a specified mean, which depends on the station. Each service area has one or more servers. Customers who cannot be served immediately upon arrival wait in line, to be served in order of arrival at that station. Each station separately can be modelled as an $M/M/s$ queue (birth and death process with Poisson arrivals, exponential service times and s servers). Thus somebody who only wants a drink might go directly to the drink station and then directly to the cashier, and depart before many people who arrived in the cafeteria earlier. We might be interested in determining how many servers we need at each area. The entire system of 7 queues can also be modelled as a CTMC, but now it has special structure, so that it is possible to determine the limiting probabilities of the states for the entire system.

1. key fact 1 (another miracle provided by the exponential distribution): the limiting numbers of customers at the different stations at time t are stochastically independent. Letting ∞ denote time in the long run, we have

$$P(X_1(\infty) = k_1, \ldots, X_7(\infty) = k_7) = P(X_1(\infty) = k_1) \times \cdots \times P(X_7(\infty) = k_7),$$

where we need to compute the probabilities $P(X_i(\infty) = k_i)$ properly.

2. key fact 2: the (net) arrival rate at each separate queue in the long run can be obtained by solving the traffic rate equations; see (6.12) -(6.14) in the notes.

3. key fact 3: given the net arrival rate at each queue, each separate queue is an $M/M/s$ queue, and thus can be analyzed as a birth-and-death process. That yields the formulas $P(X_i(\infty) = k_i)$. We plug these into the product form above in step 9.

4. key fact 4: The model above is an open network of queues, in which customers come from outside, mover around inside and then eventually leave. Another related model of interest is a **closed network of queues**. The mode has a fixed population of customers. Closed models are often useful to describe computer systems and manufacturing systems that tend to process a fixed number of jobs at any one time, with new jobs replacing old ones when the old ones have completed their required processing.

II. Basic Theory

These are the main ideas:

1. reverse-time CTMC. (See (6.1)-(6.5).)

2. time-reversible CTMC. (See (6.6).)

3. detailed (or local) balance equations. (See (6.6).)

4. main example: birth-and-death processes. (See Theorem 6.3.)

5. the departure process from an $M/M/s$ queue. (See Theorems 6.5 and 6.6.)

6. the limiting probabilities for two $M/M/1$ queues in series. (See Theorem 6.7.)

7. an open network of queues. (See Theorem 6.9. You will *not* be responsible for the details on pages 35-37 or any proofs anywhere.)