COLUMBIA BUSINESS SCHOOL

Spring 1999
Professor Costis Maglaras
B9801-33 Stochastic Processing Networks
409 Uris Hall

Part I: Reversibility and Product-Form Networks

Summary. The first part of the course is a brief *tour d'horizon* of the descriptive theory of product form networks. This is the class of queueing networks for which exact solutions and performance characterizations exist. In this section we introduce some of the basic definitions, network models, and key concepts to be used later on. In detail, section I.1 contains some background material on Markov Processes, and section I.2 inroduces the concept of reversibility, which is central in the development of the theory of product form networks. Section I.3 covers Burke's "Output Theorem," and gives a first glimse to simple product form results. From that point onwards, we turn to product form networks. Section I.4 describes and analyzes open queueing networks of reversible queues. Section I.5 generalizes these results to the class of "Kelly networks" of quasi-reversible queues. Section I.6 describes networks with symmetric queues that allow for general service time dsirtibutions. Finally, section I.7 introduces the corresponding class of closed queueing networks.

The primary reference is the book by Kelly [Kel79]. Important results will be refered to using the labels found in [Kel79]. Much of the material we cover is based on lecture notes originally developed by Professor Michael Harrison.

I.1 Preliminaries: Markov Processes

1. Let $X(t)$ be a stochastic process taking values on a discrete (countable) state space $\mathcal{S}$. The time parameter t takes values on a set $\mathcal{T}$. For a discrete time process $\mathcal{T}$ is the set of integers $\mathbf{Z}$, that is $t = 0, 1, 2, \ldots$, while for a continuous time process $\mathcal{T}$ will be the real line $\mathbf{R}$.

2. If $(X(t_1), X(t_2), \ldots, X(t_n))$ and $(X(t_1+\tau), X(t_2+\tau), \ldots, X(t_n+\tau))$ have the same probability distribution for all $t_1, t_2, \ldots, t_n, \tau \in \mathcal{T}$ then the process $X(t)$ is said to be *stationary.*

3. The stochastic process $X(t)$ is a *Markov* process (MP) if for any $t_1 < t_2 < \cdots < t_n < t_{n+1}$

$$P(X(t_{n+1}) = x_{n+1} | X(t_1) = x_1, X(t_2) = x_2, \ldots, X(t_n) = x_n) = P(X(t_{n+1}) = x_{n+1} | X(t_n) = x_n); \tag{I.1}$$

i.e., the future is independent of the past given the present or else, all information contained in the past evolution of the process that is useful in predicting its future behavior is contained in the current state of the process.

Fact I.1.1 *The stochastic process $X(t)$ is Markov, if for any $t_1 < t_2 < \cdots < t_m < \cdots < t_n$ conditional on $X(t_m)$ (the present), $(X(t_1), X(t_2), \ldots, X(t_{m-1})$ (the past) is independent of $(X(t_{m+1}), \ldots, X(t_n))$ (the future).*

4. A MP is *time homogeneous* if $P(X(t+\tau)=k|X(t)=j)$ does not depend on t.

5. A MP $X(t)$ *irreducible* if every pair of states (k,j) communicates $(k,j\in\mathcal{S})$.

6. Hereafter, when we say $X=\{X(t),t\geq 0\}$ is a MP, it is understood that

   (a) the state space $\mathcal{S}$ is discrete (countable),

   (b) the process is time homogeneous,

   (c) irreducible (i.e., all pairs of states in $\mathcal{S}$ communicate)

   (d) X stays in each state a positive length of time, and

   (e) only finitely many transitions occur in a finite amount of time.

   The last two conditions are additional assumptions imposed in order to exlude some continuous time Markov processes that exhibit very erratic behavior.

7. Define the *transition rate* from state j to state k by

$$q(j,k)=\lim_{h\to 0}\frac{1}{h}P(X(t+h)=k|X(t)=j), \tag{I.2}$$

   for $k\neq j$, with $q(j,j)=0$. Let $q(j)=\sum_{k\in\mathcal{S}}q(j,k)$.

   Fact I.1.2 *Under the assumptions imposed on X, starting from any state j, X stays there an exponentially distributed length of time with mean $1/q(j)$; the probability that the next state visited is k is $p(j,k)\equiv q(j,k)/q(j)$. We call $p(j,k)$ and $\{p(j,k)\}$ the transition probabilities and transition matrix for the* embedded jump chain *respectively. Note that $\sum_{k\in\mathcal{S}}p(j,k)=1$.*

8. An *equilibrium distribution* for X is defined as a set of positve numbers $\{\pi(j),j\in\mathcal{S}\}$ that sum to 1 $(\sum_{j\in\mathcal{S}}\pi(j)=1)$ and satisfy the *equilibrium equations*

$$\pi(j)\sum_{k\in\mathcal{S}}q(j,k)=\sum_{k\in\mathcal{S}}\pi(k)q(k,j),\quad j\in\mathcal{S}. \tag{I.3}$$

   The LHS of (I.3) is equal to $\pi(j)q(j)$ which is the "probability flux" out of state j. The RHS (I.3) is equal to the "probability flux" into state j. In equilibrium, the flux out of any state j should be equal to the flux into state j; this is precisely (I.3) which is also refered to as the *global balance equations*.

   Fact I.1.3 *If an equilibrium distribution exists, then it is unique, and it is moreover*

   (a) the unique stationary distribution of X,

   (b) the limit distribution of X for every starting state j, that is

$$\lim_{t\to\infty}P(X(t)=k|X(0)=j)=\pi(k),\quad j\in\mathcal{S},$$

(c) the long-run occupancy distribution of X, almost surely for every starting state; that is

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T \mathbf{1}_k(t)dt = \pi(k) \quad a.s.,$$

where $\mathbf{1}_k(t)$ is the indicator function ($\mathbf{1}_k(t) = 1$ if $X(t) = 1$, and zero otherwise).

Fact I.1.4 *If there exists a set of numbers $\{\pi(j)\}$ that satisfy (I.3) that sum up to ∞, then there is no equilibrium distribution, and X has no stationary distribution. In this case, $P(X(t) = k|X(0) = j) \to 0$ as $t \to \infty$ for all $j, k \in \mathcal{S}$, and the limiting fraction of time spent in any state k is also zero (almost surely, for any starting state j).*

9. As an example try analyzing the simplest example possible of the $M/M/1$ queue: start from the MP description for the queue, verify all assumptions described in (6), derive the stationary distribution, the expected waiting time and the expected (total) delay per arriving job, the expected backlog in the system. How do these measures vary with the load (or traffic intensity $\rho = \lambda/\mu$, where λ and μ are the arrival and service rates respectively)? This is the simplest example of a birth-death (B&D) process with birth rate λ (arrivals) and death rate μ (departures).

I.2 Reversibility

1. The property of *reversibility* plays an important role in the descriptive theory of queueing networks. Intuitively speaking, if we take a trace of a reversible process and run it backwards in time the resulting process is statistically indistinguishable from the original process. That is, the behavior of a reversible process remains the same when the direction of time is reversed.

Definition I.2.1 *X is said to be* reversible *if $(X(t_1), \ldots, X(t_n))$ has the same distribution with $(X(\tau - t_1), \ldots, X(\tau - t_n))$.*

Proposition I.2.1 *A reversible process is stationary.*

Proof.

$$\begin{aligned}(X(t_1), \ldots, X(t_n)) &\sim (X(\tau - t_1), \ldots, X(\tau - t_n)) \\ &\sim (X(-t_1), \ldots, X(-t_n)) \\ &\sim (X(\tau + t_1), \ldots, X(\tau + t_n));\end{aligned}$$

the first step follows from the definition of reversibility, the second by setting $\tau = 0$, and the third by the reversibility of the "time reversed" process $X(-t)$. ◇

2. **Theorem 1 (Thm. 1.3 [Kel79])** *A stationary MP X is reversible if and only if there exists a probability distribution $\{\pi(j), j \in \mathcal{S}\}$ satisfying the* detailed balance equations

$$\pi(j)q(j,k) = \pi(k)q(k,j) \quad \text{for all } j,k \in \mathcal{S}. \tag{I.4}$$

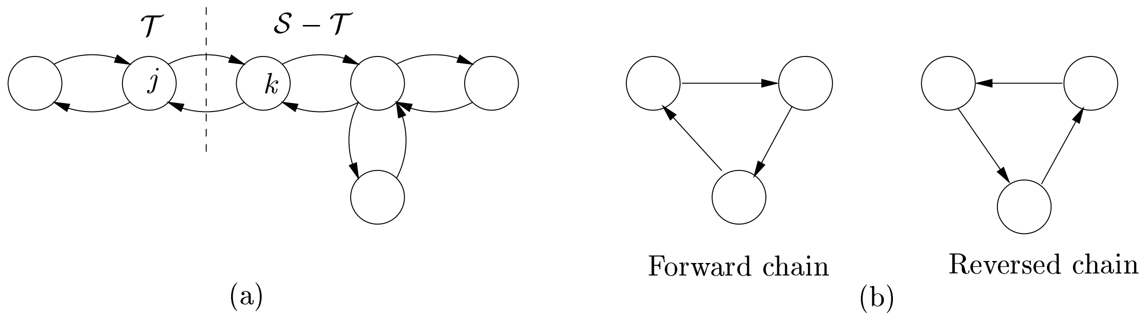

Figure 1: Examples: (a) reversible chain – simple network with tree structure; (b) non-reversible chain – the forward chain has clockwise transitions whereas the reversed chain has anticlockwise transitions; from [BG92].

Remark. Any such collection $\{\pi(j)\}$ must be the equilibrium distribution of X; sum (I.4) over $k \in \mathcal{S}$ to obtain the equilibrium equations (I.3).

3. Examples of reversible processes.

Definition I.2.2 *The* graph G *associated with the MP* X *has a vertex for each state* $j \in \mathcal{S}$ *and an edge between* j *and* k *if either* $q(j,k)$ *or* $q(k,j)$ *is positive. (Irreducibility* $\Rightarrow$ *graph is connected.) A* cut *is a division of* $\mathcal{S}$ *into complementary sets* $\mathcal{T}$ *and* $\mathcal{S}-\mathcal{T}$.

Proposition I.2.2 *If* G *is a tree, then* X *is reversible. (Thus the standard B&D process is reversible.)*

Proof. Introduce a cut along the edge (j,k) that divides $\mathcal{S}$ into $\mathcal{T}$ and $\mathcal{S}-\mathcal{T}$. In equilibrium, the flow balance equations along this cut are

$$\sum_{i\in\mathcal{T},\ l\in\mathcal{S}-\mathcal{T}} \pi(i)q(i,l) = \sum_{i\in\mathcal{T},\ l\in\mathcal{S}-\mathcal{T}} \pi(l)q(l,i). \tag{I.5}$$

Since G is a tree, it follows that $q(i,l)=0$ for all $i \neq j$, $l \neq k$, and thus (I.5) reduces to (I.4). $\diamond$

Corollary I.2.1 *The queue length process associated with an* $M/M/1$ *queue is reversible.*

Proof. The queue length process of an $M/M/1$ queue is a B&D process. $\diamond$

Proposition I.2.3 *If the transition rates of a reversible MP with state space* $\mathcal{S}$ *and equilibrium distribution* $\{\pi(j)\}$ *are altered by changing* $q(j,k)$ *to* $cq(j,k)$ *for* $j\in\mathcal{T}$ *and* $k\in\mathcal{S}-\mathcal{T}$*, then the new MP is reversible in equilibrium and has equilibrium distribution*

$$\pi'(j) = \begin{cases} B\pi(j) & \forall j \in \mathcal{T} \\ Bc\pi(j) & \forall j \in \mathcal{S}-\mathcal{T}, \end{cases}$$

where B *is some normalization constant.*

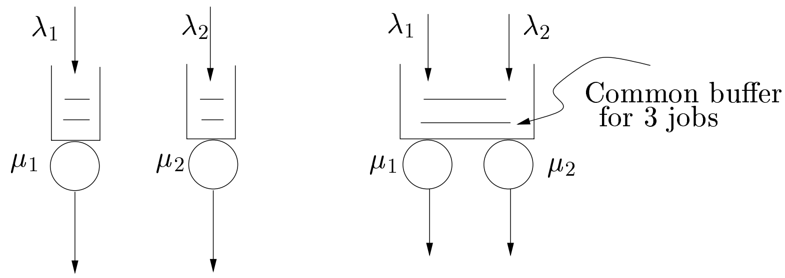

Figure 2: Two $M/M/1$ queues in parallel operating with common waiting room of size 4.

Proof. $\{\pi'(j)\}$ and $\{q'(j,k)\}$ satisfy the detailed balance equations (I.4). (This is also true for $c = 0$.) $\diamond$

Corollary I.2.2 *If the state space of X is simply truncated (or restricted) to stay in $\mathcal{T}$ (this is the case $c = 0$), then the new equilibrium distribution is*

$$\pi'(j) = \frac{\pi(j)}{\sum_{k \in \mathcal{T}} \pi(k)}, \quad j \in \mathcal{T}.$$

Example: Two independent $M/M/1$ queues are operating in parallel with infinite waiting rooms. Now impose a common waiting room of size 4. (That is, at any time there can be at most 3 jobs waiting for service and each server can be processing at most one job.)

Let n_1 and n_2 denote the # of jobs in queues one and two respectively. First we need to show that the two-dimensional queue length process $(n_1(t), n_2(t))$ is reversible. From independence of the two queues we have that

$$\pi(n_1, n_2) = \pi_1(n_1) \cdot \pi(n_2), \quad \text{where } \pi(n_i) = (1 - \rho_i)\rho_i^{n_i}.$$

It is now easy to verify conditions (I.4) to show that $(n_1(t), n_2(t))$ is reversible. Imposing the common waiting room, the truncated state space is shown in Figure 3. Using the last corollary we get that $\pi'(n_1, n_2) = B\pi(n_1, n_2)$ for $(n_1, n_2) \in A$.

4. **Theorem 2 (Thm. 1.12 [Kel79])** *If $X(t)$ is a stationary MP with transition rates $\{q(j,k\}$ and equilibrium distribution $\{\pi(j)\}$, then the* reversed *process $X(\tau - t)$ is a stationary MP with the same equilibrium distribution and transition rates*

$$q'(j,k) = \frac{\pi(k)}{\pi(j)} q(k,j), \quad j,k \in \mathcal{S}. \tag{I.6}$$

That is, the probability flux from k to j for the original process is equal to the flux from j to k for the time reversed process. This is intuitive since every forward transition from j to k corresponds to a backwards transition from k to j. In equilibrium, the number of forward transitions from j to k is equal to $\pi(j)q(j,k)$ and the number of backward transitions from k to j are equal to $\pi(k)q(k,j)$. Finally, for the process to be reversible we further require that $q'(j,k) = q(j,k)$. In this case, (I.6) reduces to (I.4).

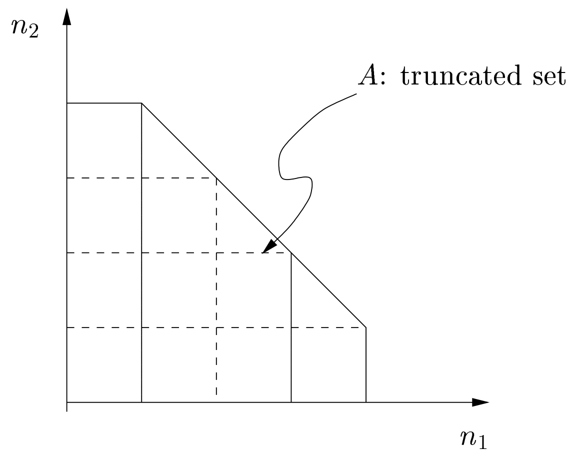

Figure 3: Truncated state space – common waiting room of size 4

5. **Theorem 3 (Thm. 1.13 [Kel79])** *Let X be a stationary MP with transition rates $\{q(j,k)\}$. If $\{q'(j,k), j,k \in \mathcal{S}\}$ are non-negative numbers and $\{\pi(j), j \in \mathcal{S}\}$ are positive numbers summing to one such that*
$$\pi(j)q(j,k) = \pi(k)q'(k,j), \quad j,k \in \mathcal{S}, \tag{I.7}$$
and
$$q'(j) = q(j), \tag{I.8}$$
then $\{\pi(j)\}$ is the equilibrium distribution of the process and $\{q'(j,k)\}$ are the transition rates of the reversed process.

   Theorem 1.13 will prove to be very useful in the discussion of complicated MP where checking the equilibrium conditions becomes very tedious, but for which guessing by inspection the transition rates for the reversed process is simple.

I.3 The Output Theorem

The following two results, taken from Kelly [Kel79] sections 2.1 and 2.2 respectively, foreshadow much more general results to come, but they are easy to prove.

1. Burke's *Output Theorem* for the $M/M/1$ queue

   Theorem 4 (Thm. 2.1 [Kel79]) *The output process from a stationary $M/M/1$ queue is a Poisson process, and the number of jobs in the queue at time t_0 is independent of the departure process prior to t_0.*

   Proof. The queue length process $n(t)$ is reversible (B&D process), so the reversed process $n(-t)$ is also an $M/M/1$ queue length process. Let $n^*(t) \equiv n(t_0 - t)$ for $0 \le t \le t_0$ denote the reversed process. Let $\{D(t), 0 \le t \le t_0\}$ denote the departure process from the $M/M/1$ queue and

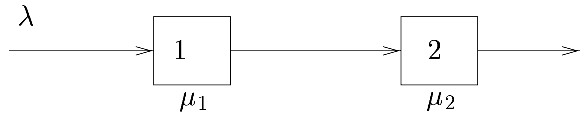

Figure 4: Two queues in tandem

$\{A^*(t), t \leq t_0\}$ be the # of arrivals in the time reversed queue length process n^* during $[0, t]$. That is,

$$A^*(t) = D(t_0) - D(t_0 - t), \quad 0 \leq t \leq t_0.$$

Now n^* is a stationary $M/M/1$ queue length process so (i) $A^*(t)$ is Poisson and (ii)

$$\{A^*(t), 0 \leq t \leq t_0\} \quad \text{is independent of } n^*(0),$$

which implies that the departure process is Poisson and that $\{D(t), 0 \leq t \leq t_0\}$ is independent of $n(t_0)$. $\diamond$

At first glance this result appears surprising. The output from the $M/M/1$ queue occur at a rate μ while the server is busy and then at zero rate whenever the server is idle. Nevertheless, observing the output of the queue in stationarity (over a long time span) given no information about the server rate μ, and in fact the output process is Poisson of rate λ.

Moreover, one would expect that a stream of close departure would suggest a busy system with a large backlog. Burke's Theorem shows that is not true. Note, however, that Burke's Theorem makes no claims about the backlog in the system before a stream of closely spaced departures. In this case, the number in the queue would tend to be large, in accordance to intuition.

2. The "Output Theorem" applies to more general queues. In particular, it is true for any queue with a Poisson arrival process where the number in the queue behaves like a B&D process; e.g., the $M/M/s$ queue.

3. The "Output Theorem" forms the basis for the theory of product form networks. It allows us to "decouple" the dependence between various queues (or nodes) in a network and greatly simplify analysis.

   The simplest queueing network that we can consider is a series of simple queues in tandem. Consider for simplicity a network of two queues in tandem shown in Figure 4. There is one exogeneous Poisson arrival stream of rate λ and the service times at each server i are exponentially distributed with mean $1/\mu_i$ for $i = 1, 2$ and independent of each other. Assume that $\rho_i = \lambda/\mu_i < 1$ for $i = 1, 2$.

   Corollary I.3.1 *The steasy state distribution is*

$$\pi(n_1, n_2) = \prod_{i=1}^{2} (1 - \rho_i)\rho_i^{n_i}, \quad n_1, n_2 = 0, 1, \ldots.$$

   This is our first product form result. Proof?

Resource allocation: Let's assume that the service discipline at each station is First-In-First-Out (FIFO). Consider the problem of allocating resources in this two station tandem network in order to minimize the expected throughput time seen by any arriving job (this is the total time the job will spend in the system). Let W_i be the waiting time encountered at each queue and $W = W_1 + W_2$. Our problem is to choose the processing rates μ_1 and μ_2 to

$$\begin{array}{ll} \min & \mathbf{E}[W] \\ s.t. & c_1\mu_1 + c_2\mu_2 \le C. \end{array}$$

The last constraint is a "total budget constraint."

First, observe that W_1 is independent of W_2, and thus the waiting time encountered at each of the queues is that corresponding to an $M/M/1$ queue in equilibrium (why?). (Note however that D_1 is not independent of D_2!) That is,

$$\mathbf{E}[W_i] = \frac{1}{\mu_i - \lambda}, \quad \text{for } \mu_i > \lambda, \quad i = 1, 2.$$

So the resource allocation problem is

$$\begin{array}{ll} \min & \frac{1}{\mu_1-\lambda} + \frac{1}{\mu_2-\lambda} \\ s.t. & c_1\mu_1 + c_2\mu_2 \le C. \end{array}$$

The Lagrangian for this problem is

$$\mathcal{L}(\mu_1, \mu_2, \alpha) = \frac{1}{\mu_1 - \lambda} + \frac{1}{\mu_2 - \lambda} + \alpha(C - c_1\mu_1 - c_2\mu_2),$$

and the first order optimality conditions are

$$\begin{array}{lll} \frac{1}{(\mu_1 - \lambda)^2} - \alpha c_1 = 0 & \Rightarrow & \mu_1 = \lambda + \frac{1}{\sqrt{\alpha c_1}} \\ \frac{1}{(\mu_2 - \lambda)^2} - \alpha c_2 = 0 & \Rightarrow & \mu_2 = \lambda + \frac{1}{\sqrt{\alpha c_2}}. \end{array}$$

Observing that the budget contraint will be binding (why?), we solve for the multiplier α to get

$$\sqrt{\alpha} = \frac{\sum_{i=1}^{2} \sqrt{c_i}}{C - (c_1 + c_2)\lambda}.$$

The optimal resource allocations μ_i can now be determined. This is the so called "square root" capacity assignment, first derived by Kleinrock. Roughly speaking, the optimal assignment first allocates just enough capacity at each station in order to satisfy its arrival rate, and then allocates the excess capacity among stations in proportion to the square root of their weighted arrival rates. This result can be extended to more general networks that admit product form solutions.

4. The "Output Theorem" (that is the Poisson nature and independence of the departure processes from each queue) can be extended to the case of acyclic (or feedforward) networks. Roughly speaking, in these networks jobs are flowing in a downstream direction, hence the name *feedforward.* The results of this section can be extended to acyclic networks using a similar induction argument as one traverses further downstream in the network.

I.4 Open Networks of Reversible nodes

Our discussion is based on section 3.1 of Kelly [Kel79].

1. *The model.* We have a network of J service stations and K job classes ($K \geq J$). Each class $k = 1, \ldots, K$ is served at a unique station $s(k) \in \{1, \ldots, J\}$. Define $C(j) = \{k : s(k) = j\}$ to be the set of job classes that are served at station j; we call this set the *constituency* of station j. In general the mapping from classes to stations is many-to-one; that is, many classes may be served at the same station.

   External arrivals into classes $1, \ldots, K$ occur according to independent Poisson processes at average rates $\nu_1, \ldots, \nu_K$. After completing service at station $s(k)$, a class k job will become (instantaneously) a class l job with probability $P(k, l)$, where $P = [P(k, l)]$ is a $K \times K$ transient *switching* or *routing* matrix of the queueing network; note that routing is Markovian, i.e., depends only on current job class designation and is independent of all past. (Transience of the switching matrix implies that any job that enters the network will eventually leave the system. A matrix P with this property is called sub-stochastic. Its spectral radius is less than one.)

2. *Effective arrival rates.* Since P is transient there exists a unique solution $\alpha(1), \ldots, \alpha(k)$ to the following linear system of *traffic equations*

$$\alpha(l) = \nu(l) + \sum_{k=1}^{K} \alpha(k) P(k, l), \quad \text{for } l = 1, \ldots, K, \tag{I.9}$$

   or in matrix form

$$\alpha = \nu + P'\alpha \;\Rightarrow\; \alpha = (I - P')^{-1}\nu. \tag{I.10}$$

   Interpret $\alpha(k)$ as the total arrival rate into class k, including both external arrivals and internal transitions.

3. *Jackson and Kelly networks.* An important special case of this class of networks are the so called Jackson networks, that have a one-to-one relation between job classes and stations. In this case, $J = K$ and each station serves jobs in FIFO. In Jackson networks, jobs that transition to a certain station "loose" their identity with respect to their future routes. Consider the example of Figure 5.

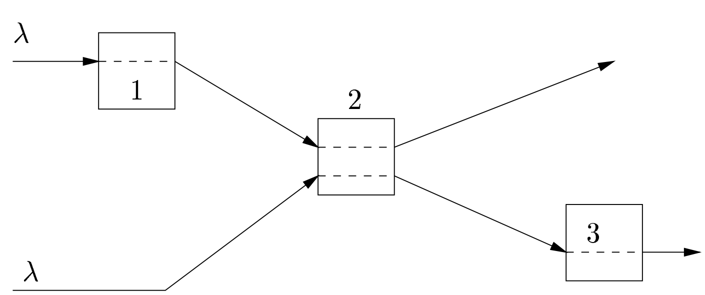

Figure 5: Example of a non-Jackson network

Clealry, the routing is not Markovian since the future route of a job completing service at station 2 depends on its previous history. To approximate it with a Jackson network we set: $P(1,2) = 1$, $P(1,3) = 1$, and $P(2,3) = 1/2$. We use $P(2,3) = 1/2$ to approximate the original network because on the average half the jobs through station 2 will go to station 3. It turns out that this setup gives the correct product form solution (unlike what we would expect so far.)

Within the class of networks considered here, one can incorporate Markovian switching among *classes* with a many-to-one relationship between classes and stations, provided that all classes served at any given station have identical service requirements. In the previous example one would need to introduce different class designations for jobs arriving to station 2 after a visit at station 1 to the jobs that will next visit station 3 as shown below:

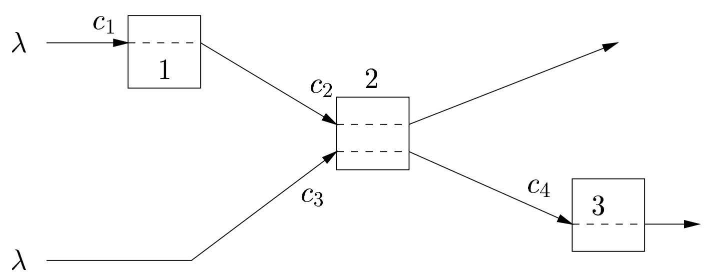

4. *Deterministic vs. Markovian routing.* In contrast the setup that allows for Markovian switching, Kelly assumes that jobs of I different types arrive according to independent Poisson processes, and jobs of each type follow a determnistic route (sequence of stations visited) through the network. The two formulations are (essentially) equivalent. Here is an example of how one would model within Kelly's deterministic routing an example with Markovian switching:

   First note that random splitting of a Poisson stream give independent Poisson streams. The derived network (after the input has been split) is equivalent to the original one if we do not use information about the splitted input stream in the station's service rate and scheduling discipline.

   In general, the general Kelly network with Markovian switching among classes can be reduced to one with multiple types and deterministic routes as follows: just define one type for each possible route through the network (lots of them!!). If the original switching matrix P had feedback then we would get countably infinite types. Even in this case we can get good finite approximations. Therefore, the restriction to Markovian switching among classes is really no restriction at all, because the number of "classes" or "types" can be arbitrarily large.

   Of course, the assumption of independent Poisson inputs is restrictive, as are the following assumptions about the "servive mechanisms" at the various stations.

5. *Service mechanism.*

   (a) The "service requirement" of each job at each stage of its route is exponentially distributed with unit mean, independent of type and all previous history; that is, all classes served at a given station have identical service requirements.

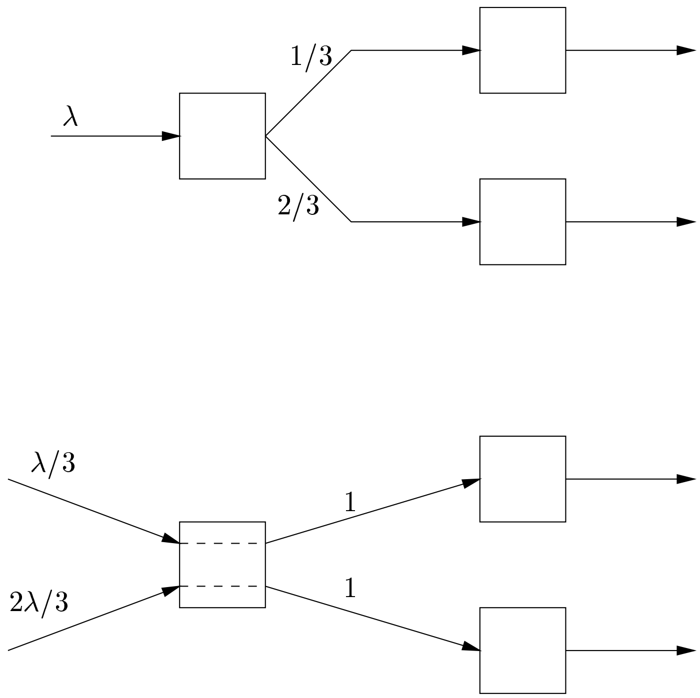

Figure 6: Detereministic modeling of Markovian switching

(b) When the total queue length at station j is n_j, service effort there is supplied at a total rate of $\phi_j(n_j)$.

(c) A proportion $\gamma_j(l, n_j)$ of this effort is devoted to the customer in position l of the queue associated with station j; when this job's service requirement is completed, jobs in positions $l+1, \ldots, n_j$ all move up by one position. (That is, the server is allowed to split its effort between the different jobs in the queue.)

(d) When a job arrives at queue j, it moves into position l ($l = 1, \ldots, n_j + 1$) with probability $\delta_j(l, n_j + 1)$; jobs previously in positions $l, l+1, \ldots, n_j$ all move back one position. (That is, jobs enter the queue in a randomized fashion; of course, in this way we could model the common special cases where a job would join at the end or at the front of the queue.)

The "departure rate" for the job in position l at station j is $\phi_j(n_j)\gamma_j(l, n_j)$.

Remark. Note that neither the positioning of a new arrival nor the allocation of service are allowed to change according to the type (or class) of the jobs involved, and the **service requirements of all job classes at a given station are identical**.

6. *Examples of scheduling disciplines.* How would you model these policies in our framework?

- *FIFO discipline.*
- *Processor sharing:* Equal allocation of capacity to all jobs queued at the station.
- ∞*-server (delay) node.*

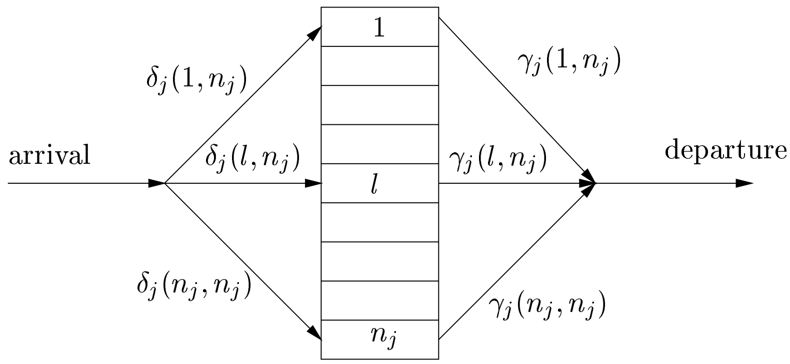

Figure 7: Schematic model of the basic queue for sec. I.4

7. *State descriptor.* To describe the state of the system at any given time, we denote by

n_j	the total queue length at station j
$c_j(l)$	the class designation of the job in position l.

Then $\mathbf{c}_j = (c_j(1), \ldots, c_j(n_j))$ describes the state at queue j and $\mathbf{C} = (\mathbf{c}_1, \ldots, \mathbf{c}_J)$ is the Markov state descriptor for the total network.

From the verbal descriptions given so far you should be able to construct the state space $\mathcal{S}$ for the Markov chain describing the evolution of the system, and the transition intensities $q(\mathbf{C}, \mathbf{D})$ for $\mathbf{C}, \mathbf{D} \in \mathcal{S}$. (The transition intensities $q(\mathbf{C}, \mathbf{D})$ will obviously depend on the particular scheduling discipline used.)

8. Given that a departure occurs from queue j, the probability it is the job in position l is $\gamma_j(l, n_j)$. Given that a job has just arrived at queue j, bringing the total length to n_j, the probability that it has joined in position l is $\delta_j(l, n_j)$. Let

$$\begin{aligned} \alpha_j(k) &= \begin{cases} \alpha(k) & \text{if } k \in C(j) \\ 0 & \text{otherwise,} \end{cases} \\ a_j &= \sum_{k=1}^{K} \alpha_j(k) = \text{ total arrival rate into station } j, \end{aligned}$$

so

$$\frac{\alpha_j(k)}{a_j} = \text{ fraction of all arrivals at station } j \text{ who are of class } k.$$

Let

$$p_j(n_j) = b_j \prod_{l=1}^{n_j} \frac{a_j}{\phi_j(l)},$$

where b_j is the appropriate normalization constant that makes this a probability distribution.

This is the equilibrium distribution for a B&D process with constant arrival rate a_j and death rate $\phi_j(l)$ when the population size is l. It is also the equilibrium distribution for the total queue length at station j.

Let

$$\pi_j(\mathbf{c}_j) = p_j(n_j) \prod_{l=1}^{n_j} \frac{\alpha_j(c_j(l))}{a_j};$$

the last term describes the particular job configuration in the queue.

9. **Theorem 5 (Thm. 3.1 [Kel79])** *The equilibrium distribution of the open queueing network described above is*

$$\pi(\mathbf{C}) = \prod_{j=1}^{J} \pi_j(\mathbf{c}_j), \quad \mathbf{C} \in \mathcal{S}.$$

Remark. Thus, in equilibrium, (i) the states of the different queues are independent, (ii) the total queue length at station j has distribution $\{p_j(\cdot)\}$, and (iii) the probability that the job in position l in queue j is of class k is $\alpha_j(k)/a_j$, independent of everything else.

Proof. (Sketch) Let's define the reversed network. First, let

$$\nu'(k) = \alpha(k) \left(1 - \sum_{l=1}^{K} P(k,l)\right) \quad \text{for } k = 1, \ldots, K$$

and

$$P'(k,l) = \frac{\alpha(l)}{\alpha(k)} P(l,k) \quad \text{for } k, l = 1, \ldots, K;$$

$\alpha(l)P(l,k)$ is the number of k arrivals due to l transitions, and thus $P'(k,l)$ will be the fraction of k arrivals that is due ot l transitions, or in terms of the reversed process $P'(k,l)$ is the fraction of class k jobs that will transition to class l jobs upon completion of their service.

Now cosnider a related open network in which the external arrival rate to class k is $\nu'(k)$, the switching matrix is $P' = [P'(k,l)]$, the level of service at station j is $\phi_j(n_j)$ as before, and the rates $\gamma_j(l,n_j)$ and $\delta_j(l,n_j)$ are reversed. Let $\{q'(\mathbf{C},\mathbf{D})\}$ be the transition intensity function for this related network. It can be verified that the hypothesized equilibrium dstribution $\pi(\cdot)$ satisfies

$$\pi(\mathbf{C})q(\mathbf{C},\mathbf{D}) = \pi(\mathbf{D})q'(\mathbf{D},\mathbf{C}), \quad \forall\ \mathbf{C},\mathbf{D} \in \mathcal{S},$$

and moreover

$$q(\mathbf{C}) = q'(\mathbf{C}) = \sum_{j=1}^{J} \phi_j(n_j) + \sum_{k=1}^{K} \nu(k) \quad \forall\ \mathbf{C} \in \mathcal{S}.$$

The desired result now follows from Theorem 1.13 of Kelly, and in addition we have that the time-reversed transition intensity function for our original network is $q(\cdot)$. $\diamond$

Remark. The proof of this theorem –as for some of the results to follow– is algrebraic and hinges on the verification fo the conditions of Theorem 1.13. It is not apparent how the probabilistic assumptions made so far are used in obtaining this result. In fact, the essential feature that results in this product form behavior is that the individual queues (nodes) as decribed so far are reversible and satisfy the "Output Theorem" (exercize). A probabilistic argument that exploits this reversible structure and the concequences of the "Output Theorem" can be found in Walrand [Wal83].

10. Following our previous remark, this theorem contains as a special case "Jackson's Theorem" that established the product form nature for the family of Jackson networks.

11. The main assumptions (and limitations) of these models are:

    (a) Multiple classes flowing through the network and Markovian switching; this allows for feedback (i.e., not acyclic structure any more).

    (b) The assumption of Poisson arrivals and exponential service times.

    (c) The scheduling disciplines allowed that cannot differentiate between different classes served at each station.

    (d) All classes served at any give station have identical service requirements.

    The basic topology of this network will dominate most of this course, but the restrictive assumptions (b)-(d) will be relaxed very soon.

I.5 Open Networks of Quasi-Reversible Nodes

Our discussion is based on section 3.2 of Kelly [Kel79]. The routing structure described in section I.4 is quite general for most purposes (i.e., adequate in modelling realistic applications) but the structure of each individual queue was quite restrictive. In this section we will describe a more general class of queues (or nodes– this is the term I will use hereafter) that still achieve the same behavior described by the product form distribution. This class of nodes will be referred to as *quasi-reversible* (Q-R).

1. *The node.* We start by describing the basic structure of a node. Later nodes of a network will be indexed by $j = 1, \ldots, J$. In that setting, the notation used below to describe a node will be augmented by hanging a j as a subscript or a functional argument in order to specify which node we are talking about.

   There is a finite set $\mathcal{C}$ of classes that are served at the node, and its state at any given time is decsribed by some $x \in \mathcal{S}$ (a countable set). (That state descriptor may contain information about the number of jobs in the queue, the class of each job in the queue their position etc..) Associated with each state x is a non-negative vector $n(x) = \{n_c(x), c \in \mathcal{C}\}$ whose c^{th} component tells how many jobs of class c are currently present at the node. Let $\mathcal{S}(c, x)$ be the set of states $x' \in \mathcal{S}$ with $n(x') = n(x) + e_c$ (where e_c is a unit vector with a 1 in the c^{th} position and zeros elsewhere); i.e., the set of x''s that differ from x by one class c arrival.

   A node is characterized by (or defined as) a transition intensity function $q(x, x'|\alpha)$ that has property (I.11) below. Here $\alpha = \{\alpha(c), c \in \mathcal{C}\}$ is a non-negative vector that *parametrizes* what really is a *family of transition intensity functions*; think of α as the vector of intensity parameters for a collection of independent Poisson arrival processes.

$$
\begin{aligned}
&q(x, x'|\alpha) = 0 \ \text{ except in the following cases} && \text{(I.11)}\\
&a)\ \ n(x') = n(x) && [\text{“internal transition”}]\\
&b)\ \ x \in \mathcal{S}(c, x') \text{ for some } c && [\text{“departure transition”}]\\
&c)\ \ x' \in \mathcal{S}(c, x) \text{ for some } c && [\text{“arrival transition”}]
\end{aligned}
$$

In cases a) and b), $q(x, x'|\alpha)$ is actually independent of α, and we will just write $q(x, x')$ hereafter. In case c) we have

$$q(x, x'|\alpha) = \alpha(c)p^c(x, x') \quad \text{where} \sum_{x' \in \mathcal{S}(c,x)} p^c(x, x') = 1,$$

and $p^c(x, x')$ is the probability that a class c arrival will steer the state from x to x'. An example of internal transitions is in the case where the service requirement for some job class follows an Erlang distribution (this is a sum of -say k- IID exponentials). In this case the state descriptor needs to keep track of the stage information. When the stage changes without service completion, we have an internal transition.

To repeat, (I.11) defines a node (not necessarily Q-R). Note that for a given α the node has a stationary distribution $\{\pi(x|\alpha), x \in \mathcal{S}\}$ and if we denote by $X(t)$ the associated MP, then the reversed process $X(-t)$ describes the equilibrium behavior of another node (one must check that (I.11) is satisfied) with transition function

$$q'(x, x'|\alpha) = \frac{\pi(x'|\alpha)}{\pi(x|\alpha)} q(x', x|\alpha) \tag{I.12}$$

and the same stationary distribution $\{\pi(x|\alpha)\}$. In each $n(x(t))$ can change only by a departure or an arrival.

2. *Quasi-reversibility.*

Definition I.5.1 *A node is said to be Q-R if, for each α such that the stationary distribution exists,*

$$\sum_{x' \in \mathcal{S}(c,x)} q'(x, x'|\alpha) = \alpha(c) \quad \forall\ x \in \mathcal{S},\ c \in \mathcal{C}. \tag{I.13}$$

Condition (I.13) gives a precise mathematical articulation of what we mean by quasi-reversibility, but the following interpretation will certainly help your intuition. First, (I.13) really just says that

$$\sum_{x' \in \mathcal{S}(c,x)} q'(x, x'|\alpha) \quad \text{is independent of } x \quad \forall\ c \in \mathcal{C}, \tag{I.14}$$

because (I.14) says that the class c arrival intensity is constant for the reversed process, and the only constant consistent with equilibrium is $\alpha(c)$, the arrival rate for class c in the original process and hence the departure rate for class c when time is reversed.

Now (I.14) is equivalent to saying that arrivals after time t in the reversed system are independent of $X(t)$ which is equivalent to saying that departures up to time t in the original system are independent of the state $X(t)$.

To summarize, quasi-reversibility is an **input-output** property: if you drive the node with a vector Poisson input process and initialize with the stationary distribution, then the vector output is statistically idenstical to the input process, and moreover departures up to time t are independent of the state $X(t)$. That is, *a node is Q-R if the Output Theorem holds.*

3. *"Partial balance" and quasi-reversibility.* Substitute the definition (I.11) of $q'(\cdot,\cdot)$ into (I.13) and you see that (I.13) is equivalent to

$$\pi(x) \sum_{x' \in \mathcal{S}(c,x)} q(x,x'|\alpha) = \sum_{x' \in \mathcal{S}(c,x)} \pi(x') q(x',x|\alpha), \tag{I.15}$$

which are the so called *partial balance equations.* The intuition is the following: starting at state x, a class c arrival can steer the state to x', where $x' \in \mathcal{S}(c,l)$. Condition (I.15) requires that the probability flux from state x to the set $\mathcal{S}(c,x)$ is equal to the probability flux from $\mathcal{S}(c,x)$ back to x; i.e., flux into x due to class c arrivals is equal to the flux out of x' due to class c departures. A similar argument explains the case of transitions due to class c departures.

It remains to consider what happens during "internal transitions" from x to x' for which $x' \neq x$ but $n(x') = n(x)$. First recall that the state transitioning from x can either stay in the set of interest (i.e., $x' \neq x$ and $n(x') = n(x)$) or it will end up in some set $\mathcal{S}(c,x)$ considered above. Since the total probability flux out of x is equal to the total probability flux into x (this follows from the equilibrium conditions), using the partial balance equations we can conclude that also the probability flux out of x to this set (i.e., $X' \neq x$ and $n(x') = n(x)$) will equal the total probability flux from this set into x.

Note that any node that satisfies the detailed balance equations also satisfies the partial balance equations (just add (I.4) over $x' \in \mathcal{S}(c,x)$), and thus reversibility implies quasi-reversibility.

4. *The network.* Nodes are indexed by $j = 1,\ldots,J$. Input and routing are exactly as in section I.4: external arrivals to classes $k = 1,\ldots,K$ via independent Poisson processes at rates $\nu(1),\ldots,\nu(k)$; upon completing service at station $s(k)$, a class k job switches class according to the transient matrix $[P(k,l)]$. Define total arrival rates $\alpha(k)$ and time reversed routing matrix $P' = [P'(k,l)]$ as in the previous section. Let $\pi_j(x) = \pi_j(x|\alpha)$ be tha stationary distribution of node j when it is driven by Poisson inputs at rates $\{\alpha(c), c \in \mathcal{C}_j\}$.

   The "network" is a MP with generic state $x = (x_1,\ldots,x_J)$, where $x_j \in \mathcal{S}_j$ for all j. Its transition intensities are defined as follows

   (a) Suppose that $c \in \mathcal{C}_j$, $x_j \in \mathcal{S}_j(c,x')$, and $x_i = x'{}_i$ for all other nodes i. Then, letting $P(c,0) \equiv 1 - \sum_{l=1}^{K} P(c,l)$, we have that $q(x,x') = q_j(x_j,x'_j)P(c,0)$. (These transitions correspond to class c service completions where the job then leaves the network.)

   (b) Suppose that $c \in \mathcal{C}_j$, $x'_j \in \mathcal{S}_j(c,x_j)$, and $x_i = x'_i$ for all other nodes i. Then $q(x,x') = \nu(c)p^c_j(x_j,x'_j)$. (These transitions correspond to external arrivals into class c.)

   (c) Suppose $n_j(x_j) = n_j(x'_j)$ and $x_i = x'_i$ for all other nodes i. Then $q(x,x') = q_j(x_j,x'_j)$. (These transitions correspond to "internal state changes" at node i.)

   (d) Suppose that $x_i \in \mathcal{S}_i(c,x'_i)$ and $x'_j \in \mathcal{S}_j(c',x_j)$, where $c \in \mathcal{C}_i$ and $c' \in \mathcal{C}_j$ and $x_k = x'_k$ for all other nodes k. Then

   $$q(x,x') = q_i(x_i,x'_i)P(c,c')p^c_j(x_j,x'_j).$$

   (These transitions correspond to service completions of class c jobs at node i where the job next witches to class c'. Note that one can have $i = j$ here, meaning that classes c and c' are served at the same node i. This last formula is similar to that of case (b) only with a different effective arrival rate for class c' arrivals due to class c completions starting at state x_i and transitioning into state x_i'.

(e) Otherwise, $q(x, x') = 0$.

5. *Kelly's Theorem.*

Theorem 6 (Thm. 3.7 [Kel79]) *If the vector of total arrival rates is such that the stationary distribution $\pi_j(\cdot, \alpha)$ exists for every node j, then the unique stationary distribution for the network is*

$$\pi(x) = \pi_1(x_1|\alpha) \cdots \pi_J(x_J|\alpha).$$

In order to prove this theorem one would proceed as follows. First, guess the time-reversed transition intensity function $q'(x, x')$ for the network to be as follows in cases (a)-(e) above:

(a) $q'(x', x) = q_j'(x_j', x_j|\alpha)$

(b) $q'(x', x) = q_j(x_j', x_j|\alpha)\left(1 - \sum_{l=1}^{K} P'(c, l)\right)$

(c) $q'(x', x) = q_j'(x_j', x_j|\alpha)$

(d) $q'(x', x) = q_j'(x_j', x_j|\alpha) P'(c', c) \dfrac{q_i'(x_i', x_i|\alpha)}{\sum_{y \in \mathcal{S}_i(c, x_i')} q_i'(x_i', y|\alpha)}$

(e) $q'(x', x) = 0$ otherwise.

Now try to verify the conditions of Theorem 1.13; i.e., first show that $\pi(x)q(x, x') = \pi(x')q'(x', x)$ for all $x, x' \in \mathcal{S}$ (this requires a separate argument for each of the cases (a)-(d) above), and second, verify that

$$\sum_{x' \in \mathcal{S}} q(x, x') = \sum_{x' \in \mathcal{S}} q'(x, x').$$

One would expect that the statement of this theorem is true for acyclic (or feedforward) network, but how about networks with feedback that are allowed in this model? In fact, the flows in the feedback paths are no longer Poisson. For a detailed probabilistic argument of this result see Walrand [Wal83].

6. *Stability.* This is our first look at a stability result. The existence of a stationary distribution (and consequently of the product form solution) is guaranteed provided that the traffic intensity at each station is less than 1. That is,

$$\rho_i \equiv \frac{\sum_{c \in \mathcal{C}_i} \alpha(c)}{\mu_i} < 1 \quad \text{for } i = 1, \ldots, J.$$

That is, each node in isolation should have enough processing capacity in order to satisfy all of its effective arriving traffic (of rate $\sum_{c \in \mathcal{C}_i} \alpha(c)$).

7. *Examples of Q-R nodes.* Kelly's Theorem says that a network of quasi-reversible nodes admits a product form solution. So far, however, we have not specified an example of a Q-R node.

The setup of section I.4 requires essentially that all job classes served at a node have the same exponential service distribution. With that restriction, the service discipline is quite general. It turns out (in homework?) that this node is indeed Q-R. The next section shows another example of a Q-R node.

I.6 Symmetric Queues

In this section we describe another example of a Q-R node, where the service time requirements are essentially arbitrary but the scheduling discipline is very specific (e.g., FIFO is rulled out).

1. Independent Poisson inputs with rates $\{\nu(c), c \in \mathcal{C}\}$. Jobs of class c require $w(c)$ independent *stages* of service; the duration of each is exponential with mean $d(c)$. Each class has its own Erlang service time distribution, and we will generalize that to mixtures of Erlang distributions, which are in fact dense in the set of all distributions.

2. Let $c(l)$ be the class of the job in position l, and $u(l)$ the stage of service this job has reached, $1 \leq u(l) \leq w(c(l))$. Set $\mathbf{c}(l) = (c(l), u(l))$. The state of the node is $\mathbf{c} = (\mathbf{c}(1), \ldots, \mathbf{c}(n))$, where n is the number of jobs present. The queue discipline for a symmetric node is as follows:

   (a) Total service effort is supplied at rate $\phi(n)$

   (b) Fraction $\gamma(l, n)$ of that effort goes to job in position l; $\gamma(1, n) + \cdots + \gamma(n, n) = 1$. When a job in position l completes his final stage and departs, jobs in positions $l+1, \ldots, n$ move to positions $l, \ldots, n-1$ respectively.

   (c) When a new job arrives, bringing the total population to $n+1$, he moves into position l with probability $\gamma(l, n+1)$; jobs previously in positions $l, \ldots, n$ move to $l+1, \ldots, n+1$.

   This specializes the discipline of section I.4 by requiring that $\gamma = \delta$ (hence the name symmetric). Again emphasize that the positioning of new arrivals and division of effort **cannot depend on class**.

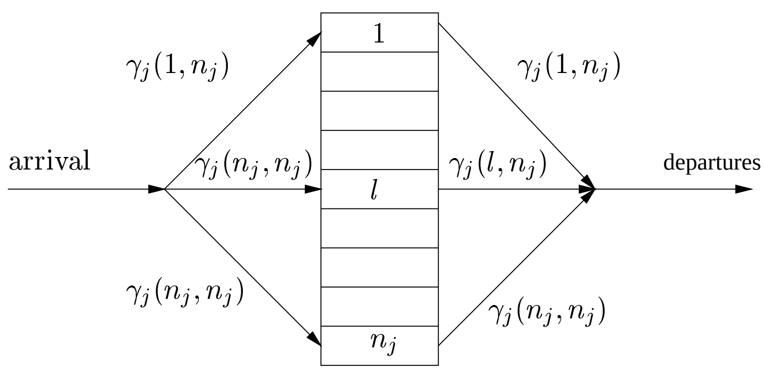

Figure 8: Symmetric queue

3. *Examples.* Think of total service requirement of a job as $S = S_1 + \cdots + S_{w(c)}$ (a random variable). The job departs when the cumulative service effort provided reaches S. Examples of disciplines that can be modeled in this framework are:

   (a) Stack (LIFO, preemptive-resume)

   (b) Server-sharing

(c) Delay node (∞-server queue)

(d) Multi-server node with no waiting room ($M/G/S/S$) – pure Loss system.

4. *Equilibrium distribution.* For a specification of the equilibrium distribution, first define

$$
\begin{aligned}
a(c) &= w(c)d(c) \qquad \forall\, c \in \mathcal{C} \\
&= \text{avg. service requirement (total service effort required) by jobs of class } c,
\end{aligned}
$$

$$
\begin{aligned}
a &= \sum_{c\in\mathcal{C}} \nu(c)a(c) = \text{ ``total workload,''} \\
b &= \left(\sum_{n=0}^{\infty} \prod_{l=1}^{n} \frac{a}{\phi(l)} \right)^{-1}.
\end{aligned}
$$

Remember $\mathbf{c} = (\mathbf{c}(1), \ldots, \mathbf{c}(n))$ and $\mathbf{c}(l) = (c(l), u(l))$. Then,

$$\pi(\mathbf{c}) = b \left(\prod_{l=1}^{n} \frac{a}{\phi(l)} \right) \left(\prod_{l=1}^{n} \frac{\nu(c(l))a(c(l))}{a} \frac{1}{w(c(l))} \right).$$

Also,

$$P\left(n(t) = n\right) = b \prod_{l=1}^{n} \frac{a}{\phi(l)}.$$

This is the same as if there were a single job class with exponential service time distribution (mean m) Poisson input (rate ν) and $\nu m = a$. The stationary distribution of n depends only on the total workload a and the service rate function $\phi(\cdot)$.

Given $n(t) = n$, the descriptions of (c, u) of the jobs in positions $1, \ldots, n$ are independent. The probability that a given job is of class c is $\nu(c)a(c)/a$, and given c, the various stage descriptions $u = 1, \ldots, w(c)$ are equally likely.

5. *FIFO.* First note that the equilibrium distribution described above depends on $w(c)$ and $d(c)$ only through their product $a(c)$ (i.e., the total workload).

   Consider now a single server queue, with Poisson arrivals, under the FIFO discipline, each job requiring k successive exponential service stages. This is the $M/E_k/1$ queue (E_k is the Erlang with k stages). It can be shown that the stationary distribution of this queue depends both on the mean and the variance of the service requirement, which is not the same as in the case of the symmetric queue. Why? Well, the reason is that FIFO cannot be put into a symmetric queue format: if we need to be processing the first job in the queue, it must be that $\gamma(1, n) = 1$, but this implies that new arrivals are placed at the head of the queue (and not the end as FIFO requires).

6. *Mixture of Eralng distributions.* Now suppose that each arrival of class c is assigned a *refined class descignation* (c, z), $z \in \mathcal{Z}$, with probability $p(c, z)$, and then the service requirement is Erlang with $w(c, z)$ stages and mean requirement of $d(c, z)$ per stage. Random thinning of

Poisson streams gives independent Poisson streams, so this situation maps into that already described. Let

$$\begin{aligned} a(c) &= \sum_{z \in \mathcal{Z}} p(c,z) w(c,z) d(c,z), \\ a &= \sum_{c \in \mathcal{C}} \nu(c) a(c). \end{aligned}$$

Summing up the equilibrium distributions of the refined state descriptions, we find that

- Marginal distribution of n is as before
- Given n, gross class designations $c(1), \ldots, c(n)$are independent, and the probability that any given job is of (gross) class c is $\nu(c)a(c)/a$, as before.
- Given the (gross) class designation c of a job in any position, the conditional distribution of total service effort received thus far by that job is

  $$F_c^*(x) = \frac{1}{a(c)} \int_0^x [1 - F_c(y)] dy,$$

  where F_c is the the distribution of total service requirement for jobs of (gross) class c (a mixture of Erlang distributions), and F_c^* is the *residual lifetime distribution* – familiar from renewal theory.
- Method of stages to approximate arbitrary distributions; i.e., general distributions are approximated as mixtures of Erlang distributions that in turn are just sums of exponentials.

7. Kelly proves that π is the equilibrium distribution by guessing the structure of the reversed process (the obvious one) and verifying as usual, using Theorem 1.13. Since the time-reversed system has Poisson input streams, independent of current state, the node is $Q - R$.

   Now we see that very general networks of $Q - R$ nodes can be built up. Poisson inputs and nodes of various types, including

   - FIFO, multi-server, and all classes served at node have ientical exponential service time distribution.
   - Processor sharing node where each class has its own arbitrary service distribution.
   - A delay node where each class has its own arbitrary delay distribution (jobs of a given type may visit many times with different distributions).

I.7 Closed Networks

1. *Closed network description.* Now there are no external arrivals, and we assume that the switching matrix $P = [P(k,l)]$ is stochastic (row sums equal to one). There is an initial population of size N, and these jobs circulate endlessly through the class designations $k = 1, \ldots, K$. As before we have J service stations and each class k is served at a unique station $s(k)$. Each station is assumed to be quasi-reversible in equilibrium, and all the notation of sections I.4-I.6 remains in

force. In particular, $\pi_j(\cdot, \alpha)$ denotes the stationary distribution for station j when it is driven by independent Poisson inputs with arrival rates $\{\alpha(k), k \in \mathcal{C}(j)\}$.

To avoid trivialities, we assume that the switching matrix P to be such that no job class k is transient. Thus the job classes can be partitioned into "communicating classes" -or "cells"- and renumbered so that P has block-diagonal structure (each block representing an irreducible communicating class or cell). That is, jobs whose initial class designation belongs to a particular cell can never visit any class outside that cell, and all classes within a cell communicate. We denote by $M(r)$ the initial population of cell r, so $M(1) + \cdots M(R) = N$, where R is the # of cells.

2. *Stationary distribution.* To describe the stationary distribution of the closed network we need a positive vector $\alpha = (\alpha(1), \ldots, \alpha(k))$ satisfying the traffic equations

$$\alpha(l) = \sum_{k=1}^{K} \alpha(k) P(k, l) \quad \text{for } l = 1, \ldots, K. \tag{I.16}$$

The solution of (I.16) is defined only up to a set of R positive constant (one for each cell); the α's for any given cell can be thought of as an arbitrary positive rescaling of the equilibrium distribution for that cell. We assume there exists a choice of the scaling constants such that the stationary distribution $\pi_j(\cdot, |\alpha)$ exists for each station j of the network. That choice of α is considered fixed hereafter.

Let us denote by $\mathcal{S}$ the set of all state descriptions $x = (x_1, \ldots, x_J)$ for the network. Then,

$$\sum_{j=1}^{J} \sum_{c \in C_r} n_j(c, x_j) = M(r) \quad \text{for all } r = 1, \ldots, R, \tag{I.17}$$

where C_r is the set of job classes belonging to cell r. This is the state space for our multi-cell closed queueing network.

3. *Main result.*

Theorem 7 *The unique equilibrium distribution for the closed network is*

$$\pi(x) = B(\alpha, M(1), \ldots, M(R)) \prod_{j=1}^{J} \pi_j(x_j|\alpha), \quad x \in \mathcal{S},$$

where $B(\cdot)$ is the normalization constant that makes $\pi(\cdot)$ a probability distribution on $\mathcal{S}$.

This theorem is proves as usual by "guessing" the time-reversed transition intensity function $q'(x', x)$ and the using Theorem 1.13.

- The stationary distribution is that of an open network model restricted to the hyperplanes of (I.17).
- A remarkable (that is, not at all obvious) "corollary" of this theorem is that the equilibrium distribution $\pi(\cdot)$ ultimately obtained is unaffected by the R arbitrary scaling constants in our choice of α; the effect of those scaling constants is cancelled out by the normalization constant $B(\cdot)$.

4. *Examples.*

- The system has one CPU (shared server) and N users. Each user alternate between "thinking" time (randomly distributed) at the client and "processing" time at the CPU or server.

$$R_i = \text{mean think time for user } i \text{ general distribution, and}$$
$$S_i = \text{mean processing time (service time) for tasks generated by user } i.$$

For details and extensions on this model see section 4.3 in Kelly [Kel79].

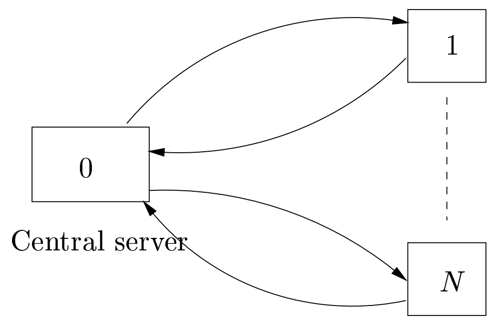

Figure 9: Timesharing computer system

- This is the classical "machine interference problem." There are N unreliable machines under the care of a single operator. Machines have exponential running times (i.e., times to failure) with mean R, and the operator repairs the failed machines on a FIFO basis with exponential service times with mean S. Let n_0 be the # of failed machines (queued at the operator) and n_1 be the # of operating machines; $n_1 = N - n_0$.

  Modeling the system as a closed network one can see which of the assumptions are really necessary and exploit possible generalizations; see section 4.2 in Kelly [Kel79]. For example,

  (a) The running distribution of each machine does not have to be exponential – a general distribution will still lead to a quasi-reversible node.

  (b) Multiple repairmen can be included, but still with a common exponential repair time distribution.

  (c) Each machine can have its own running time distribution with mean R_i.

  (d) Different stages of machine usage with different failure modes can be modeled; e.g., a recently repaired machine is likely to require a "tune-up" shortly thereafter and then take a much longer time until it completely fails again.

  It should be apparent that the one assumption that is harder to relax is that of exponential repair time distribution at the operator. (Why? What is the structure of the repair node?)

I.8 Notes and References

Basic references. A good reference for Markov chains is the first half of Karlin's book [Kar75]. The main reference for the material covered here is the book by Kelly [Kel79]. Historically, the

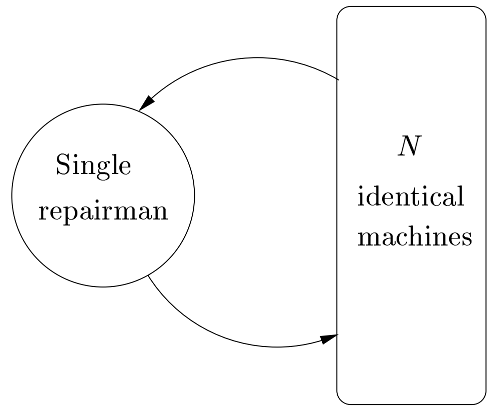

Figure 10: Single repairman

first important results in the area of product form networks are due to Jackson [Jac57, Jac63]; the motivation there was in manufacturing systems. Later on, important references can be found in the work of Kleinrock [Kle64, Kle75, Kle76], where the motivation was first in communication networks and later, in computer systems. Another major reference on this material, motivated by computer system applications, is the paper by Baskett, Chandy, Muntz, and Palacios [BCMP75]. Around the same time Kelly introduced the concept of quasi-reversibility. A probabilistic argument explaining the nature of product form results can be found in Walrand [Wal83]. A good exposition of the material can be found in the book by Walrand [Wal88].

Appliations. The paper by Solberg [Sol77] describes the applications of closed queueing network models in manufacturing systems. The "BCMP" paper [BCMP75] describes applications in computer systems. An extensive overview of applications of queueing network models to manufacturing systems can be found in the review article by Buzacott and Shanthikumar [BS92] and in the book by Gershwin [Ger94].

Current status. The network models considered here have been generalized in several directions, but some of the restrictive probabilistic and structural assumptions are still in place. An updated reference on this area is the forthcoming book by Chao, Miyazawa, and Pinedo [CMP99] that describes several extensions to the class of models discussed here that still fall in the category of "product form" networks.

References

[BCMP75] F. Baskett, M. Chandy, R. Muntz, and F. Palacios. Open, closed and mixed networks of queues of different classes of customers. *Journal of the Association for Computing Machinery*, 22:248–260, April 1975.

[BG92] D. Bertsekas and R. Gallager. *Data Networks.* Prentice-Hall, Englewood Cliffs, N.J., 1992.

[BS92] J. A. Buzacott and J. G. Shanthikumar. Design of manufacturing models using queueing models. *Queueing Systems*, 12:135–214, 1992.

[CMP99] X. Chao, M. Miyazawa, and M. Pinedo. *Queueing networks: Customers, Signals, and Product Form Solutions*. John Wiley & Sons, 1999. Forthcoming.

[Ger94] S. B. Gershwin. *Manufacturing Systems Engineering*. Prentice Hall, 1994.

[Jac57] J. R. Jackson. Networks of waiting lines. *Math. Oper. Res.*, 5:518–521, 1957.

[Jac63] J. R. Jackson. Jobshop-like queueing systems. *Management Science*, 10:131–142, 1963.

[Kar75] S. Karlin. *A first course in stochastic processes*. Academic Press, New York, N.Y., 1975.

[Kel79] F. P. Kelly. *Reversibility and Stochastic Networks*. John Wiley & Sons, New York, N.Y., 1979.

[Kle64] L. Kleinrock. *Communication Nets: Stochastic Message Flow and Delay*. McGraw-Hill, New York, N.Y., 1964.

[Kle75] L. Kleinrock. *Queueing Systems*, volume 1. Wiley, New York, N.Y., 1975.

[Kle76] L. Kleinrock. *Queueing Systems*, volume 2. Wiley, New York, N.Y., 1976.

[Sol77] J. J. Solberg. A mathematical model of computerized manufacturing systems. In *4th International Conference on Production Research, Tokyo*, 1977.

[Wal83] J. Walrand. A probabilistic look at networks of quasi-reversible queues. *IEEE Trans. Information Theory*, 29(6):825–831, November 1983.

[Wal88] J. Walrand. *An introduction to Queueing Networks*. Prentice-Hall, Englewood Cliffs, N.J., 1988.

I.9 Exercises – Homework #1

1. Consider an open Jackson network (one-to-one correspondence between job classes and service stations) with J stations, exponential service time distribution at each station, FIFO queue discipline at each station, and possibly multiple servers at each station. We say that the *routing chain is reversible* if the matrix $P'(i,j)$, defined in the lecture notes in section I.4 is identical to $P(i,j)$.

   (a) Show that if the routing chain is reversible in equilibrium, then the vector queue length process is reversible in equilibrium.

   (b) Suppose that the routing matrix satisfies $P(i,1) = 1$ for $i > 1$ and all external arrivals are into station 1. (This is common in modelling flexible manufacturing systems, where node 1 ia an infinite-server station representing a recirculating conveyor that carries workpieces on carts from one operation to another.) Show that the routing chain is reversible in equilibrium.

   (c) Suppose that the rouitng chain has the property described in (b). Now consider a modified system where each station $i = 2, \ldots, J$ has finite capacity $b(i)$. In terms of the queue length process, this is interpreted to mean simply that movement of jobs from station 1 to station i is eliminated when the total queue length at station i is $b(i)$. What is the equilibrium distribution for this modified system?

2. Consider an $M/M/1$ queue with arrival rate λ, mean service time $1/m$ and last-in-first-out preemptive-resume discipline: if a new job arrives while you are being served, that job immediately displaces you, and your service must begin all over when you next gain access to the server. Think of the service time of a job as a random variable S associated with the job rather than the server: once the job has occupied the server for S uninterrupted minutes it departs the system. How would you describe the state of the system in order to make it a Markov process? How would your answer differ if service time variability were associated entirely with the server rather than the job? (In the latter case, the time required to complete service of any given job is exponentially distributed with parameter m, independent of any previous experience of that job.)

3. Consider the open queueing network model described in section I.4. Specialize to the case where every station has a single server and FIFO service discipline, but then generalize by allowing jobs of class k to have an arbitrary service time distribution $F_k(x)$. Networks having this structure are called standard shops.

   Now consider the apparently more general case of a standard shop with server breakdown. To be specific, suppose that server j breaks down after providing an exponentially distributed amount of service: the failure rate (reciprocal of the mean time to failure) is b_j when server j is working. Server breakdown during a period of idleness is impossible.

   A job who is being served at the time of a breakdown at station j is not harmed; repair takes a random length of time with distribution $G_j(y)$, and then the service resumes where it left off. The service may be interrupted repeatedly by breakdowns.

(a) Under what conditions will the system have an equilibrium distribution? Assuming these conditions are met, what is the long-run fraction of time that server j spends working, being repaired, and simply idle? (You need not justify your answers.)

(b) Explain how a standard shop with server breakdown can be reduced to one without breakdowns. What are the data of the equivalent system?

4. Consider a service node of the type described in section I.4. Show that such a node is quasi-reversible.

5. In Figure 11, there are 4 stations A,B,C,D and two types of jobs 1 and 2. Type 1 jobs visit station A, then B, then C, then leave the system with probability p and with probability $1-p$ visit station D and then exit. Type 2 jobs visit A, then B, then D and exit. All stations use FIFO. The arrivals of type i jobs are Poisson with rate λ_i, $i=1,2$. Each station is a single server queue with exponential service times. The service rates are respectively μ_A, μ_B and μ_C. Station D's rate dependes on the type: it is $\mu_{D,1}$ for type 1 and $\mu_{D,2}$ for type 2.

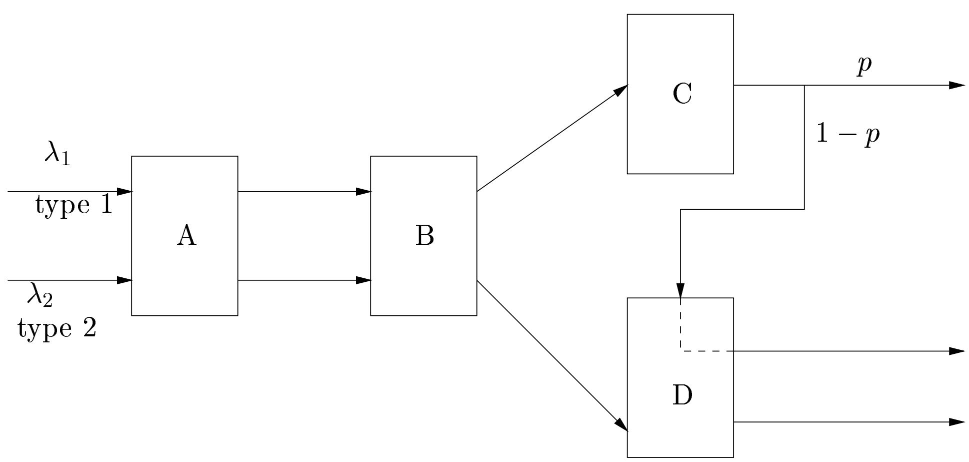

Figure 11: Four station network of ex.5

(a) Is the joint distribution of the number of jobs in stations A,B,C of product form? If so find it.

(b) Is the joint distribution of the number of jobs in stations A, B, C, D of product form? If so find it.

(c) Find the distribution of jobs in station D.

(d) Find the total expected time type 1 jobs spent in the system