Notes for ECE 467
Communication Network Analysis

Bruce Hajek

December 15, 2006

Contents

Preface

This is the latest draft of notes I have used for the graduate course *Communication Network Analysis*, offered by the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign. The notes describe many of the most popular analytical techniques for design and analysis of computer communication networks, with an emphasis on performance issues such as delay, blocking, and resource allocation. Topics that are *not* covered in the notes include the Internet protocols (at least not explicitly), simulation techniques and simulation packages, and some of the mathematical proofs. These are covered in other books and courses.

The topics of these notes form a basis for understanding the literature on performance issues in networks, including the Internet. Specific topics include

- The basic and intermediate theory of queueing systems, along with stability criteria based on drift analysis and fluid models
- The notion of effective bandwidth, in which a constant bit rate equivalent is given for a bursty data stream in a given context
- An introduction to the calculus of deterministic constraints on traffic flows
- The use of penalty and barrier functions in optimization, and the natural extension to the use of utility functions and prices in the formulation of dynamic routing and congestion control problems
- Some topics related to performance analysis in wireless networks, including coverage of basic multiple access techniques, and transmission scheduling
- The basics of dynamic programming, introduced in the context of a simple queueing control problem
- The analysis of blocking and the reduced load fixed point approximation for circuit switched networks.

Students are assumed to have already had a course on computer communication networks, although the material in such a course is more to provide motivation for the material in these notes, than to provide understanding of the mathematics. In addition, since probability is used extensively, students in the class are assumed to have previously had two courses in probability. Some prior exposure to the theory of Lagrange multipliers for constrained optimization and nonlinear optimization algorithms is desirable, but not necessary.

I'm grateful to students and colleagues for suggestions and corrections, and am always eager for more. Bruce Hajek, December 2006

Chapter 1

Countable State Markov Processes

1.1 Example of a Markov model

Consider a two-stage pipeline as pictured in Figure 1.1. Some assumptions about it will be made in order to model it as a simple discrete time Markov process, without any pretension of modeling a particular real life system. Each stage has a single buffer. Normalize time so that in one unit of time a packet can make a single transition. Call the time interval between k and $k+1$ the kth "time slot," and assume that the pipeline evolves in the following way during a given slot.

If at the beginning of the slot, there are no packets in stage one, then a new packet arrives to stage one with probability a, independently of the past history of the pipeline and of the outcome at state two.

If at the beginning of the slot, there is a packet in stage one and no packet in stage two, then the packet is transfered to stage two with probability d_1.

If at the beginning of the slot, there is a packet in stage two, then the packet departs from the stage and leaves the system with probability d_2, independently of the state or outcome of stage one.

These assumptions lead us to model the pipeline as a discrete-time Markov process with the state space $\mathcal{S} = \{00, 01, 10, 11\}$, transition probability diagram shown in Figure 1.2 (using the notation $\bar{x} = 1 - x$) and one-step transition probability matrix P given by

$$P = \begin{pmatrix} \bar{a} & 0 & a & 0 \\ \bar{a}d_2 & \bar{a}\bar{d}_2 & ad_2 & a\bar{d}_2 \\ 0 & d_1 & \bar{d}_1 & 0 \\ 0 & 0 & d_2 & \bar{d}_2 \end{pmatrix}$$

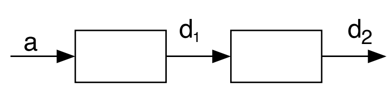

Figure 1.1: A two-stage pipeline

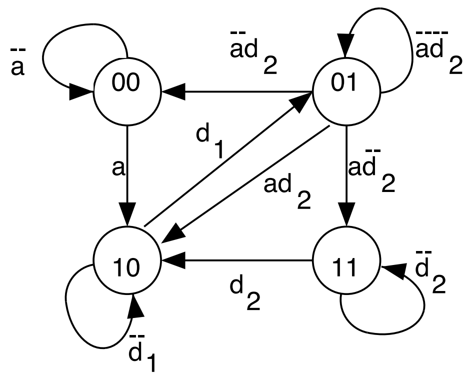

Figure 1.2: One-step transition probability diagram for example.

The rows of P are probability vectors. (In these notes, probability vectors are always taken to be row vectors, and more often than not, they are referred to as probability distributions.). For example, the first row is the probability distribution of the state at the end of a slot, given that the state is 00 at the beginning of a slot. Now that the model is specified, let us determine the throughput rate of the pipeline.

The equilibrium probability distribution $\pi = (\pi_{00}, \pi_{01}, \pi_{10}, \pi_{11})$ is the probability vector satisfying the linear equation $\pi = \pi P$. Once π is found, the throughput rate η can be computed as follows. It is defined to be the rate (averaged over a long time) that packets transit the pipeline. Since at most two packets can be in the pipeline at a time, the following three quantities are all clearly the same, and can be taken to be the throughput rate.

The rate of arrivals to stage one

The rate of departures from stage one (or rate of arrivals to stage two)

The rate of departures from stage two

Focus on the first of these three quantities. Equating long term averages with statistical averages yields

$$\begin{aligned}
\eta &= P[\text{an arrival at stage 1}] \\
&= P[\text{an arrival at stage 1}|\text{stage 1 empty at slot beginning}]P[\text{stage 1 empty at slot beginning}] \\
&= a(\pi_{00} + \pi_{01}).
\end{aligned}$$

Similarly, by focusing on departures from stage 1, obtain $\eta = d_1\pi_{10}$. Finally, by focusing on departures from stage 2, obtain $\eta = d_2(\pi_{01} + \pi_{11})$. These three expressions for η must agree.

Consider the numerical example $a = d_1 = d_2 = 0.5$. The equation $\pi = \pi P$ yields that π is proportional to the vector $(1, 2, 3, 1)$. Applying the fact that π is a probability distribution yields that $\pi = (1/7, 2/7, 3/7, 1/7)$. Therefore $\eta = 3/14 = 0.214\ldots$.

By way of comparison, consider another system with only a single stage, containing a single buffer. In each slot, if the buffer is empty at the beginning of a slot an arrival occurs with probability a, and if the buffer has a packet at the beginning of a slot it departs with probability d. Simultaneous arrival and departure is not allowed. Then $\mathcal{S} = \{0, 1\}$, $\pi = (d/(a+d), a/(a+d))$ and the throughput rate is $ad/(a+d)$. The two-stage pipeline with $d_2 = 1$ is essentially the same as the one-stage system. In case $a = d = 0.5$, the throughput rate of the single stage system is 0.25, which as expected is somewhat greater than that of the two-stage pipeline.

1.2 Definition, Notation and Properties

Having given an example of a discrete state Markov process, we now digress and give the formal definitions and some of the properties of Markov processes. Let $\mathbb{T}$ be a subset of the real numbers $\mathbb{R}$ and let $\mathcal{S}$ be a finite or countably infinite set. A collection of $\mathcal{S}$–valued random variables $(X(t) : t \in \mathbb{T})$ is a *discrete-state Markov process* with state space $\mathcal{S}$ if

$$P[X(t_{n+1}) = i_{n+1}|X(t_n) = i_n, \ . \ . \ . \ , X(t_1) = i_1] = P[X(t_{n+1}) = i_{n+1}|X(t_n) = i_n] \tag{1.1}$$

whenever

$$\begin{cases} t_1 < t_2 < \ . \ . \ . \ < t_{n+1} \text{ are in } \mathbb{T}, \\ i_i, i_2, ..., i_{n+1} \text{ are in } \mathcal{S}, \text{ and} \\ P[X(t_n) = i_n, \ . \ . \ . \ , X(t_1) = i_1] > 0. \end{cases} \tag{1.2}$$

Set $p_{ij}(s,t) = P[X(t) = j|X(s) = i]$ and $\pi_i(t) = P[X(t) = i]$. The probability distribution $\pi(t) = (\pi_i(t) : i \in \mathcal{S})$ should be thought of as a row vector, and can be written as one once $\mathcal{S}$ is ordered. Similarly, $H(s,t)$ defined by $H(s,t) = (p_{ij}(s,t) : i, j \in \mathcal{S})$ should be thought of as a matrix. Let e denote the column vector with all ones, indexed by $\mathcal{S}$. Since $\pi(t)$ and the rows of $H(s,t)$ are probability vectors for $s, t \in \mathbb{T}$ and $s \le t$, it follows that $\pi(t)e = 1$ and, $H(s,t)e = e$.

Next observe that the marginal distributions $\pi(t)$ and the transition probabilities $p_{ij}(s,t)$ determine all the finite dimensional distributions of the Markov process. Indeed, given

$$\begin{cases} t_1 < t_2 < \ . \ . \ . \ < t_n \ \text{ in } \mathbb{T}, \\ i_i, i_2, ..., i_n \ \in \mathcal{S} \end{cases} \tag{1.3}$$

one writes

$$\begin{aligned} &P[X(t_1) = i_1, \ldots, X(t_n) = i_n] = \\ &\qquad P[X(t_1) = i_1, \ldots, X(t_{n-1}) = i_{n-1}]P[X(t_n) = i_n|X(t_1) = i_1, \ldots, X(t_{n-1}) = i_{n-1}] \\ &= \quad P[X(t_1) = i_1, \ldots, X(t_{n-1}) = i_{n-1}]p_{i_{n-1}i_n}(t_{n-1}, t_n) \end{aligned}$$

Application of this operation $n - 2$ more times yields that

$$P[X(t_1) = i_1, X(t_2) = i_2, \ . \ . \ . \ , X(t_n) = i_n] = \pi_{i_1}(t_1)p_{i_1i_2}(t_1, t_2) \ . \ . \ . \ p_{i_{n-1}i_n}(t_{n-1}, t_n), \tag{1.4}$$

which shows that the finite dimensional distributions of X are indeed determined by $(\pi(t))$ and $(p_{ij}(s,t))$. From this and the definition of conditional probabilities it follows by straight substitution that

$$\begin{aligned} &P[X(t_j) = i_j, \text{ for } 1 \le j \le n + l|X(t_n) = i_n] = \\ &\quad P[X(t_j) = i_j, \text{ for } 1 \le j \le n|X(t_n) = i_n]P[X(t_j) = i_j, \text{ for } n \le j \le n + l|X(t_n) = i_n] \end{aligned} \tag{1.5}$$

whenever $P[X(t_n) = i_n] > 0$. Property (1.5) is equivalent to the Markov property. Note in addition that it has no preferred direction of time, simply stating that the past and future are conditionally independent given the present. It follows that if X is a Markov process, the time reversal of X defined by $\tilde{X}(t) = X(-t)$ is also a Markov process.

A Markov process is *time homogeneous* if $p_{ij}(s,t)$ depends on s and t only through $t-s$. In that case we write $p_{ij}(t-s)$ instead of $p_{ij}(s,t)$, and $H_{ij}(t-s)$ instead of $H_{ij}(s,t)$.

Recall that a random process is *stationary* if its finite dimensional distributions are invariant with respect to translation in time. Referring to (1.4), we see that a time-homogeneous Markov process is stationary if and only if its one dimensional distributions $\pi(t)$ do not depend on t. If, in our example of a two-stage pipeline, it is assumed that the pipeline is empty at time zero and that $a \neq 0$, then the process is not stationary (since $\pi(0) = (1,0,0,0) \neq \pi(1) = (1-a,0,a,0)$), even though it is time homogeneous. On the other hand, a Markov random process that is stationary is time homogeneous.

Computing the distribution of $X(t)$ by conditioning on the value of $X(s)$ yields that $\pi_j(t) = \sum_i P[X(s)=i, X(t)=j] = \sum_i \pi_i(s)p_{ij}(s,t)$, which in matrix form yields that $\pi(t) = \pi(s)H(s,t)$ for $s,t \in \mathbb{T}, s \leq t$. Similarly, given $s < \tau < t$, computing the conditional distribution of $X(t)$ given $X(s)$ by conditioning on the value of $X(\tau)$ yields

$$H(s,t) = H(s,\tau)H(\tau,t) \qquad s,\tau,t \in \mathbb{T},\ s < \tau < t. \tag{1.6}$$

The relations (1.6) are known as the Chapman-Kolmogorov equations.

If the Markov process is time-homogeneous, then $\pi(s+\tau) = \pi(s)H(\tau)$ for $s, s+\tau \in \mathbb{T}$ and $\tau \geq 0$. A probability distribution π is called an *equilibrium* (or invariant) distribution if $\pi H(\tau) = \pi$ for all $\tau \geq 0$.

Repeated application of the Chapman-Kolmogorov equations yields that $p_{ij}(s,t)$ can be expressed in terms of transition probabilities for s and t close together. For example, consider Markov processes with index set the integers. Then $H(n,k+1) = H(n,k)P(k)$ for $n \leq k$, where $P(k) = H(k,k+1)$ is the one-step transition probability matrix. Fixing n and using forward recursion starting with $H(n,n) = I$, $H(n,n+1) = P(n)$, $H(n,n+2) = P(n)P(n+1)$, and so forth yields

$$H(n,l) = P(n)P(n+1)\cdots P(l-1)$$

In particular, if the process is time-homogeneous then $H(k) = P^k$ for all k for some matrix P, and $\pi(l) = P^{l-k}\pi(k)$ for $l \geq k$. In this case a probability distribution π is an equilibrium distribution if and only if $\pi P = \pi$.

In the next section, processes indexed by the real line are considered. Such a process can be described in terms of $p(s,t)$ with $t-s$ arbitrarily small. By saving only a linearization, the concept of generator matrix arises naturally.

1.3 Pure-Jump, Time-Homogeneous Markov Processes

Let $\mathcal{S}$ be a finite or countably infinite set, and let $\triangle \notin \mathcal{S}$. A *pure-jump function* is a function $x : \mathbb{R}_+ \rightarrow \mathcal{S} \cup \{\triangle\}$ such that there is a sequence of times, $0 = \tau_0 < \tau_1 < \ \ldots$, and a sequence of states, $s_0, s_1, \ \ldots$ with $s_i \in \mathcal{S}$, and $s_i \neq s_{i+1}$, $i \geq 0$, so that

$$x(t) = \begin{cases} s_i & \text{if } \tau_i \leq t < \tau_{i+1} \quad i \geq 0 \\ \triangle & \text{if } t \geq \tau^* \end{cases} \tag{1.7}$$

where $\tau^* = \lim_{i\to\infty} \tau_i$. If τ^* is finite it is said to be the explosion time of the function x. The example corresponding to $\mathcal{S} = \{0,1,\ldots\}$, $\tau_i = i/(i+1)$ and $s_i = i$ is pictured in Fig. 1.3. Note

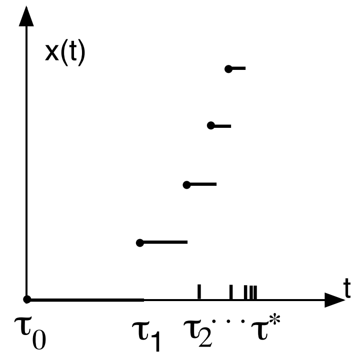

Figure 1.3: Sample pure-jump function with an explosion time

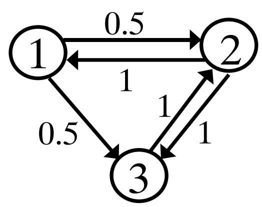

Figure 1.4: Transition rate diagram for a continuous time Markov process

that $\tau^* = 1$ for this example. A *pure-jump Markov process* $(X_t : t \geq 0)$ is a Markov process such that, with probability one, its sample paths are pure-jump functions.

Let $Q = (q_{ij} : i, j \in \mathcal{S})$ be such that

$$\begin{array}{ll} q_{ij} \geq 0 & i, j \in \mathcal{S}, \quad i \neq j \\ q_{ii} = -\sum_{j \in \mathcal{S}, j \neq i} q_{ij} & i \in \mathcal{S}. \end{array} \tag{1.8}$$

An example for state space $\mathcal{S} = \{1, 2, 3\}$ is

$$Q = \begin{pmatrix} -1 & 0.5 & 0.5 \\ 1 & -2 & 1 \\ 0 & 1 & -1 \end{pmatrix},$$

which can be represented by the transition rate diagram shown in Figure 1.4. A pure-jump, time-homogeneous Markov process X has *generator matrix* Q if

$$\lim_{h \searrow 0} (p_{ij}(h) - I_{\{i=j\}})/h = q_{ij} \qquad i, j \in \mathcal{S} \tag{1.9}$$

or equivalently

$$p_{ij}(h) = I_{\{i=j\}} + hq_{ij} + o(h) \qquad i, j \in \mathcal{S} \tag{1.10}$$

where $o(h)$ represents a quantity such that $\lim_{h \to 0} o(h)/h = 0$.

For the example this means that the transition probability matrix for a time interval of duration h is given by

$$\begin{pmatrix} 1-h & 0.5h & 0.5h \\ h & 1-2h & h \\ 0 & h & 1-h \end{pmatrix} + \begin{pmatrix} o(h) & o(h) & o(h) \\ o(h) & o(h) & o(h) \\ o(h) & o(h) & o(h) \end{pmatrix}$$

The first term is a stochastic matrix, owing to the assumptions on the generator matrix Q.

Proposition 1.3.1 *Given a matrix Q satisfying (1.8), and a probability distribution $\pi(0) = (\pi_i(0) : i \in \mathcal{S})$, there is a pure-jump, time-homogeneous Markov process with generator matrix Q and initial distribution $\pi(0)$. The finite-dimensional distributions of the process are uniquely determined by $\pi(0)$ and Q.*

The proposition can be proved by appealing to the space-time properties in the next section. In some cases it can also be proved by considering the forward-differential evolution equations for $\pi(t)$, which are derived next. Fix $t > 0$ and let h be a small positive number. The Chapman-Kolmogorov equations imply that

$$\frac{\pi_j(t+h) - \pi_j(t)}{h} = \sum_{i \in \mathcal{S}} \pi_i(t) \left(\frac{p_{ij}(h) - I_{\{i=j\}}}{h} \right). \tag{1.11}$$

Consider letting h tend to zero. *If* the limit in (1.9) is uniform in i for j fixed, then the limit and summation on the right side of (1.11) can be interchanged to yield the forward-differential evolution equation:

$$\frac{\partial \pi_j(t)}{\partial t} = \sum_{i \in \mathcal{S}} \pi_i(t) q_{ij} \tag{1.12}$$

or $\frac{\partial \pi(t)}{\partial t} = \pi(t)Q$. This equation, known as the Kolmogorov forward equation, can be rewritten as

$$\frac{\partial \pi_j(t)}{\partial t} = \sum_{i \in \mathcal{S}, i \neq j} \pi_i(t) q_{ij} - \sum_{i \in \mathcal{S}, i \neq j} \pi_j(t) q_{ji}, \tag{1.13}$$

which states that the rate change of the probability of being at state j is the rate of "probability flow" into state j minus the rate of probability flow out of state j.

1.4 Space-Time Structure

Let $(X_k : k \in \mathbb{Z}_+)$ be a time-homogeneous Markov process with one-step transition probability matrix P. Let T_k denote the time that elapses between the k^{th} and $k+1^{th}$ jumps of X, and let $X^J(k)$ denote the state after k jumps. See Fig. 1.5 for illustration. More precisely, the *holding times* are defined by

$$T_0 = \min\{t \geq 0 : X(t) \neq X(0)\} \tag{1.14}$$
$$T_k = \min\{t \geq 0 : X(T_0 + \ \ldots \ + T_{k-1} + t) \neq X(T_0 + \ \ldots \ + T_{k-1})\} \tag{1.15}$$

and the *jump process* $X^J = (X^J(k) : k \geq 0)$ is defined by

$$X^J(0) = X(0) \quad \text{and} \quad X^J(k) = X(T_0 + \ \ldots \ + T_{k-1}) \tag{1.16}$$

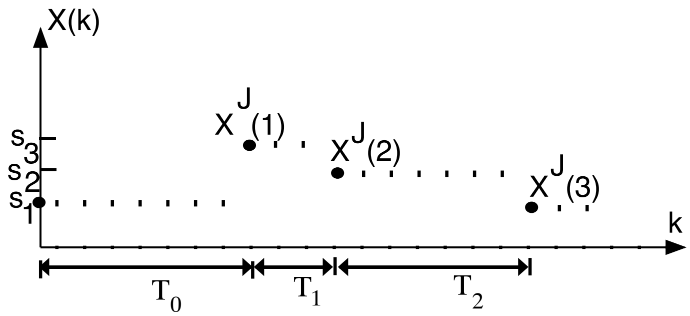

Figure 1.5: Illustration of jump process and holding times.

Clearly the holding times and jump process contain all the information needed to construct X, and vice versa. Thus, the following description of the joint distribution of the holding times and the jump process characterizes the distribution of X.

Proposition 1.4.1 *Let $X = (X(k) : k \in \mathbb{Z}_+)$ be a time-homogeneous Markov process with one-step transition probability matrix P.*

(a) The jump process X^J is itself a time-homogeneous Markov process, and its one-step transition probabilities are given by $p^J_{ij} = p_{ij}/(1-p_{ii})$ for $i \neq j$, and $p^J_{ii} = 0$, $i, j \in \mathcal{S}$.

(b) Given $X(0)$, $X^J(1)$ is conditionally independent of T_0.

(c) Given $(X^J(0), \ldots, X^J(n)) = (j_0, \ldots, j_n)$, the variables $T_0, \ \ldots, T_n$ are conditionally independent, and the conditional distribution of T_l is geometric with parameter $p_{j_l j_l}$:

$$P[T_l = k | X^J(0) = j_0, \ \ldots, X^J(n) = j_n] = p^{k-1}_{j_l j_l}(1 - p_{j_l j_l}) \quad 0 \leq l \leq n, \ k \geq 1.$$

Proof. Observe that if $X(0) = i$, then

$$\{T_0 = k, X^J(1) = j\} = \{X(1) = i, X(2) = i, \ldots, X(k-1) = i, X(k) = j\},$$

so

$$P[T_0 = k, X^J(1) = j | X(0) = i] = p^{k-1}_{ii} p_{ij} = [(1-p_{ii})p^{k-1}_{ii}]p^J_{ij} \tag{1.17}$$

Because for i fixed the last expression in (1.17) displays the product of two probability distributions, conclude that *given* $X(0) = i$,

T_0 has distribution $((1-p_{ii})p^{k-1}_{ii} : k \geq 1)$, the geometric distribution of mean $1/(1-p_{ii})$

$X^J(1)$ has distribution $(p^J_{ij} : j \in \mathcal{S})$ (i fixed)

T_0 and $X^J(1)$ are independent

More generally, check that

$$P[X^J(1) = j_1, \ldots, X^J(n) = j_n, T_o = k_0, \ldots, T_n = k_n | X^J(0) = i] =$$
$$p^J_{ij_1} p^J_{j_1 j_2} \ldots p^J_{j_{n-1} j_n} \prod_{l=0}^{n} (p^{k_l - 1}_{j_l j_l}(1 - p_{j_l j_l}))$$

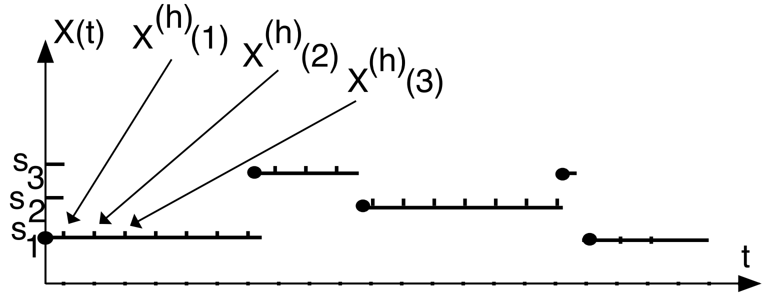

Figure 1.6: Illustration of sampling of a pure-jump function

This establishes the proposition. ■

Next we consider the space-time structure of time-homogeneous continuous-time pure-jump Markov processes. Essentially the only difference between the discrete- and continuous-time Markov processes is that the holding times for the continuous-time processes are exponentially distributed rather than geometrically distributed. Indeed, define the holding times $T_k, k \geq 0$ and the jump process X^J using (1.14)-(1.16) as before.

Proposition 1.4.2 *Let $X = (X(t) : t \in \mathbb{R}_+)$ be a time-homogeneous, pure-jump Markov process with generator matrix Q. Then*

(a) The jump process X^J is a discrete-time, time-homogeneous Markov process, and its one-step transition probabilities are given by

$$p^J_{ij} = \begin{cases} -q_{ij}/q_{ii} & \text{for } i \neq j \\ 0 & \text{for } i = j \end{cases} \tag{1.18}$$

(b) Given $X(0)$, $X^J(1)$ is conditionally independent of T_0.

(c) Given $X^J(0) = j_0, \ldots, X^J(n) = j_n$, the variables $T_0, \ldots, T_n$ are conditionally independent, and the conditional distribution of T_l is exponential with parameter $-q_{j_l j_l}$:

$$P[T_l \geq c | X^J(0) = j_0, \ldots, X^J(n) = j_n] = exp(cq_{j_l j_l}) \qquad 0 \leq l \leq n.$$

Proof. Fix $h > 0$ and define the "sampled" process $X^{(h)}$ by $X^{(h)}(k) = X(hk)$ for $k \geq 0$. See Fig. 1.6. Then $X^{(h)}$ is a discrete time Markov process with one-step transition probabilities $p_{ij}(h)$ (the transition probabilities for the original process for an interval of length h). Let $(T^{(h)}_k : k \geq 0)$ denote the sequence of holding times and $(X^{J,h}(k) : k \geq 0)$ the jump process for the process $X^{(h)}$.

The assumption that with probability one the sample paths of X are pure-jump functions, implies that *with probability one:*

$$\lim_{h \to 0} (X^{J,h}(0), X^{J,h}(1), \ldots, X^{J,h}(n), hT^{(h)}_0, hT^{(h)}_1, \ldots, hT^{(h)}_n) = (X^J(0), X^J(1), \ldots, X^J(n), T_0, T_1, \ldots, T_n) \tag{1.19}$$

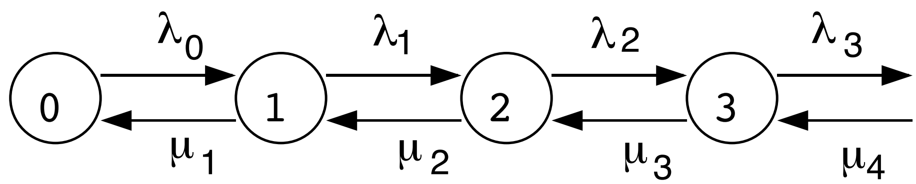

Figure 1.7: Transition rate diagram of a birth-death process

Since convergence with probability one implies convergence in distribution, the goal of identifying the distribution of the random vector on the righthand side of (1.19) can be accomplished by identifying the limit of the distribution of the vector on the left.

First, the limiting distribution of the process $X^{J,h}$ is identified. Since $X^{(h)}$ has one-step transition probabilities $p_{ij}(h)$, the formula for the jump process probabilities for discrete-time processes (see Proposition 1.4.1, part a) yields that the one step transition probabilities $p_{ij}^{J,h}$ for $X^{(J,h)}$ are given by

$$\begin{aligned} p_{ij}^{J,h} &= \frac{p_{ij}(h)}{1-p_{ii}(h)} \\ &= \frac{p_{ij}(h)/h}{(1-p_{ii}(h))/h} \to \frac{q_{ij}}{-q_{ii}} \quad \text{as } h \to 0 \end{aligned} \tag{1.20}$$

for $i \neq j$, where the limit indicated in (1.20) follows from the definition (1.9) of the generator matrix Q. Thus, the limiting distribution of $X^{J,h}$ is that of a Markov process with one-step transition probabilities given by (1.18), establishing part (a) of the proposition. The conditional independence properties stated in (b) and (c) of the proposition follow in the limit from the corresponding properties for the jump process $X^{J,h}$ guaranteed by Proposition 1.4.1. Finally, since $\log(1+\theta) = \theta + o(\theta)$ by Taylor's formula, we have for all $c \geq 0$ that

$$\begin{aligned} P[hT_l^{(h)} > c | X^{J,h}(0) = j_0, \ldots, X^{J,h} = j_n] &= (p_{j_l j_l}(h))^{\lfloor c/h \rfloor} \\ &= \exp(\lfloor c/h \rfloor \log(p_{j_l j_l}(h))) \\ &= \exp(\lfloor c/h \rfloor (q_{j_l j_l} h + o(h))) \\ &\to \exp(q_{j_l j_l} c) \quad \textit{as } h \to 0 \end{aligned}$$

which establishes the remaining part of (c), and the proposition is proved. ■

Birth-Death Processes A useful class of countable state Markov processes is the set of birth-death processes. A (continuous time) birth-death process with parameters $(\lambda_0, \lambda_2, \ldots)$ and $(\mu_1, \mu_2, \ldots)$ (also set $\lambda_{-1} = \mu_0 = 0$) is a pure-jump Markov process with state space $\mathcal{S} = \mathbb{Z}_+$ and generator matrix Q defined by $q_{kk+1} = \lambda_k, q_{kk} = -(\mu_k + \lambda_k)$, and $q_{kk-1} = \mu_k$ for $k \geq 0$, and $q_{ij} = 0$ if $|i-j| \geq 2$. The transition rate diagram is shown in Fig. (1.7). The space-time structure of such a process is as follows. Given the process is in state k at time t, the next state visited is $k+1$ with probability $\lambda_k/(\lambda_k + \mu_k)$ and $k-1$ with probability $\mu_k/(\lambda_k + \mu_k)$. The holding time of state k is exponential with parameter $\lambda_k + \mu_k$.

The space-time structure just described can be used to show that the limit in (1.9) is uniform in i for j fixed, so that the Kolmogorov forward equations are satisfied. These equations are:

$$\frac{\partial \pi_k(t)}{\partial t} = \lambda_{k-1}\pi_{k-1}(t) - (\lambda + \mu)\pi_k(t) + \mu_{k+1}\pi_{k+1}(t) \tag{1.21}$$

1.5 Poisson Processes

A *Poisson process with rate* λ is a birth-death process $N = (N(t) : t \geq 0)$ with initial distribution $P[N(0) = 0] = 1$, birth rates $\lambda_k = \lambda$ for all k and death rates $\mu_k = 0$ for all k. The space-time structure is particularly simple. The jump process N^J is deterministic and is given by $N^J(k) = k$. Therefore the holding times are not only conditionally independent given N^J, they are independent and each is exponentially distributed with parameter λ.

Let us calculate $\pi_j(t) = P[N(t) = j]$. The Kolmogorov forward equation for $k = 0$ is $\partial\pi_0/\partial t = -\lambda\pi_0$, from which we deduce that $\pi_0(t) = \exp(-\lambda t)$. Next the equation for π_1 is $\partial\pi_1/\partial t = \lambda\exp(-\lambda t) - \lambda\pi_1(t)$ which can be solved to yield $\pi_1(t) = (\lambda t)\exp(-\lambda t)$. Continuing by induction on k, verify that $\pi_k(t) = (\lambda t)^k \exp(-\lambda t)/k!$, so that $N(t)$ is a Poisson random variable with parameter k.

It is instructive to solve the Kolmogorov equations by another method, namely the z-transform method, since it works for some more complex Markov processes as well. For convenience, set $\pi_{-1}(t) = 0$ for all t. Then the Kolmogorov equations for π become $\frac{\partial \pi_k}{\partial t} = \lambda\pi_{k-1} - \lambda\pi_k$. Multiplying each side of this equation by z^k, summing over k, and interchanging the order of summation and differentiation, yields that the z transform $P^*(z,t)$ of $\pi(t)$ satisfies $\frac{\partial P^*(z,t)}{\partial t} = (\lambda z - \lambda)P^*(z,t)$. Solving this with the initial condition $P^*(z,0) = 1$ yields that $P^*(z,t) = \exp((\lambda z - \lambda)t)$. Expanding $\exp(\lambda z t)$ into powers of z identifies $(\lambda t)^k \exp(-\lambda t)/k!$ as the coefficient of z^n in $P^*(z,t)$.

In general, for i fixed, $p_{ij}(t)$ is determined by the same equations but with the initial distribution $p_{ij}(0) = I_{\{i=j\}}$. The resulting solution is $p_{ij}(t) = \pi_{j-i}(t)$ for $j \geq i$ and $p_{ij}(t) = 0$ otherwise. Thus for $t_0 < t_1 < \ldots < t_d = s < t$,

$$\begin{aligned}
&P[N(t) - N(s) = l | N(t_0) = i_0, \ldots, N(t_{d-1}) = i_{d-1}, N(s) = i] \\
&= P[N(t) = i + l | N(s) = i] \\
&= [\lambda(t-s)]^l \exp(-\lambda(t-s))/l!
\end{aligned}$$

Conclude that $N(t) - N(s)$ is a Poisson random variable with mean $\lambda(t-s)$. Furthermore, $N(t) - N(s)$ is independent of $(N(u) : u \leq s)$, which implies that the increments $N(t_1) - N(t_0), N(t_2) - N(t_1), N(t_3) - N(t_2), \ldots$ are mutually independent whenever $0 \leq t_0 < t_1 < \ldots$.

Turning to another characterization, fix $T > 0$ and $\lambda > 0$, let $U_1, U_2, \ldots$ be uniformly distributed on the interval $[0,T]$, and let K be a Poisson random variable with mean λT. Finally, define the random process $(\tilde{N}(t) : 0 \leq t \leq T)$ by

$$\tilde{N}(t) = \sum_{i=1}^{K} I_{\{t \geq U_i\}}. \tag{1.22}$$

That is, for $0 \leq t \leq T$, $\tilde{N}(t)$ is the number of the first K uniform random variables located in $[0,t]$. We claim that $\tilde{N}$ has the same distribution as a Poisson random process N with parameter λ, restricted to the interval $[0,T]$. To verify this claim, it suffices to check that the increments of the two