The role of the random-product technique in the theory of Markov chains on a countable state space

Brian Fralix
School of Mathematical and Statistical Sciences
Clemson University
Clemson, SC, USA
bfralix@clemson.edu

Farhad Hasankhani
Department of Industrial Engineering
Clemson University
Clemson, SC, USA
fhasank@clemson.edu

Amin Khademi
Department of Industrial Engineering
Clemson University
Clemson, SC, USA
khademi@clemson.edu

April 18, 2020

Abstract

We show how the random-product technique from [4] can be used to give a complete account of the construction of invariant measures associated with countable-state discrete-time and continuous-time Markov chains. Our hope is that this article will give the method more exposure within the probability, applied probability and stochastic operations research communities, and that it will prove to be useful to other researchers performing their own investigations. We also hope that instructors will find this material useful while teaching various types of courses in applied probability.

Keywords: Random product, Markov chains, stochastic networks, invariant measure.
2020 MSC: 60J10, 60J27, 60J28, 60K37

Contents

1 Introduction

Markov chains and processes constitute important classes of stochastic processes, as they possess just the right type of structure so that (i) they can be analyzed mathematically at a high level of sophistication, yet (ii) they can be used to model many types of random phenomena in many different fields of study. Entire courses that focus on these processes alone have been taught for many years in departments/schools of mathematics, statistics, industrial engineering, and computer science, with a theoretical and/or an applied focus at levels that can vary considerably, depending on the intended audience. Typically, calculating various performance measures associated with these processes requires only usage of undergraduate-level mathematical tools and constructs: stationary/limiting distributions can be found by solving a finite or infinite system of linear equations, while the time-dependent behavior of these processes can be analyzed through the study of a (particularly simple) finite or infinite system of linear difference equations in a discrete-time context, and a finite or infinite system of linear differential equations in the continuous-time context.

That being said, a proper analysis of a Markov process can be rather challenging: proving the existence of both a stationary distribution, as well as a unique limiting distribution can involve more sophisticated tools and ideas from differential equations, functional analysis, linear algebra, and other areas of classical probability theory (such as e.g. martingale theory). Even if one is able to discern whether or not an infinite-state Markov process has a unique stationary distribution, in many cases, it can be very difficult to determine whether or not the stationary distribution can be expressed in closed-form. For the case of Markovian stochastic networks, this simple fact has led, over the last fifty years, to a significant amount of research focusing on determining necessary and sufficient conditions for when the elements of the stationary distribution associated with the network exhibit a *product-form* (and, hence, 'closed-form') structure. Many of the ideas used in this line of research involve use of partial balance equations, and/or usage of the *time-reversal* of a Markov chain: searching for the time-reversal at the same time as a stationary distribution in an intelligent manner can sometimes lead to the discovery of a product-form stationary distribution. This material is well-known, and covered in many textbooks: see e.g. Kelly [16], Serfozo [22], and Chen and Yao [5].

Recently it was shown in [4] that each element of the stationary distribution of an ergodic continuous-time Markov chain (CTMC) can be expressed in terms of the expected value of a *random product*, whose behavior is governed by an alternative CTMC that must satisfy two essential properties connecting it to the original CTMC. There are, theoretically speaking, infinitely many ways to choose this alternative CTMC, with one possible choice being the time-reversal of the original CTMC: this choice yields random products that are actually deterministic, thanks to a well-known result often referred to as the generalized *Kolmogorov Criterion* (see [16]). What makes these random-product results appealing is that (i) typically the time-reversal is not known in advance, and (ii) using random products allows for the application of additional probabilistic techniques—such as first-step analysis—towards deriving the stationary distribution of an ergodic CTMC. In the interest of making our exposition smoother, throughout this manuscript we will refer to all past and present usage of the random-product representation as the *random-product technique*.

The random-product technique, to the best of our knowledge, first appeared in [4], and its usage up to now has primarily been geared towards deriving both stationary distributions and Laplace transforms of transition functions associated with special types of CTMCs. Papers featuring this type of analysis include [10] which features a study of a class of Markovian 'smart' polling models; [11] which focuses on the study of Markovian reentrant-line models; [19] which analyzes classical queueing models with Erlang-distributed interarrival or service distributions, and [14] which further studies a class of Markovian Bitcoin queueing models introduced in the work of Göbel et al. [13]. This technique was also used to a lesser degree in [21], which features an analysis of the time-dependent behavior of M/M/c preemptive-priority queueing models with two customer classes.

Up until now, the random-product technique has made minimal contributions to the actual theory behind discrete-time and continuous-time Markov chains, as it has mainly been used to derive performance measures for various examples of CTMCs. In [9], random products were used to verify a conjecture of Pollett [20] who presented what he correctly suspected were sufficient conditions for two CTMCs to be similar. One interesting by-product of this work is a result that shows, through a fairly elementary argument, how the Laplace transform of each transition function of a CTMC can be expressed in terms of the expected value of a modified random product, where the domain of the Laplace transform includes the open half-plane of complex numbers having positive real part. A special case of this result is contained in [4], but there the Laplace transforms are only defined on the positive portion of the real line. Another contribution of the random-product technique to the theory of Markov chains can be found in [15], where it was shown how the technique can be used to re-derive many classical results associated with what are often known as *Markov processes of G/M/1-type*. One especially appealing aspect of the random-product technique in this context is that it replaces arguments involving taboo probabilities with arguments that involve first-step analysis, which arguably makes the theory more accessible, as first-step analysis is usually taught in courses on Markov chains and processes at all levels.

This survey was written while having in mind the goal of illustrating how the random-product technique can be used to establish many classical results from the theory of both discrete-time Markov chains (DTMCs) and CTMCs, that are either directly or indirectly related to the problem of constructing invariant measures, as well as the problem of establishing both existence of an invariant measure, as well as the overall structure of the set of all invariant measures associated with a given DTMC or CTMC. Given this is a survey article, whenever we prove a result with the random-product technique, we will also compare our approach with past approaches. We hope—at a minimum—that instructors will find this material useful while teaching their own graduate (and possibly even undergraduate) courses on Markov chains, but we also feel that this is still a little-known technique, and it may very well be the case that as more researchers learn about it, more completely new results will be discovered through its use.

2 Discrete-time Markov Chains

Let $\{X_n\}_{n\geq 0}$ represent a DTMC defined on a fixed probability space $(\Omega, \mathcal{B}, \mathbb{P})$—all random elements encountered in this survey will be defined on this space—whose state space S is either finite or countably infinite. These assumptions imply that $\{X_n\}_{n\geq 0}$ satisfies the *Markov property*: for each integer $n \geq 0$, and each collection of states $x_0, x_1, \ldots, x_n, x_{n+1} \in S$,

$$\mathbb{P}(X_{n+1} = x_{n+1} \mid X_n = x_n, X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = \mathbb{P}(X_{n+1} = x_{n+1} \mid X_n = x_n). \quad (1)$$

We assume further that $\{X_n\}_{n\geq 0}$ is a time-homogeneous DTMC, meaning that for each integer $n \geq 0$, and each $x, y \in S$ (where possibly $x = y$),

$$\mathbb{P}(X_{n+1} = y \mid X_n = x) = \mathbb{P}(X_1 = y \mid X_0 = x) =: p(x, y). \quad (2)$$

In light of these assumptions, we can associate with $\{X_n\}_{n\geq 0}$ a single one-step transition probability matrix $\mathbf{P} := [p(x, y)]_{x,y\in S}$ which completely describes how the chain moves from one time step to the next (assuming we interpret the elements of the index set of $\{X_n\}_{n\geq 0}$ as time instants). In order to properly initialize the chain, we introduce the initial distribution $\boldsymbol{\alpha} := [\alpha(x)]_{x\in S}$, which we interpret as a row vector indexed by S whose elements are defined as

$$\alpha(x) := \mathbb{P}(X_0 = x), \qquad x \in S. \quad (3)$$

It is well-known that through usage of (1), (2), and (3), $\boldsymbol{\alpha}$ and $\mathbf{P}$ together fully determine the finite-dimensional distributions of $\{X_n\}_{n\geq 0}$. Indeed, for any integer $n \geq 0$, and any collection of states $x_0, x_1, \ldots, x_n \in S$,

$$\mathbb{P}(X_0 = x_0, X_1 = x_1, \ldots, X_n = x_n) = \alpha(x_0) \prod_{\ell=1}^{n} p(x_{\ell-1}, x_\ell). \tag{4}$$

We close this introduction with a remark on notation. For each state $x \in S$ we will let $\mathbb{P}_x(\cdot) := \mathbb{P}(\cdot \mid X_0 = x)$, meaning $\mathbb{P}_x(\cdot)$ denotes the conditional probability of $(\cdot)$, given $X_0 = x$. We will often slightly abuse notation by letting $\mathbb{P}_x$ denote a conditional probability, given some other DTMC is in state x at time zero, but it will always be clear from the context which chain is being used whenever we write $\mathbb{P}_x$. Finally, $\mathbb{E}_x$ will denote the conditional expectation associated with $\mathbb{P}_x$, as is often done in the literature.

2.1 The Markov and Strong Markov Property

Most applications of the random-product technique up to now have made extensive use of the Strong Markov property, which is a property shared by all time-homogeneous discrete-time Markov chains. We will not use the random-product technique itself at any point throughout our brief discussion of the Strong Markov property, but given the prevalence of this property in all previous applications of the technique, we are of the opinion that it is useful to provide a quick overview of the basics. A more in-depth discussion of these topics can be found in e.g. Chapter 1 of Asmussen [1].

For each integer $n \geq 0$, define the sub-σ-field $\mathcal{F}_n$ as

$$\mathcal{F}_n := \sigma(X_0, X_1, \ldots, X_n) \tag{5}$$

which represents the smallest σ-field containing events of the form $\{X_0 = x_0, X_1 = x_1, \ldots, X_n = x_n\}$, for any collection of states $x_0, x_1, \ldots, x_n \in S$, which follows from S being a countable set.

The collection of sub-σ-fields $\{\mathcal{F}_n\}_{n\geq 0}$ forms a *filtration*, namely the filtration induced by $\{X_n\}_{n\geq 0}$. It is also helpful to further associate with this filtration the sub-σ-field $\mathcal{F}_\infty$, where

$$\mathcal{F}_\infty := \sigma\left(\bigcup_{n\geq 0} \mathcal{F}_n\right) \tag{6}$$

which is simply the smallest σ-field containing each individual sub-σ-field $\mathcal{F}_n$, $n \geq 0$.

The concept of conditional expectation with respect to a sub-σ-field can be used to state both the Markov property and the Strong Markov property in more sophisticated, yet simpler ways. For instance, it is easy to show that the Markov property given in (1) can be strengthened as follows: for each integer $n \geq 0$, each set $B \in 2^S$, each state $x \in S$, and any collection of sets $A_0, A_1, \ldots, A_{n-1} \in 2^S$, we have

$$\mathbb{P}(X_{n+1} \in B \mid X_n = x, X_{n-1} \in A_{n-1}, \ldots, X_0 \in A_0) = p(x, B) := \sum_{y \in B} p(x, y). \tag{7}$$

This observation shows that the Markov property can alternatively be stated in terms of $\{\mathcal{F}_n\}_{n\geq 0}$ as

$$\mathbb{P}(X_{n+1} \in B \mid \mathcal{F}_n) = p(X_n, B) \tag{8}$$

$\mathbb{P}$-almost surely, for each $B \in 2^S$. More generally, it is also possible to show that for any nonnegative function $h : S^\infty \to [0, \infty)$, and each integer $n \geq 0$,

$$\mathbb{E}_x[h(X_n, X_{n+1}, X_{n+2}, \ldots) \mid \mathcal{F}_n] = \mathbb{E}_{X_n}[h(X_0, X_1, X_2, \ldots)]$$

$\mathbb{P}$-almost surely. Readers should note that $\mathbb{E}_{X_n}[\cdot]$ is a random variable, in particular a measurable function of X_n, and when $X_n = x$, $\mathbb{E}_{X_n}[\cdot] = \mathbb{E}_x[\cdot]$.

We are almost ready to state the strong Markov property: the only remaining task is to define the concept of a stopping time.

Definition 2.1 *We say that a nonnegative extended integer-valued random variable* $\eta : \Omega \to \{0, 1, 2, \ldots\} \cup \{\infty\}$ *is a* stopping time *with respect to the filtration* $\{\mathcal{F}_n\}_{n\geq 0}$ *if, for each integer* $n \geq 0$, $\{\eta \leq n\} \in \mathcal{F}_n$.

One can easily see that η is a stopping time with respect to $\{\mathcal{F}_n\}_{n\geq 0}$ if and only if $\{\eta = n\} \in \mathcal{F}_n$ for each $n \geq 0$: we omit the proof.

Definition 2.2 *Given a stopping time* η *with respect to* $\{\mathcal{F}_n\}_{n\geq 0}$*, we define the collection of sets* $\mathcal{F}_\eta$ *as follows:*

$$\mathcal{F}_\eta := \{A \in \mathcal{F}_\infty : A \cap \{\eta = n\} \in \mathcal{F}_n, \forall\, n \geq 0\}.$$

It is possible to prove that $\mathcal{F}_\eta$ is a sub-σ-field, and that $\mathcal{F}_\eta$ can intuitively be thought of as complete information about the path of $\{X_n\}_{n\geq 0}$ up to the stopping time η.

Theorem 2.1 (*Strong Markov Property*) *For each nonnegative function* $h : S^\infty \to [0, \infty)$ *and each stopping time* η *with respect to* $\{\mathcal{F}_n\}_{n\geq 0}$*, we have that on the set* $\{\eta < \infty\}$

$$\mathbb{E}[h(X_\eta, X_{\eta+1}, X_{\eta+2}, \ldots) \mid \mathcal{F}_\eta] = \mathbb{E}_{X_\eta}[h(X_0, X_1, X_2, \ldots)].$$

$\mathbb{P}$*-almost surely.*

Proof See e.g. Chapter 1 of [1]. $\diamondsuit$

2.2 The Distribution of X_n

Our next task is to analyze all of the marginal distributions of $\{X_n\}_{n\geq 0}$. For each integer $n \geq 1$, we define the n-step transition probability matrix $\mathbf{P}^{(n)} := [p^{(n)}(x, y)]_{x,y\in S}$ as follows:

$$p^{(n)}(x, y) := \mathbb{P}_x(X_n = y), \qquad x, y \in S.$$

It is both well-known, and easy to show that $\mathbf{P}^{(n)} = \mathbf{P}^n$ for each integer $n \geq 1$, which means that the probability mass function of X_n satisfies

$$\mathbb{P}(X_n = y) = \sum_{x\in S} \alpha(x) p^{(n)}(x, y) \tag{9}$$

or, in matrix terms, the row vector that represents the probability mass function of X_n is simply $\boldsymbol{\alpha}\mathbf{P}^n$. Many textbooks usually stop here, but we will find it useful to study such distributions via a generating function approach, particularly because doing so provides us with a smooth way of introducing the random-product technique.

For each $(x, y) \in S \times S$, we define, on the open unit disk $\mathbb{D}[0, 1) := \{z \in \mathbb{C} : |z| < 1\}$ centered at the origin on the complex plane $\mathbb{C}$, the generating function $\varphi_{x,y} : \mathbb{D}[0, 1) \to \mathbb{C}$ as follows:

$$\varphi_{x,y}(z) := \sum_{n=0}^{\infty} z^n p^{(n)}(x, y), \qquad z \in \mathbb{D}[0, 1). \tag{10}$$

Expression (10) is of limited use to us if we already know how to find each $p^{(n)}(x, y)$ term analytically, but often this will not be the case, so our next goal is to give an alternative expression for $\varphi_{x,y}$ that is more amenable to both analytic and numerical studies. In order to state and derive this alternative expression, we need to introduce both the return times $\{\eta_x\}_{x\in S}$ and the hitting times $\{\sigma_x\}_{x\in S}$ associated with $\{X_n\}_{n\geq 0}$, where for each $x \in S$,

$$\eta_x := \inf\{n \geq 1 : X_n = x\}, \qquad \sigma_x := \inf\{n \geq 0 : X_n = x\}.$$

Lemma 2.1 will play an important role in the development of the random-product technique. It should be noted that Equation (11) is a special case of Theorem 1.14 on page 201 of Çinlar [6].

Lemma 2.1 *For each $(x, y) \in S \times S$, we find that for each $z \in \mathbb{D}[0, 1)$,*

$$\varphi_{x,y}(z) = \frac{\mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1} z^n \mathbf{1}(X_n = y)\right]}{1 - \mathbb{E}_x\left[z^{\eta_x}\right]}. \tag{11}$$

Furthermore,

$$\varphi_{x,y}(z) = \varphi_{x,x}(z)\omega_{x,y}(z) \tag{12}$$

where $\omega_{x,y} : \mathbb{D}[0, 1) \cup \{1\} \to \mathbb{C}$ is defined as

$$\omega_{x,y}(z) = \mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1} z^n \mathbf{1}(X_n = y)\right], \quad z \in \mathbb{D}[0, 1). \tag{13}$$

Proof Clearly (12) follows from setting $y = x$ in (11), so it remains to establish (11). Observe that for each $n \geq 1$,

$$\begin{aligned} \mathbb{P}_x(X_n = y) &= \mathbb{P}_x(X_n = y, \eta_x > n) + \sum_{m=1}^{n} \mathbb{P}_x(X_n = y, \eta_x = m) \\ &= \mathbb{P}_x(X_n = y, \eta_x > n) + \sum_{m=1}^{n} \mathbb{P}_x(X_n = y \mid \eta_x = m)\mathbb{P}_x(\eta_x = m) \end{aligned} \tag{14}$$

and since $\mathbb{P}_x(X_n = y \mid \eta_x = m) = p^{(n-m)}(x, y)$ for each $m \in \{1, 2, \ldots, n\}$, we find that

$$p^{(n)}(x, y) = \mathbb{P}_x(X_n = y, \eta_x > n) + \sum_{m=1}^{n} p^{(n-m)}(x, y)\mathbb{P}_x(\eta_x = m). \tag{15}$$

Multiplying both sides of (15) by z^n for $x \in \mathbb{D}[0, 1)$ and summing appropriately further yields

$$\begin{aligned} \varphi_{x,y}(z) &= \sum_{n=0}^{\infty} z^n \mathbb{P}_x(X_n = y, \eta_x > n) + \sum_{n=1}^{\infty}\sum_{m=1}^{n} z^m z^{(n-m)} p^{(n-m)}(x, y)\mathbb{P}_x(\eta_x = m) \\ &= \mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1} z^n \mathbf{1}(X_n = y)\right] + \sum_{m=0}^{\infty} z^m \mathbb{P}_x(\eta_x = m) \sum_{n=m}^{\infty} z^{(n-m)} p^{(n-m)}(x, y) \\ &= \mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1} z^n \mathbf{1}(X_n = y)\right] + \varphi_{x,y}(z)\mathbb{E}_x\left[z^{\eta_x}\right]. \end{aligned} \tag{16}$$

Finally, solving for $\varphi_{x,y}(z)$ in (16) yields (12), which completes the proof of Lemma 2.1. ♢

Readers should glean from Lemma 2.1 that if $\omega_{x,y}(z)$ can be expressed in closed-form for each $y \in S$ (note that $\omega_{x,x}(z) = 1$ for each $z \in D[0, 1)$) and if $\sum_{y \in S} \omega_{x,y}(z)$ is tractable, then for each $y \in S$

$$\varphi_{x,y}(z) = \frac{\omega_{x,y}(z)}{(1 - z)\sum_{u \in S} \omega_{x,u}(z)} \tag{17}$$

meaning that these generating functions are known once each of the $\omega_{x,y}$ functions are known. The $\omega_{x,y}$ functions will also play an important role towards constructing invariant measures for $\{X_n\}_{n \geq 0}$, as well as sufficient conditions for when $\{X_n\}_{n \geq 0}$ is positive recurrent: readers not familiar with these ideas will find them defined shortly.

2.3 The Random-Product Technique in Discrete Time

The random-product technique can be used to provide an alternative expression for each $\omega_{x,y}$ function, and on occasion these new expressions can aid in calculating these functions analytically. Before we introduce the random-product technique in discrete-time, we will first need to come up with a classification system for the states of our time-homogeneous Markov chain $\{X_n\}_{n\geq 0}$ in order to make all random-products well-defined.

Definition 2.3 *We say that a state $x \in S$ is a Type-1 state if $\sum_{y \in S} p(y,x) > 0$. Any state that is not a Type-1 state is called a Type-2 state.*

For each $i \in \{1,2\}$, let S_i denote the set of all Type-i states. Readers should observe that a state x is a Type-2 state if and only if the only way $\{X_n\}_{n\geq 0}$ can ever visit state x is if $X_0 = x$, and even if this is the case, we must also have $X_n \in S_1$ for each $n \geq 1$.

Associated with $\{X_n\}_{n\geq 0}$ is an entire collection of DTMCs having state space S_1, where a DTMC $\{\tilde{X}_n\}_{n\geq 0}$ is in this collection if and only if its transition probability matrix $\tilde{\mathbf{P}}$ satisfies the following property: given any two states $x, y \in S_1$, where possibly $x = y$, $\tilde{p}(x,y) > 0$ if and only if $p(y,x) > 0$. The reason why we set the state space of $\{\tilde{X}_n\}_{n\geq 0}$ to be S_1 is because we want to avoid situations where, under our rule for constructing $\{\tilde{X}_n\}_{n\geq 0}$, we end up with a state x that must satisfy

$$\sum_{y \in S} \tilde{p}(x,y) = 0$$

and restricting the state space to S_1 avoids this dilemma. We further associate with $\{\tilde{X}_n\}_{n\geq 0}$ its return times $\{\tilde{\eta}_x\}_{x\in S}$ and its hitting times $\{\tilde{\sigma}_x\}_{x\in S}$, where

$$\tilde{\eta}_x := \inf\{n \geq 1 : \tilde{X}_n = x\}, \qquad \tilde{\sigma}_x := \inf\{n \geq 0 : \tilde{X}_n = x\}.$$

Our next result provides a way of expressing each $\omega_{x,y}(z)$ term in terms of $\{\tilde{X}_n\}_{n\geq 0}$, whenever $(x,y) \in S_1 \times S_1$. Observe too that $\mathbb{P}_x$ and/or $\mathbb{E}_x$ refers to conditioning on either $\{X_0 = x\}$ or $\{\tilde{X}_0 = x\}$: it will be clear from the context which event is being conditioned on in each conditional probability/expectation.

Theorem 2.2 *For each $(x,y) \in S_1 \times S_1$, we find that for each $z \in \mathbb{D}[0,1) \cup \{1\}$,*

$$\omega_{x,y}(z) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty) z^{\tilde{\sigma}_x} \prod_{\ell=1}^{\tilde{\sigma}_x} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right]. \tag{18}$$

Proof Equality (18) can be established by making use of a sum-over-paths approach, while simultaneously making use of the fact that for any collection of states $x_0, x_1, \ldots, x_n \in S_1$, $x_0, x_1, \ldots, x_n$ is a feasible path for $\{\tilde{X}_n\}_{n\geq 0}$ if and only if $x_n, \ldots, x_1, x_0$ is a feasible path under $\{X_n\}_{n\geq 0}$.

Fix a state $x \in S$, and observe first that (18) clearly holds when $y = x$, since $w_{x,x}(z) = 1$. Furthermore, for each $y \in S_1$ satisfying $y \neq x$,

$$\begin{aligned}
\mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty) z^{\tilde{\sigma}_x} \prod_{\ell=1}^{\tilde{\sigma}_x} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] &= \sum_{n=1}^{\infty} \sum_{x_0,x_1,\ldots,x_n : x_0 = y, x_1 \neq x, \ldots, x_{n-1} \neq x} z^n \left[\prod_{\ell=1}^{n} \frac{p(x_\ell, x_{\ell-1})}{\tilde{p}(x_{\ell-1}, x_\ell)}\right] \prod_{\ell=1}^{n} \tilde{p}(x_{\ell-1}, x_\ell) \\
&= \sum_{n=1}^{\infty} \sum_{x_0,x_1,\ldots,x_n : x_0 = y, x_1 \neq x, \ldots, x_{n-1} \neq x} z^n \prod_{\ell=1}^{n} p(x_\ell, x_{\ell-1}) \\
&= \sum_{n=1}^{\infty} z^n \mathbb{P}_x(\eta_x > n, X_n = y) \\
&= \mathbb{E}_x\left[\sum_{n=1}^{\eta_x - 1} z^n \mathbf{1}(X_n = y)\right] \\
&= \omega_{x,y}(z).
\end{aligned}$$

This establishes (18), and completes the proof of Theorem 2.2. $\diamondsuit$

Example Suppose $\{X_n\}_{n\geq 0}$ is such that $S = \{0, 1, 2, 3, \ldots\}$, where $\mathbf{P}$ satisfies

$$p(x, x+1) = p, \qquad p(x, x-1) = q := 1 - p, \qquad x \geq 1$$

and $p(0, 0) = 1$, with all other elements of $\mathbf{P}$ set equal to zero. This DTMC is often referred to as a simple random walk with an absorbing boundary at state zero.

We can use Theorem 2.2 to calculate $\varphi_{1,y}(z)$ for each $z \in \mathbb{D}[0, 1)$, and each state $y \in S$. Recall from (12) that for each state $y \in S$,

$$\varphi_{1,y}(z) = \varphi_{1,1}(z)\omega_{1,y}(z)$$

which means our main task reduces to the problem of first calculating $\omega_{1,y}(z)$ for each state $y \neq 1$, then normalizing to get both $\varphi_{1,1}(z)$, and $\varphi_{1,y}(z)$ for each $y \neq 1$. Observe first that, regardless of how we choose $\{\tilde{X}_n\}_{n\geq 0}$, summing over all finite-length feasible paths from state 0 to state 1 under $\{\tilde{X}_n\}_{n\geq 0}$ yields

$$\omega_{1,0}(z) = \mathbb{E}_0\left[\mathbf{1}(\tilde{\sigma}_1 < \infty) z^{\tilde{\sigma}_1} \prod_{\ell=1}^{\tilde{\sigma}_1} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] = q\sum_{n=1}^{\infty} z^n = \frac{qz}{1-z}.$$

Next, suppose $y \geq 2$, and choose the transition matrix $\tilde{\mathbf{P}}$ so that for each state $y \geq 2$,

$$\tilde{p}(y, y-1) = p(y, y-1), \qquad \tilde{p}(y, y+1) = p(y, y+1).$$

Under this particular choice for $\tilde{\mathbf{P}}$, we observe after a bit of thought that under the measure $\mathbb{P}_y$,

$$\mathbf{1}(\tilde{\sigma}_1 < \infty) \prod_{\ell=1}^{\tilde{\sigma}_1} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} = \mathbf{1}(\tilde{\sigma}_1 < \infty) r^{y-1}$$

where $r := p/q$. Then by the Strong Markov property,

$$\omega_{1,y}(z) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_1 < \infty) z^{\tilde{\sigma}_1} \prod_{\ell=1}^{\tilde{\sigma}_1} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] = r^{y-1}\mathbb{E}_1[\mathbf{1}(\sigma_0 < \infty) z^{\sigma_0}]^{y-1} = (r\mathbb{E}_1[z^{\sigma_0}])^{y-1}$$

where the last equality follows from the fact that $z \in D[0, 1)$. Finally, since

$$\sum_{y\in S} \varphi_{1,y}(z) = \frac{1}{1-z}$$

for each $z \in \mathbb{D}[0, 1)$, we conclude that

$$\frac{1}{1-z} = \varphi_{1,1}(z)\left[\frac{qz}{1-z} + \sum_{y=1}^{\infty} (r\mathbb{E}_1[z^{\sigma_0}])^{y-1}\right] = \varphi_{1,1}(z)\left[\frac{qz}{1-z} + \frac{1}{1 - r\mathbb{E}_1[z^{\sigma_0}]}\right] = \varphi_{1,1}(z)\left[\frac{qz - pz\mathbb{E}_1[z^{\sigma_0}] + 1 - z}{(1-z)(1 - r\mathbb{E}_1[z^{\sigma_0}])}\right]$$

which means

$$\varphi_{1,1}(z) = \frac{1 - r\mathbb{E}_1[z^{\sigma_0}]}{1 - pz - pz\mathbb{E}_1[z^{\sigma_0}]}.$$

Obviously the other generating functions can be calculated. For instance,

$$\varphi_{1,0}(z) = \frac{qz - pz\mathbb{E}_1[z^{\sigma_0}]}{(1-z)(1 - pz - pz\mathbb{E}_1[z^{\sigma_0}])}$$

and for each state $y \geq 2$,

$$\varphi_{1,y}(z) = \frac{(1 - r\mathbb{E}_1[z^{\sigma_0}])(r\mathbb{E}_1[z^{\sigma_0}])^{y-1}}{1 - pz - pz\mathbb{E}_1[z^{\sigma_0}]}.$$

We claim that these expressions for $\phi_{1,y}$ are essentially closed-form expressions. Using first-step analysis (i.e. conditioning on X_1), coupled with the Strong Markov property, yields

$$\mathbb{E}_1[z^{\sigma_0}] = qz + pz\mathbb{E}_2[z^{\sigma_0}] = qz + pzE_1[z^{\sigma_0}]^2$$

which shows $\mathbb{E}_1[z^{\sigma_0}]$ is a root of a quadratic polynomial, in particular,

$$\mathbb{E}_1[z^{\sigma_0}] = \frac{1 - \sqrt{1 - 4pqz^2}}{2pz}.$$

We end this example by quickly noting that similar expressions can be derived for $\varphi_{x,y}(z)$ when $x \neq 1$, although the expressions are slightly more involved when x is allowed to be an arbitrary state.

Example Our next example illustrates how it is possible to study the time-dependent behavior of a random walk on a graph. Suppose $G(V, E)$ denotes an undirected, connected graph, where V denotes its set of vertices and E its set of edges. Due to $G(V, E)$ being an undirected graph, we have that for each $x, y \in V$, $(x, y) \in E$ if and only if $(y, x) \in E$. We will also assume, without loss of generality, that the graph has no self-loops, meaning $(x, x) \notin E$ for each $x \in V$.

We further associate with this graph the function $d : V \times V \to (0, \infty)$, which satisfies the following assumptions:

(a) $d((x, y)) > 0$ for each edge $(x, y) \in E$;

(b) $d((x, y)) = d((y, x))$ for each edge $(x, y) \in E$;

(c) $d((x, y)) = 0$ for each $(x, y) \notin E$.

Throughout the rest of this example, we will express $d((x, y))$ as $d(x, y)$, in the interest of making our formulas more aesthetically pleasing. Letting, for each $x \in V$,

$$d(x) := \sum_{y \in V} d(x, y)$$

one can easily show that due to both our assumption that the graph is connected, as well as our assumptions on the function d, we must have $d(x) > 0$.

We can use this framework to construct a DTMC $\{X_n\}_{n \geq 0}$, whose transition matrix $\mathbf{P} := [p(x, y)]_{x,y \in S}$ satisfies

$$p(x, y) := \frac{d(x, y)}{d(x)}, \quad x, y \in S.$$

Our objective is to use the random-product technique to derive an expression for $\varphi_{x,y}(z)$, for any two states $x, y \in S$. Given that $d(x) > 0$ for each state $x \in S$, and $d(x, y) > 0$ if and only if $d(y, x) > 0$ for any two states $x, y \in S$, it quickly follows that $\mathbf{P}$ satisfies the property that $p(y, x) > 0$ if and only if $p(x, y) > 0$. Selecting, then, $\tilde{\mathbf{P}} = \mathbf{P}$, it follows that for each $x, y \in S$ where $x \neq y$,

$$\begin{aligned}
\omega_{x,y}(z) &= \mathbb{E}_y\left[\mathbf{1}(\sigma_x < \infty)z^{\sigma_x} \prod_{\ell=1}^{\sigma_x} \frac{p(X_\ell, X_{\ell-1})}{p(X_{\ell-1}, X_\ell)}\right] \\
&= \mathbb{E}_y\left[\mathbf{1}(\sigma_x < \infty)z^{\sigma_x} \prod_{\ell=1}^{\sigma_x} \frac{d(X_{\ell-1})d(X_\ell, X_{\ell-1})}{d(X_\ell)d(X_{\ell-1}, X_\ell)}\right] \\
&= \frac{d(y)}{d(x)}\mathbb{E}_y[\mathbf{1}(\sigma_x < \infty)z^{\sigma_x}] \\
&= \frac{d(y)}{d(x)}\mathbb{E}_y[z^{\sigma_x}]
\end{aligned}$$

for each $z \in D[0,1)$. After a small amount of algebra, we conclude that for each $x, y \in S$,

$$\varphi_{x,y}(z) = \frac{d(y)\mathbb{E}_y[z^{\sigma_x}]}{1 - \mathbb{E}_x[z^{\sigma_x}]} = \frac{d(y)\mathbb{E}_y[z^{\sigma_x}]}{(1-z)\sum_{u \in S} d(u)\mathbb{E}_u[z^{\sigma_x}]}.$$

In conclusion, for a random walk on an undirected graph, calculating the $\varphi_{x,y}$ generating functions of the random walk simplifies to the problem of calculating the generating functions of hitting times and return times of the random walk.

2.4 Stationary Distributions

The random-product technique is an especially useful tool to apply towards the construction of invariant measures and stationary distributions.

Definition 2.4 *A nonnegative row vector (i.e. a nonnegative measure)* $\boldsymbol{\mu} := [\mu(y)]_{y \in S}$ *is said to be an* invariant measure *of* $\{X_n\}_{n \geq 0}$ *if it satisfies* $\boldsymbol{\mu} = \boldsymbol{\mu}\mathbf{P}$*. Saying* $\boldsymbol{\mu}$ *is nonnegative simply means* $\mu(y) \geq 0$ *for each* $y \in S$*: likewise, saying* $\boldsymbol{\mu}$ *is positive means* $\mu(y) > 0$ *for each* $y \in S$*.*

Recurrent states can be used to construct invariant measures for a DTMC, so before we actually construct invariant measures we first need to recall how both recurrent and transient states are defined.

Definition 2.5 *We say that state* $x \in S$ *is* recurrent *if* $\mathbb{P}_x(\eta_x < \infty) = 1$*. If a state* x *is not recurrent, we say it is* transient*.*

It is also both helpful and necessary to further classify a recurrent state as being either 'positive recurrent' or 'null recurrent'.

Definition 2.6 *A recurrent state* x *is* positive recurrent *if* $\mathbb{E}_x[\eta_x] < \infty$*. If a recurrent state* x *is not positive recurrent, we say it is* null recurrent*.*

Readers should recall that if two distinct states x and y communicate, i.e. if there exists two integers $m \geq 1$, $n \geq 1$ such that $p^{(m)}(x,y) > 0$ and $p^{(n)}(y,x) > 0$, then those two states must both be either positive recurrent, null recurrent, or transient. In other words, if C_x denotes the set consisting of x, as well as all states that communicate with x, then each state in C_x must be either positive recurrent, null recurrent, or transient. Recall also that $\{X_n\}_{n \geq 0}$ is said to be *irreducible* if all states communicate.

Example It is both useful and instructive to illustrate these concepts through usage of the one-dimensional reflected random walk. This is simply a DTMC $\{X_n\}_{n \geq 0}$ with state space $S = \{0, 1, 2, \ldots\}$ and transition probability matrix $\mathbf{P} = [p(x,y)]_{x,y \in S}$ where the transition probabilities are as follows:

$$\begin{aligned} &p(0,0) = q, \\ &p(x, x+1) = p, \quad x = 0, 1, \ldots \\ &p(x, x-1) = q, \quad x = 1, 2, \ldots \end{aligned}$$

and $p(x,y) = 0$ for all other x and y. We assume p and q are real numbers satisfying $0 < p < 1$, and $p + q = 1$. We can establish positive recurrence, null recurrence, and transience of each state of this irreducible DTMC by noting the following consequences of Theorem 2.2:

$$w_{x,y}(1) = \mathbb{E}_y \left[\mathbf{1}(\tilde{\sigma}_x < \infty) \prod_{\ell=1}^{\tilde{\sigma}_x} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right] \tag{19}$$

and

$$\mathbb{E}_x[\eta_x] = \sum_{y \in S} w_{x,y}(1). \tag{20}$$

Our first task is to show that state 0 is positive recurrent when $p < q$. Set $\tilde{\mathbf{P}} = \mathbf{P}$ and choose state 0 as a reference point: then for each state $y \geq 0$,

$$\mathbb{E}_y \left[\mathbf{1}(\tilde{\sigma}_0 < \infty) \prod_{\ell=1}^{\tilde{\sigma}_0} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right] = \left(\frac{p}{q} \right)^y \mathbb{P}_y(\sigma_0 < \infty). \tag{21}$$

Summing (21) over all states $y \geq 0$ further yields

$$\mathbb{E}_0 [\eta_0] = \sum_{y=0}^{\infty} \left(\frac{p}{q} \right)^y \mathbb{P}_y(\sigma_0 < \infty) < \infty \tag{22}$$

since $p < q$, and $0 \leq \mathbb{P}_y(\sigma_0 < \infty) \leq 1$ for each $y \geq 0$. This proves $\mathbb{E}_0 [\eta_0] < \infty$, i.e. state 0 is positive recurrent.

The same choice for $\tilde{\mathbf{P}}$ can be used to show state 0 is null recurrent when $p = q$. Formula (21) further reveals that, in this case, for each state $y \geq 1$,

$$w_{0,y}(1) = \mathbb{P}_y(\sigma_0 < \infty) = \mathbb{P}_1(\sigma_0 < \infty)^y \tag{23}$$

where the last equality follows from the Strong Markov property. Summing (23) over each state $y \geq 0$ (clearly (23) holds also for the case where $y = 0$) we get

$$\mathbb{E}_0 [\eta_0] = \sum_{y=0}^{\infty} \mathbb{P}_1(\sigma_0 < \infty)^y. \tag{24}$$

From Equality (24), it follows that $\mathbb{P}_1(\sigma_0 < \infty) = 1$. If this were not the case, we would have that $\mathbb{P}_0(\eta_0 < \infty) < 1$ and $\mathbb{E}_0[\eta_0] = \infty$, yet if $\mathbb{P}_1(\sigma_0 < \infty) < 1$, formula (24) implies $\mathbb{E}_0 [\eta_0] < \infty$, a contradiction. This proves $\mathbb{P}_1(\sigma_0 < \infty) = 1$, meaning $\mathbb{P}_0(\eta_0 < \infty) = 1$ and by (24), $\mathbb{E}_0 [\eta_0] = \infty$, so state 0 is null recurrent.

It remains to show that each state is transient when $p > q$: this portion of the argument requires analyzing $w_{0,y}(1)$ under two different choices of $\tilde{\mathbf{P}}$. On the one hand, if we choose $\tilde{\mathbf{P}} = \mathbf{P}$, we get

$$\mathbb{E}_0 \left[\sum_{n=0}^{\eta_0 - 1} \mathbf{1}(X_n = y) \right] = \left(\frac{p}{q} \right)^y \mathbb{P}_y(\sigma_0 < \infty). \tag{25}$$

On the other hand, if we choose $\tilde{\mathbf{P}}$ such that for each state $y \geq 1$,

$$\tilde{p}(y, y-1) = p, \quad \tilde{p}(y, y+1) = q \tag{26}$$

we instead get

$$\mathbb{E}_0 \left[\sum_{n=0}^{\eta_0 - 1} \mathbf{1}(X_n = y) \right] = \mathbb{P}_y(\tilde{\sigma}_0 < \infty) = 1, \tag{27}$$

this is immediate by the first case we explored in this example, because $\frac{q}{p} < 1$. Hence,

$$1 = \left(\frac{p}{q} \right)^y \mathbb{P}_y(\sigma_0 < \infty),$$

which implies for each state $y \geq 1$ that

$$\mathbb{P}_y(\sigma_0 < \infty) = \left(\frac{q}{p} \right)^y < 1$$

from which we get $\mathbb{P}_0(\eta_0 < \infty) < 1$, i.e state 0 is transient.

We are almost ready to show how the random-product technique can be used to construct invariant measures with recurrent states. Many of the results in this section will rely on an alternative DTMC $\{\overline{X}_n\}_{n \geq 0}$ constructed from the original DTMC $\{X_n\}_{n \geq 0}$, where this alternative DTMC is constructed from a known positive invariant measure $\boldsymbol{\mu}$ of $\{X_n\}_{n \geq 0}$.

Definition 2.7 *Given a DTMC $\{X_n\}_{n\geq 0}$ with state space S and positive invariant measure $\boldsymbol{\mu}$ satisfying $\mu(y) < \infty$ for each $y \in S$, we define $\{\overline{X}_n\}_{n\geq 0}$ as the $\boldsymbol{\mu}$-dual of $\{X_n\}_{n\geq 0}$, where the state space of $\{\overline{X}_n\}_{n\geq 0}$ is S, and its transition matrix is $\overline{\mathbf{P}}$, whose elements satisfy*

$$\overline{p}(x, y) := \frac{\mu(y)p(y, x)}{\mu(x)}.$$

Readers should notice that $\overline{\mathbf{P}}$ is a transition probability matrix because $\boldsymbol{\mu}$ is a positive invariant measure: indeed, for each $x \in S$,

$$\sum_{y \in S} \overline{p}(x, y) = \frac{1}{\mu(x)} \sum_{y \in S} \mu(y)p(y, x) = \frac{\mu(x)}{\mu(x)} = 1.$$

Ideally it would be best to incorporate the symbol $\boldsymbol{\mu}$ somewhere onto $\{\overline{X}_n\}_{n\geq 0}$, but doing so makes, in our opinion, the notation unwieldy. It will always be made clear which positive invariant measure we are referring to anytime we write $\{\overline{X}_n\}_{n\geq 0}$. Furthermore, we can also associate with $\{\overline{X}_n\}_{n\geq 0}$ the hitting-time and return-time random variables $\{\overline{\sigma}_x\}_{x\in S}$, $\{\overline{\eta}_x\}_{x\in S}$, defined as

$$\overline{\sigma}_x := \inf\{n \geq 0 : \overline{X}_n = x\}, \qquad \overline{\eta}_x := \inf\{n \geq 1 : \overline{X}_n = x\}.$$

Notice too that $\{\overline{X}_n\}_{n\geq 0}$ is a possible choice for $\{\tilde{X}_n\}_{n\geq 0}$ anytime we make use of the random-product technique, and we will often make this choice when we use this technique to prove a result. Readers that are already familiar with the theory of Markov chains may, at first glace, say that $\{\overline{X}_n\}_{n\geq 0}$ is the transition matrix of the *time-reversal* of $\{X_n\}_{n\geq 0}$, but it may not be the case that $\boldsymbol{\mu}$ has finite total mass: in fact, in our proofs we will usually want to refrain from further assuming

$$\sum_{y \in S} \mu(y) < \infty.$$

Proposition 2.1 *Suppose $\{X_n\}_{n\geq 0}$ is irreducible and has a positive invariant measure $\boldsymbol{\mu}$. Then its $\boldsymbol{\mu}$-dual $\{\overline{X}_n\}_{n\geq 0}$ is also irreducible, and its return-time random variables $\{\overline{\eta}_x\}_{x\in S}$ satisfy*

$$\mathbb{P}(\overline{\eta}_x = n \mid \overline{X}_0 = x) = \mathbb{P}(\eta_x = n \mid X_0 = x)$$

for each integer $n \geq 1$.

Proof Clearly $\{\overline{X}_n\}_{n\geq 0}$ is irreducible: this is true because if $x = x_0, x_1, \ldots, x_n = y$ is a feasible path from state x to state y under $\mathbf{P}$, meaning

$$\prod_{\ell=1}^{n} p(x_{\ell-1}, x_\ell) > 0$$

then the reverse path $y = x_n, x_{n-1}, \ldots, x_1, x_0 = x$ is a feasible path from state y to state x under $\overline{\mathbf{P}}$, since

$$\prod_{\ell=1}^{n} \overline{p}(x_\ell, x_{\ell-1}) = \prod_{\ell=1}^{n} \frac{\mu(x_{\ell-1})p(x_{\ell-1}, x_\ell)}{\mu(x_\ell)} = \frac{\mu(x)}{\mu(y)} \prod_{\ell=1}^{n} p(x_{\ell-1}, x_\ell) > 0.$$

Next, observe that for each integer $n \geq 1$,

$$\begin{aligned}
\mathbb{P}_x(\overline{\eta}_x = n) &= \sum_{x_0, x_1, \ldots, x_n : x_0 = x, x_n = x, x_1, \ldots, x_{n-1} \neq x} \prod_{\ell=1}^{n} \overline{p}(x_{\ell-1}, x_\ell) \\
&= \sum_{x_0, x_1, \ldots, x_n : x_0 = x, x_n = x, x_1, \ldots, x_{n-1} \neq x} \prod_{\ell=1}^{n} p(x_\ell, x_{\ell-1}) = \mathbb{P}_x(\eta_x = n)
\end{aligned}$$

which proves the claim. $\diamondsuit$

The next result is an immediate consequence of Proposition 2.1.

Corollary 2.1 *Let $\boldsymbol{\mu}$ be a positive invariant measure of $\{X_n\}_{n\geq 0}$. Then for each state $x \in S$, x is recurrent with respect to $\{X_n\}_{n\geq 0}$ if and only if it is recurrent with respect to the μ-dual $\{\overline{X}_n\}_{n\geq 0}$. Furthermore, x is positive recurrent with respect to $\{X_n\}_{n\geq 0}$ if and only if it is positive recurrent with respect to $\{\overline{X}_n\}_{n\geq 0}$, and x is null recurrent with respect to $\{X_n\}_{n\geq 0}$ if and only if it is null recurrent with respect to $\{\overline{X}_n\}_{n\geq 0}$.*

We are now ready to show how a recurrent state can be used to construct an invariant measure. Given a recurrent state $x \in S$, define the measure $\boldsymbol{\mu}^{(x)} := [\mu^{(x)}(y)]_{y\in S}$ as

$$\mu^{(x)}(y) := \omega_{x,y}(1) = \mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1} \mathbf{1}(X_n = y)\right]. \tag{28}$$

It is worth pointing out that the total mass of $\mu^{(x)}$ over S is the expected amount of time it takes $\{X_n\}_{n\geq 0}$ to return to state x, given $X_0 = x$, i.e.

$$\sum_{y\in S} \mu^{(x)}(y) = \mathbb{E}_x[\eta_x]. \tag{29}$$

It is well-known that $\boldsymbol{\mu}^{(x)}$ is an invariant measure of $\{X_n\}_{n\geq 0}$: this is typically proven by making use of the *cycle-trick*, see e.g. the texts of Durrett [7], [8]. Our usage of the random-product technique will allow us to replace the cycle-trick with a first-step analysis argument, once we are able to re-express each $w_{x,y}(1)$ in terms of the return-time random variables $\tilde{\eta}_x$ instead of the hitting-time random variables $\tilde{\sigma}_x$ (this is the objective of the next proposition). Once we have established that $\boldsymbol{\mu}^{(x)}$ is an invariant measure, we will show that each positive, finite invariant measure $\boldsymbol{\mu}$ must be a positive scalar multiple of $\boldsymbol{\mu}^{(x)}$. This will be done by again using the random-product technique, combined with properties of the $\boldsymbol{\mu}$-dual $\{\overline{X}_n\}_{n\geq 0}$.

Proposition 2.2 *Suppose x is a recurrent state. Then for each state $y \in S$,*

$$\omega_{x,y}(1) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right]. \tag{30}$$

Proof The statement is trivial for the cases where $y \neq x$, since under $\mathbb{P}_y$, $\tilde{\eta}_x = \tilde{\sigma}_x$. Next,

$$\begin{aligned}
\mathbb{E}_x\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x} \frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] &= \sum_{n=1}^{\infty} \sum_{x_0,\ldots,x_n : x_0 = x, x_n = x, x_1,\ldots,x_{n-1}\neq x} \left[\prod_{\ell=1}^{n} \frac{p(x_\ell, x_{\ell-1})}{\tilde{p}(x_{\ell-1}, x_\ell)}\right] \prod_{\ell=1}^{n} p(x_\ell, x_{\ell-1}) \\
&= \sum_{n=1}^{\infty} \sum_{x_0,\ldots,x_n : x_0 = x, x_n = x, x_1,\ldots,x_{n-1}\neq x} \prod_{\ell=1}^{n} p(x_\ell, x_{\ell-1}) \\
&= \sum_{n=1}^{\infty} \mathbb{P}_x(\eta_x = n) = \mathbb{P}_x(\eta_x < \infty) = 1 = \omega_{x,x}(1)
\end{aligned}$$

which proves the claim. $\diamondsuit$

Theorem 2.3 *Fix a recurrent state $x \in S$. Then, the row vector $\boldsymbol{\mu}^{(x)}$ satisfies $\boldsymbol{\mu}^{(x)}\mathbf{P} = \boldsymbol{\mu}^{(x)}$. Furthermore, if $\{X_n\}_{n\geq 0}$ is both irreducible and positive recurrent, then all positive invariant measures $\boldsymbol{\mu}$ of $\{X_n\}_{n\geq 0}$ are scalar multiples of each other.*

Proof Given x is recurrent and $X_0 = x$, we may assume without loss of generality that $\{X_n\}_{n\geq 0}$ is both irreducible and recurrent, because otherwise, under the assumption that $X_0 = x$, we can think of the state space of the chain as simply the communicating class C_x containing x.

Our goal is to use Proposition 2.2 to show that $\boldsymbol{\mu}^{(x)}$ is an invariant measure for $\{X_n\}_{n\geq 0}$: doing so requires us to show that for every $y \in S$,

$$\mu^{(x)}(y) = \sum_{z\in S} \mu^{(x)}(z) p(z, y). \tag{31}$$

Indeed, from Proposition 2.2, we can say that for each state $y \in S$,

$$\mu^{(x)}(y) = \omega_{x,y}(1) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x}\frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right]. \tag{32}$$

Conditioning on $\tilde{X}_1$ within the right-hand-side of (32) yields

$$\begin{aligned}
\mu^{(x)}(y) &= \sum_{z\in S}\mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x}\frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \mid \tilde{X}_1 = z\right]\tilde{p}(y,z) \\
&= \sum_{z\in S}\mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=2}^{\tilde{\eta}_x}\frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \mid \tilde{X}_1 = z\right]\frac{p(z,y)}{\tilde{p}(y,z)}\tilde{p}(y,z) \\
&= \sum_{z\in S}\mu^{(x)}(z)p(z,y),
\end{aligned}$$

which establishes (31), thus proving $\boldsymbol{\mu}^{(x)}$ is an invariant measure. It is also easy to use the balance equations to show that $\mu^{(x)}(y)$ is both positive and finite for each state $y \in C_x$.

It remains to show that when $\{X_n\}_{n\geq 0}$ is both irreducible and recurrent, all positive invariant measures of $\{X_n\}_{n\geq 0}$ are scalar multiples of each other. In order to prove this claim, we fix an arbitrarily chosen state $x \in S$, and show that each positive invariant measure must be a positive scalar multiple of $\boldsymbol{\mu}^{(x)}$. Let $\boldsymbol{\mu}$ be another positive invariant measure of $\{X_n\}_{n\geq 0}$, and let $\{\overline{X}_n\}_{n\geq 0}$ be its $\boldsymbol{\mu}$-dual. Then for each state $y \neq x$,

$$\mu^{(x)}(y) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{\mu(\tilde{X}_{\ell-1})}{\mu(\tilde{X}_\ell)}\right] = \frac{\mu(y)}{\mu(x)}\mathbb{P}_y(\tilde{\sigma}_x < \infty) = \frac{\mu(y)}{\mu(x)}$$

or, equivalently

$$\mu(y) = \mu(x)\mu^{(x)}(y) \tag{33}$$

which proves $\boldsymbol{\mu}$ is a positive scalar multiple of $\boldsymbol{\mu}^{(x)}$, as (33) is trivially true when $y = x$. $\diamondsuit$

The next theorem is also a well-known result that provides sufficient conditions for $\{X_n\}_{n\geq 0}$ to be positive recurrent.

Theorem 2.4 *If $\{X_n\}_{n\geq 0}$ is irreducible, and has a positive stationary distribution $\boldsymbol{\pi} := (\pi(y))_{y\in S}$, then it is positive recurrent.*

Proof Fix a state $x \in S$, and recall that the total mass of the measure $\boldsymbol{\mu}^{(x)}$ satisfies

$$\sum_{y\in S}\mu^{(x)}(y) = \sum_{y\in S}\mathbb{E}_x\left[\sum_{n=0}^{\eta_x - 1}\mathbf{1}(X_n = y)\right] = \mathbb{E}_x[\eta_x]. \tag{34}$$

Letting now $\{\overline{X}_n\}_{n\geq 0}$ be the $\boldsymbol{\pi}$-dual of $\{X_n\}_{n\geq 0}$, choose $\{\tilde{X}_n\}_{n\geq 0}$ to be $\{\overline{X}_n\}_{n\geq 0}$: for each state $y \neq x$

$$\mu^{(x)}(y) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{p(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{p}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] = \mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{\pi(\tilde{X}_{\ell-1})}{\pi(\tilde{X}_\ell)}\right] = \frac{\pi(y)}{\pi(x)}\mathbb{P}_y(\tilde{\sigma}_x < \infty). \tag{35}$$

Combining observations (34) and (35) gives

$$\mathbb{E}_x[\eta_x] = \mu^{(x)}(x) + \sum_{y\neq x}\mu^{(x)}(y) = 1 + \sum_{y\neq x}\frac{\pi(y)}{\pi(x)}\mathbb{P}_y(\tilde{\sigma}_x < \infty) \leq \frac{\pi(x)}{\pi(x)} + \sum_{y\neq x}\frac{\pi(y)}{\pi(x)} = \frac{1}{\pi(x)} < \infty,$$

which proves state x (and all states, since x was chosen arbitrarily) is positive recurrent. $\diamondsuit$

Typically Theorem 2.4 is proven with a 'proof-by-contradiction' argument, by first assuming that all states are transient: under this assumption, we have that for each integer $n \geq 1$, and each state $y \in S$ that

$$\pi(y) = \sum_{z \in S} \pi(z) p^{(n)}(z, y).$$

Letting $n \to \infty$ on the right-hand-side, while at the same time applying the dominated convergence theorem, coupled with the fact that $\lim_{n\to\infty} p^{(n)}(z, y) = 0$ due to each state being transient, yields $\pi(y) = 0$ for each state $y \in S$, which is a contradiction. This proves each state is recurrent, and by Theorem 2.3, we conclude that all states are positive recurrent.

2.5 Time-Reversals

Theorem 2.4 showed that if an irreducible DTMC $\{X_n\}_{n\geq 0}$ has a positive invariant measure $\boldsymbol{\mu}$ with finite total mass, we can normalize $\boldsymbol{\mu}$ to get the unique stationary distribution $\boldsymbol{\pi}$ of $\{X_n\}_{n\geq 0}$, since $\boldsymbol{\pi}$ must be a positive scalar multiple of $\boldsymbol{\mu}$. Not only that, the $\boldsymbol{\mu}$-dual has the same transition probability matrix as the $\boldsymbol{\pi}$-dual.

The $\boldsymbol{\pi}$-dual $\{\overline{X}_n\}_{n\geq 0}$ has already proven itself to be a very useful DTMC, given the role it has played in a few of our proofs, but in many texts the $\boldsymbol{\pi}$-dual is referred to as the *time-reversal* of $\{X_n\}_{n\geq 0}$. One also typically says that an irreducible, positive recurrent DTMC $\{X_n\}_{n\geq 0}$ is *reversible* if $\overline{\mathbf{P}} = \mathbf{P}$, i.e. if the transition probability matrix of the $\boldsymbol{\pi}$-dual of $\{X_n\}_{n\geq 0}$ is equal to the transition probability matrix of $\{X_n\}_{n\geq 0}$.

Theorem 2.5 (*Generalized Kolmogorov's Criterion*) *Suppose $\{X_n\}_{n\geq 0}$ is both irreducible and positive recurrent with stationary distribution $\boldsymbol{\pi}$, and let $\{\tilde{X}_n\}_{n\geq 0}$ be another DTMC whose state space is S, and whose transition matrix is $\tilde{\mathbf{P}}$, where the elements of $\tilde{\mathbf{P}}$ satisfy the following property: for each $x, y \in S$, $\tilde{p}(x, y) > 0$ if and only if $p(y, x) > 0$. Then the following statements are equivalent:*

- (a) *$\{\tilde{X}_n\}_{n\geq 0}$ is the time-reversal of $\{X_n\}_{n\geq 0}$.*
- (b) *The transition probabilities found in $\tilde{\mathbf{P}}$ satisfy the following criterion: for each state $x \in S$, and each feasible path with respect to $\tilde{\mathbf{P}}$ of the form $x = x_0, x_1, \ldots, x_{n-1}, x_n = x$ such that $x_1 \neq x, x_2 \neq x, \ldots, x_{n-1} \neq x$, we have*

$$\prod_{\ell=1}^{n} \tilde{p}(x_{\ell-1}, x_\ell) = \prod_{\ell=1}^{n} p(x_\ell, x_{\ell-1}). \tag{36}$$

- (c) *The transition probabilities found in $\tilde{\mathbf{P}}$ satisfy the following ratio criterion: for any two distinct states $x, y \in S$, we have that*

$$\prod_{\ell=1}^{n} \frac{p(x_\ell, x_{\ell-1})}{\tilde{p}(x_{\ell-1}, x_\ell)} \tag{37}$$

  depends only on y and x, for any feasible path $y = x_0, x_1, \ldots, x_{n-1}, x_n = x$ under $\tilde{\mathbf{P}}$ satisfying the property that $x_1 \neq x, x_2 \neq x, \ldots, x_{n-1} \neq x$.

Property (b) is what is typically referred to as the generalized Kolmogorov's criterion, but in light of the random-product representation, property (c) appears to be the more natural condition to check in order to show $\tilde{\mathbf{P}}$ is the transition matrix of the time-reversal. Property (c) also gives us another way of trying to guess what the transition matrix of the time-reversal should look like: simply look for a transition matrix $\tilde{\mathbf{P}}$ where the random products governed by $\tilde{\mathbf{P}}$ become constant with probability one.

3 Continuous-time Markov Chains

Our objective in this section is to illustrate how the random-product representation can be used to establish the existence of invariant measures associated with a CTMC. This continuous-time setting will require a somewhat deeper study of these issues, and along the way, we will also show how the random-product representation can be used to show that the transition functions of (the minimal representation of) a CTMC satisfy the Kolmogorov Forward Equations.

3.1 Constructing a CTMC

It helps to first recall how to construct a continuous-time Markov chain: we follow the construction found in Chapter 2 of Asmussen [1]. Suppose first that $\{X_n\}_{n\geq 0}$ is a DTMC, having a countable state space $S \cup \{\partial\}$ and transition matrix $\mathbf{P} := [p(x,y)]_{x,y\in S\cup\{\partial\}}$. The state ∂ should be thought of as an outside absorbing state that will play an important role in a few of our derivations. Given that ∂ is an absorbing state, we have $p(\partial,\partial) = 1$, but readers should note that S could contain other absorbing states as well.

We further associate with $\{X_n\}_{n\geq 0}$ a collection of sojourn rates $\{q(x)\}_{x\in S\cup\{\partial\}} \subset [0,\infty)$, where $q(\partial) = 0$, as well as a sequence of i.i.d. exponentially distributed random variables $\{E_n\}_{n\geq 0}$ with rate 1, that are independent of $\{X_n\}_{n\geq 0}$. Together, $\{X_n\}_{n\geq 0}$, $\{q(x)\}_{x\in S\cup\{\partial\}}$ and $\{E_n\}_{n\geq 0}$ induce the set of transition times $\{T_n\}_{n\geq 0}$: define $T_0 := 0$, and for each integer $n \geq 1$, define

$$T_n := \sum_{k=0}^{n-1} \frac{E_k}{q(X_k)}. \tag{38}$$

It could be the case that $\sup_{n\geq 0} T_n < \infty$ with a positive probability, in which case we say that we have an *explosion* at time

$$T_\infty := \sup_{n\geq 0} T_n. \tag{39}$$

Having said all of this, we can define a continuous-time process $\{X(t); t \geq 0\}$ as follows:

$$X(t) = \begin{cases} \sum_{n=0}^{\infty} X_n \mathbf{1}_{\{T_n \leq t < T_{n+1}\}}, & t < T_\infty; \\ \Delta, & t \geq T_\infty. \end{cases} \tag{40}$$

where Δ is referred to as an external 'cemetery state', which can only be reached by $\{X(t); t \geq 0\}$ if the process explodes in finite time. Two important observations can be gleaned from this construction: (i) $\{X_n\}_{n\geq 0}$ is a discrete-time Markov chain, and (ii) conditional on $\{X_n\}_{n\geq 0}$, $\{T_{n+1} - T_n\}_{n\geq 0}$ are independent exponentially distributed random variables, where $T_{n+1} - T_n$ has rate $q(X_n)$.

It is possible to show that $\{X(t); t \geq 0\}$ is a Markov process, in that it satisfies the Markov property: for each integer $n \geq 0$, each collection of times $0 \leq t_1 < t_2 < \ldots < t_n < t_{n+1}$, and each collection of states $x_1, x_2, \ldots, x_n, x_{n+1}$, we have

$$\mathbb{P}(X(t_{n+1}) = x_{n+1} \mid X(t_n) = x_n, \ldots, X(t_1) = x_1) = \mathbb{P}(X(t_{n+1}) = x_{n+1} \mid X(t_n) = x_n). \tag{41}$$

Moreover, the transition probabilities of $\{X(t); t \geq 0\}$ are also time-homogeneous, meaning

$$\mathbb{P}(X(t_{n+1}) = x_{n+1} \mid X(t_n) = x_n) = \mathbb{P}(X(t_{n+1} - t_n) = x_{n+1} \mid X(0) = x_n). \tag{42}$$

Note too that $\{X(t); t \geq 0\}$ also satisfies a continuous-time version of the Strong Markov property: we will not need to make use of this result in our analysis, but we refer readers interested in learning more about the Strong Markov property to Chapter 1, Section 8 of [1].

We can use the sojourn rates $\{q(x)\}_{x\in S}$ and the transition matrix $\mathbf{P}$ to define the rate matrix $\mathbf{Q} := [q(x,y)]_{x,y\in S}$, where for each state $x \in S$,

$$q(x,x) := -q(x) \tag{43}$$

and for any two distinct states $x, y \in S$,

$$q(x,y) := q(x)p(x,y). \tag{44}$$

It is clear from our construction that $\mathbf{Q}$ is *stable*, meaning $|q(x,x)| = q(x) < \infty$ for each $x \in S$. Typically people also assume that $\mathbf{Q}$ is *conservative*, meaning

$$\sum_{y \in S} q(x,y) = 0$$

which amounts to assuming $p(x,\partial) = 0$ for each $x \in S$, but we will not need to make this assumption in our analysis.

3.2 Transition Functions

Further associated with $\{X(t); t \geq 0\}$ is a collection of transition functions $\{p_{x,y}\}_{x,y \in S}$, where for each $x, y \in S$ (where possibly $x = y$) $p_{x,y} : [0,\infty) \to [0,1]$ is defined as follows:

$$p_{x,y}(t) := \mathbb{P}(X(t) = y \mid X(0) = x), \quad t \geq 0. \tag{45}$$

These transition functions can be used to construct a collection of transition matrices $\{\mathbf{P}(t)\}_{t \geq 0}$, where for each $t \geq 0$,

$$\mathbf{P}(t) := [p_{x,y}(t)]_{x,y \in S}. \tag{46}$$

It is well-known that a sample path approach can be used to derive a system of integral equations satisfied by the transition functions of a CTMC. These facts are summarized in the following theorem: we omit the proof, as it is standard.

Theorem 3.1 *Suppose $x, y \in S$, where possibly $x = y$. Then for each $t > 0$,*

$$p_{x,y}(t) = e^{-q(x)t}\mathbf{1}(x = y) + \int_0^t \left[\sum_{z \neq x} q(x,z)p_{z,y}(t-s)\right] e^{-q(x)s} ds. \tag{47}$$

Furthermore, $p_{x,y}$ is continuous on $[0,\infty)$, differentiable on $(0,\infty)$, and

$$\lim_{h \downarrow 0} \frac{p_{x,y}(h) - \mathbf{1}(x = y)}{h} = q(x,y). \tag{48}$$

It is not difficult to use (47) to show that the transition functions satisfy what are known as the Kolmogorov Backward equations.

Theorem 3.2 (*Kolmogorov Backward Equations*) *For each $x, y \in S$, where possibly $x = y$,*

$$p'_{x,y}(t) = \sum_{z \in S} q(x,z)p_{z,y}(t) \tag{49}$$

for each $t > 0$.

Proof Equation (47) can also be expressed as

$$p_{x,y}(t)e^{q(x)t} = \mathbf{1}(x = y) + \int_0^t \left[\sum_{z \neq x} q(x,z)p_{z,y}(s)\right] e^{q(x)s} ds. \tag{50}$$

Taking derivatives of both sides of (50) and multiplying throughout by $e^{-q(x)t}$ yields (49). $\diamondsuit$

It is useful to restate the Kolmogorov Backward Equations in terms of the transition matrices of $\{X(t); t \geq 0\}$: doing so reveals that for each $t > 0$,

$$\mathbf{P}^{'}(t) = \mathbf{Q}\mathbf{P}(t)$$

with initial condition $\mathbf{P}(0) = \mathbf{I}$, and when S is further assumed to be finite, the transition matrices can be expressed in terms of the matrix exponential $e^{\mathbf{Q}t}$, meaning

$$\mathbf{P}(t) = e^{\mathbf{Q}t} := \sum_{k=0}^{\infty} \mathbf{Q}^k \frac{t^k}{k!}.$$

The transition functions of $\{X(t); t \geq 0\}$ also satisfy the Kolmogorov Forward Equations.

Theorem 3.3 (*Kolmogorov Forward Equations*) *For each $x, y \in S$, where possibly $x = y$,*

$$p^{'}_{x,y}(t) = \sum_{z \in S} p_{x,z}(t) q(z, y) \tag{51}$$

for each $t > 0$.

Again, the Forward Equations can alternatively be stated in matrix form as $\mathbf{P}^{'}(t) = \mathbf{P}(t)\mathbf{Q}$.

It is well-known that Theorem 3.3 is harder to establish than Theorem 3.2 when S is not finite, and in most textbooks extra assumptions are imposed on $\{X(t); t \geq 0\}$ in order to prove 3.3 rigorously. Later we will show that the random-product technique provides a way of deriving Theorem 3.3 rigorously, but it is important to remind readers that we are only considering chains that, upon explosion, reach a cemetery state and stay there from the explosion point onwards.

3.3 Laplace Transforms and the Random-Product Technique

Earlier we showed how the random-product technique can be used to study generating functions associated with sequences of n-step transition probabilities, so it is not too surprising that the random-product technique can also be used to study the Laplace transform of each transition function of $\{X(t); t \geq 0\}$. For each $x, y \in S$, where possibly $x = y$, let $\pi_{x,y}$ denote the Laplace transform of the transition function $p_{x,y}$. Clearly $\pi_{x,y}$ is well-defined on the open complex half-plane $\mathbb{C}_+ := \{\alpha \in \mathbb{C} : Re(\alpha) > 0\}$ as

$$\pi_{x,y}(\alpha) := \int_0^{\infty} e^{-\alpha t} p_{x,y}(t) dt, \quad \alpha \in \mathbb{C}_+. \tag{52}$$

Our first step towards analyzing each Laplace transform $\pi_{x,y}$ is to first express it in terms of hitting times associated with $\{X(t); t \geq 0\}$, just as we did for each generating function $\varphi_{x,y}$ in the previous section. Recall that we associate with $\{X_n\}_{n \geq 0}$ both the hitting times $\{\sigma_x\}_{x \in S}$ and the recurrence times $\{\eta_x\}_{x \in S}$, where for each state $x \in S$,

$$\eta_x := \inf\{n \geq 1 : X_n = x\}, \quad \sigma_x := \inf\{n \geq 0 : X_n = x\}.$$

A similar set of hitting and recurrence times can be associated with $\{X(t); t \geq 0\}$. For each state $x \in S$, we define the random variables τ_x and θ_x as

$$\tau_x := \inf\{t > 0 : X(t-) \neq X(t) = x\}, \quad \theta_x := \inf\{t \geq 0 : X(t) = x\}$$

where $X(t-) := \lim_{s \uparrow t} X(s)$ is the left-hand-limit of the CTMC at time t.

Lemma 3.1, given below, expresses the Laplace transforms of the transition functions of $\{X(t); t \geq 0\}$ in terms of hitting times associated with $\{X(t); t \geq 0\}$: this result appeared in [9], but it is also a special case of Theorem 4.16 on page 257 of [6]. The proof of this result is analogous to the proof of Lemma 2.1, so we omit it.

Lemma 3.1 *Fix a state $x \in S$. Then for each $y \in S$, and each $\alpha \in \mathbb{C}_+$,*

$$\pi_{x,y}(\alpha) = \frac{\mathbb{E}_x\left[\int_0^{\tau_x} e^{-\alpha t}\mathbf{1}(X(t) = y)dt\right]}{1 - \mathbb{E}_x[e^{-\alpha\tau_x}]} \tag{53}$$

and in particular,

$$\pi_{x,x}(\alpha) = \frac{1}{(q(x)+\alpha)(1 - \mathbb{E}_x[e^{-\alpha\tau_x}])}. \tag{54}$$

Moreover,

$$\pi_{x,y}(\alpha) = \pi_{x,x}(\alpha)w_{x,y}(\alpha) \tag{55}$$

where

$$w_{x,y}(\alpha) := (q(x)+\alpha)\mathbb{E}_x\left[\int_0^{\tau_x} e^{-\alpha t}\mathbf{1}(X(t) = y)dt\right]. \tag{56}$$

Lemma 3.2 itself has nothing to do with the random-product representation, but will be needed to establish other results later.

Lemma 3.2 *Suppose $\{X(t); t \geq 0\}$ is regular, and let $x \in S$. Then for each $\alpha \in \mathbb{C}_+$,*

$$1 = \sum_{y \in S} \alpha\pi_{y,y}(\alpha)\mathbb{E}_x[e^{-\theta\tau_y}]. \tag{57}$$

Proof Recall from (54) that for each $y \in S$,

$$\alpha\pi_{y,y}(\alpha) = \frac{\alpha}{(q(y)+\alpha)(1 - \mathbb{E}_y[e^{-\alpha\tau_y}])}. \tag{58}$$

Next, let $\tau_y^{(n)}$ denote the nth time $\{X(t); t \geq 0\}$ reaches state y: then by the Strong Markov property, coupled with (54), we get

$$\begin{aligned}
\mathbb{E}_x\left[\int_0^{\infty} \alpha e^{-\alpha t}\mathbf{1}(X(t) = y)dt\right] &= \sum_{n=1}^{\infty} \mathbb{E}_x\left[\int_{\tau_y^{(n)}}^{\tau_y^{(n+1)}} \alpha e^{-\alpha t}\mathbf{1}(X(t) = y)dt\right] \\
&= \sum_{n=1}^{\infty} \mathbb{E}_x[e^{-\alpha\tau_y^{(n)}}]\frac{\alpha}{(q(y)+\alpha)} \\
&= \sum_{n=1}^{\infty} \mathbb{E}_x[e^{-\alpha\theta_y}]\mathbb{E}_y[e^{-\alpha\tau_y}]^{n-1}\frac{\alpha}{(q(y)+\alpha)} \\
&= \frac{\alpha\mathbb{E}_x[e^{-\alpha\theta_y}]}{(q(y)+\alpha)(1 - \mathbb{E}_y[e^{-\alpha\tau_y}])} \\
&= \alpha\pi_{y,y}(\alpha)\mathbb{E}_x[e^{-\alpha\theta_y}].
\end{aligned} \tag{59}$$

Finally, since $\{X(t); t \geq 0\}$ is regular, it follows that

$$1 = \sum_{y \in S}\mathbb{E}_x\left[\int_0^{\infty} \alpha e^{-\alpha t}\mathbf{1}(X(t) = y)dt\right] = \sum_{y \in S}\alpha\pi_{y,y}(\alpha)\mathbb{E}_x[e^{-\alpha\theta_y}] \tag{60}$$

proving the lemma. $\diamondsuit$

Our objective now is to introduce a random-product representation for each $w_{x,y}(\alpha)$ term. Doing so involves constructing an alternative CTMC $\{\tilde{X}(t); t \geq 0\}$ with state space $S \cup \{\partial\}$ and generator $\tilde{\mathbf{Q}}$ which satisfies the following criteria:

(i) For each pair of states $(x, y) \in S \times S$ satisfying $x \neq y$, $\tilde{q}(x, y) > 0$ if and only if $q(y, x) > 0$.

(ii) For each state $x \in s$, $\tilde{q}(x, x) = q(x, x)$.

It is important to emphasize here that these results are not applied to the external absorbing state ∂, as it is not a member of S. We will want ∂ to play the role of external absorbing state for both $\{X(t); t \geq 0\}$ and $\{\tilde{X}(t); t \geq 0\}$, and statement (ii) above determines the value taken by $\tilde{q}(x, \partial)$ for each state $x \in S$.

It is also important to realize that in some situations, it may be impossible to select $\tilde{\mathbf{Q}}$ so that it satisfies these conditions. For instance, what if there is a state $x \in S$ such that $q(x) = 0$, yet $\sum_{y \neq x} q(y, x) > 0$? This situation causes problems: assuming x satisfies $q(x) = 0$, yet $\sum_{y \neq x} q(y, x) > 0$ we get

$$0 = q(x) = \tilde{q}(x) \geq \sum_{y \neq x} \tilde{q}(x, y) > 0$$

a contradiction. A similar situation arises if there exists some state $x \in S$ that satisfies $q(x) > 0$, yet $\sum_{y \neq x} q(y, x) = 0$. This case happens if state x is a state that can only be reached by $\{X(t); t \geq 0\}$ if $X(0) = x$. While we could simply set $\tilde{q}(x, \partial) = q(x)$ in order to get around the problem, there may be cases where we also want to force $\tilde{q}(x, \partial)$ to satisfy $\tilde{q}(x, \partial) = 0$, which would again make it impossible to find a $\tilde{\mathbf{Q}}$ that satisfies the needed conditions.

These two nuisance cases must be accounted for in order to precisely introduce the random-product technique, and this can be done if we now construct a new CTMC $\{X^e(t); t \geq 0\}$ with an enlarged state space S^e containing $S \cup \{\partial\}$, and with a new generator matrix $\mathbf{Q}^e := [q^e(x, y)]_{x,y \in S^e}$.

Here is how S^e is constructed. For each state $x \in S$ that satisfies $q(x) = 0$, yet $\sum_{y \neq x} q(y, x) > 0$, we introduce within S^e a new state x^*—a 'copy' of x—and we set

$$q^e(x, x^*) = 1, \quad q^e(x^*, x) = 1.$$

Furthermore, for each state $y \neq x$ that satisfies $q(y, x) > 0$, we define $q^e(y, x) := q(y, x)/2$, and $q^e(y, x^*) := q(y, x)/2$. It is easy to see that for each state $y \neq x, x^*$,

$$\mathbb{P}_y(X^e(t) = x) = \mathbb{P}_y(X^e(t) = x^*)$$

for each real $t \geq 0$.

Next, for each $x \in S$ satisfying $q(x) > 0$, yet $\sum_{y \neq x} q(y, x) = 0$, we introduce within S^e a copy x^* of x, and we set, for each state $y \neq x, x^*$

$$q^e(x, y) = q(x, y), \qquad q^e(x^*, y) = q(x, y), \qquad q^e(x, x^*) = q^e(x^*, x) = 1.$$

Under these conditions, for each state $y \neq x, x^*$,

$$\mathbb{P}_x(X^e(t) = y) = \mathbb{P}_{x^*}(X^e(t) = y)$$

and furthermore,

$$\mathbb{P}_x(X^e(t) = x^*) = \mathbb{P}_{x^*}(X^e(t) = x), \qquad \mathbb{P}_x(X^e(t) = x) = \mathbb{P}_{x^*}(X^e(t) = x^*)$$

for each $t \geq 0$.

Theorem 3.4 provides us with a random-product representation for each $w_{x,y}(\alpha)$ term. While it may be necessary to extend X to X^e in order to ensure these random-product expressions are well-defined, in the interests of making the result simpler to state we assume in Theorem 3.4 that X does not need to be extended to X^e.

Theorem 3.4 *For each $\alpha \in \mathbb{C}_+ \cup \{0\}$, we have that for each $x, y \in S$ satisfying $x \neq y$,*

$$w_{x,y}(\alpha) = \mathbb{E}_y \left[\mathbf{1}(\tilde{\sigma}_x < \infty) e^{-\alpha \tilde{\theta}_x} \prod_{\ell=1}^{\tilde{\sigma}_x} \frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right]. \tag{61}$$

Proof The proof of this result can also be found in [9], but we provide it here for convenience. Fix two states $x, y \in S$, where $x \neq y$, and observe that for each $\alpha \in \mathbb{C}_+ \cup \{0\}$,

$$\begin{aligned}
&\mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)e^{-\alpha\tilde{\theta}_x}\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] \\
&= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n : x_0=y, x_n=x, x_1,\ldots,x_{n-1}\neq x}\left[\prod_{\ell=1}^{n}\frac{\tilde{q}(x_{\ell-1})}{\tilde{q}(x_{\ell-1})+\alpha}\right]\left[\prod_{\ell=1}^{n}\frac{q(x_\ell, x_{\ell-1})}{\tilde{q}(x_{\ell-1}, x_\ell))}\right]\left[\prod_{\ell=1}^{n}\frac{\tilde{q}(x_{\ell-1}, x_\ell)}{\tilde{q}(x_{\ell-1})}\right] \\
&= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n : x_0=y, x_n=x, x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{q(x_\ell, x_{\ell-1})}{q(x_{\ell-1})+\alpha} \\
&= \frac{q(x)+\alpha}{q(y)+\alpha}\sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n : x_0=y, x_n=x, x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{q(x_\ell, x_{\ell-1})}{q(x_\ell)+\alpha} \\
&= (q(x)+\alpha)\sum_{n=1}^{\infty}\mathbb{E}_x[e^{-\alpha T_n}\mathbb{E}[1-e^{-\alpha(T_{n+1}-T_n)} \mid X_0,\ldots,X_n, T_1,\ldots,T_n]\mathbf{1}(X_n = y, \eta_x > n)] \\
&= (q(x)+\alpha)\mathbb{E}_x\left[\sum_{n=1}^{\eta_x-1}\int_{T_n}^{T_{n+1}}e^{-\alpha t}\mathbf{1}(X(t)=y)dt\right] \\
&= (q(x)+\alpha)\mathbb{E}_x\left[\int_0^{\tau_x}e^{-\alpha t}\mathbf{1}(X(t)=y)dt\right] \\
&= w_{x,y}(\alpha)
\end{aligned} \tag{62}$$

which establishes (61). ◊

Our next result shows that the transition functions of $\{X(t); t \geq 0\}$ satisfy the Kolmogorov Forward equations. It should also be noted that this is the only place where we need to make use of the extended CTMC X^e. While usage of X^e is admittedly not so elegant, it is arguably more elementary than the point process approaches considered in e.g. Brémaud [3], and Last and Brandt [18].

Theorem 3.5 *For each state $x \in S$, we have*

$$\alpha\pi_{x,y}(\alpha) - \mathbf{1}(x = y) = \sum_{z\in S}\pi_{x,z}(\alpha)q(z, y). \tag{63}$$

Equation (63) is a Laplace-transform version of (51), so by the uniqueness property of Laplace transforms, in order to prove (51) it suffices to prove (63).

Proof Assume first that X coincides with its extended version X^e, and suppose $x, y \in S$ are two distinct states. Then from Theorem 3.4, we have that for each $\alpha \in \mathbb{C}_+ := \{\alpha \in \mathbb{C} : Re(\alpha) > 0\}$,

$$\pi_{x,y}(\alpha) = \pi_{x,x}(\alpha)\mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)e^{-\alpha\tilde{\theta}_x}\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right].$$

Next, note that conditioning on $\tilde{X}_1$ yields

$$\mathbb{E}_y\left[\mathbf{1}(\tilde{\sigma}_x < \infty)e^{-\alpha\tilde{\theta}_x}\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] = \sum_{z\neq y}\frac{q(z,y)}{q(y)+\alpha}\mathbb{E}_z\left[\mathbf{1}(\tilde{\sigma}_x < \infty)e^{-\alpha\tilde{\theta}_x}\prod_{\ell=1}^{\tilde{\sigma}_x}\frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right]$$

and multiplying both sides by $\pi_{x,x}(\alpha)$ and $(q(y)+\alpha)$ further gives

$$(q(y)+\alpha)\pi_{x,y}(\alpha) = \sum_{z\neq y}\pi_{x,z}(\alpha)q(z,y)$$

or, equivalently

$$\alpha\pi_{x,y}(\alpha) = \sum_{z\in S} \pi_{x,z}(\alpha)q(z,y)$$

which establishes (63) for the case where $x \neq y$.

It remains to consider the case where $x = y$. First, recall from Lemma 3.1 that for each state $z \neq x$,

$$\pi_{x,z}(\alpha) = \pi_{x,x}(\alpha)(q(x)+\alpha)\mathbb{E}_x\left[\int_0^{\tau_x} e^{-\alpha t}\mathbf{1}(X(t)=z)dt\right].$$

Using a technique analogous to the 'cycle-trick' gives

$$\begin{aligned}
\sum_{z\neq x}\pi_{x,z}(\alpha)q(z,x) &= \pi_{x,x}(\alpha)(q(x)+\alpha)\sum_{z\neq x}\mathbb{E}_x\left[\int_0^{\tau_x} e^{-\alpha t}\mathbf{1}(X(t)=z)dt\right]q(z,x)\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\sum_{z\neq x}\mathbb{E}_x\left[\sum_{n=1}^{\infty}\mathbf{1}(X_n=z,\eta_x>n)\int_{T_n}^{T_{n+1}} e^{-\alpha t}\right]q(z,x)\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\sum_{z\neq x}\mathbb{E}_x\left[\sum_{n=1}^{\infty}\mathbf{1}(X_n=z,\eta_x>n)e^{-\alpha T_n}(1-e^{-\alpha(T_{n+1}-T_n)})\right]q(z,x)\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\sum_{z\neq x}\mathbb{E}_x\left[\sum_{n=1}^{\infty}\mathbf{1}(X_n=z,\eta_x>n)e^{-\alpha T_n}\right]\frac{q(z,x)}{q(z)+\alpha}\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\sum_{z\neq x}\mathbb{E}_x\left[\sum_{n=1}^{\infty}\mathbf{1}(X_1\neq x,\ldots,X_{n-1}\neq x,X_n=z,X_{n+1}=x)e^{-\alpha T_{n+1}}\right]\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\mathbb{E}_x\left[\sum_{n=1}^{\infty}\mathbf{1}(X_1\neq x,\ldots,X_{n-1}\neq x,X_n\neq x,X_{n+1}=x)e^{-\alpha T_{n+1}}\right]\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)\mathbb{E}_x[e^{-\tau_x}]\\
&= \pi_{x,x}(\alpha)(q(x)+\alpha)-\pi_{x,x}(\alpha)(q(x)+\alpha)(1-\mathbb{E}_x[e^{-\tau_x}]).
\end{aligned}$$

Applying now Lemma 3.1 gives

$$\sum_{z\in S}\pi_{x,z}(\alpha)q(x,z) = \alpha\pi_{x,x}(\alpha)-1$$

which is (63) for the case where $x = y$.

It remains to prove that the transition functions of $\{X(t); t\geq 0\}$ satisfy the Kolmogorov Forward Equations, when S^e is strictly larger than S. Letting $\{p^{(e)}_{x,y}\}_{x,y\in S^e}$ denote the transition functions of $\{X^e(t); t\geq 0\}$, we find that

$$\begin{aligned}
p'_{x,x}(t) = \frac{d}{dt}\left[p^{(e)}_{x,x}(t)+p^{(e)}_{x,x^*}(t)\right] &= \sum_{y\in S^e} p^{(e)}_{x,y}(t)q^e(y,x) + \sum_{y\in S^e} p_{x,y}(t)q^e(y,x^*)\\
&= p^{(e)}_{x,x}(t)q^e(x,x)+p^{(e)}_{x,x^*}(t)q^e(x^*,x)+p^{(e)}_{x,x}(t)q^e(x,x^*)+p^{(e)}_{x,x^*}(t)q^e(x^*,x^*)\\
&= p_{x,x}(t)q^e(x,x^*)+p_{x,x}(t)q^e(x,x)\\
&= p_{x,x}(t)q(x,x)
\end{aligned}$$

where readers should remember that $q^e(x,x) = q^e(x^*,x^*) = q(x,x) - q^e(x,x^*)$. The other equations for $\{X(t); t\geq 0\}$ can be established in an analogous manner: we omit the details. $\diamondsuit$

3.4 Stationary Distributions and Invariant Measures

Our next objective is to use the random-product technique to construct invariant measures and stationary distributions associated with $\{X(t); t \geq 0\}$. The study of stationary distributions of CTMCs is made complicated by the fact that it can be difficult to find a solution $\mathbf{p}$ that satisfies $\mathbf{p} = \mathbf{p}\mathbf{P}(t)$ for each real $t \geq 0$, due to the difficulty involved with calculating $\mathbf{P}(t)$. Fortunately, the random-product technique provides us with a simple way of comparing the set of all positive solutions ν of $\boldsymbol{\nu} = \boldsymbol{\nu}\mathbf{P}(t)$ for each $t \geq 0$, with the set of all positive solutions $\mathbf{p}$ of the (much simpler) linear system $\mathbf{pQ} = \mathbf{0}$.

Definition 3.1 *Given a state $x \in S$, we say that x is* transient *if $\mathbb{P}_x(\tau_x < \infty) < 1$. Likewise, we say that x is* recurrent *if $\mathbb{P}_x(\tau_x < \infty) = 1$.*

Most readers know that $\mathbb{P}_x(\eta_x < \infty) = \mathbb{P}_x(\tau_x < \infty)$ for each state $x \in S$, meaning we can determine all of the transient and recurrent states for $\{X(t); t \geq 0\}$ by instead finding all of the transient and recurrent states of its embedded DTMC $\{X_n\}_{n\geq 0}$. Similarly, we say that $\{X(t); t \geq 0\}$ is irreducible if and only if its embedded DTMC $\{X_n\}_{n\geq 0}$ is irreducible, and the communicating classes of $\{X(t); t \geq 0\}$ are simply the communicating classes of $\{X_n\}_{n\geq 0}$.

Definition 3.2 *A recurrent state $x \in S$ is said to be* positive recurrent *if $\mathbb{E}_x[\tau_x] < \infty$, and* null recurrent *if $\mathbb{E}_x[\tau_x] = \infty$.*

In order to link the set of all positive solutions $\mathbf{m}$ of $\mathbf{m} = \mathbf{mP}(t)$, for each $t \geq 0$, with the set of all possible solutions of $\mathbf{mQ} = \mathbf{0}$, we need to define two types of invariant measures. A few of our proofs will also require the use of a related object, known in the literature as a subinvariant measure.

Definition 3.3 *A row vector $\mathbf{m} := [m(y)]_{y\in S}$ is an* invariant measure with respect to $\mathbf{Q}$ *if $\mathbf{mQ} = \mathbf{0}$, and a* subinvariant measure with respect to $\mathbf{Q}$ *if $\mathbf{mQ} \leq \mathbf{0}$. Furthermore, $\mathbf{m}$ is an* invariant measure with respect to $\mathbf{P}$ *if for each $t \geq 0$, $\mathbf{m} = \mathbf{mP}(t)$, and an invariant measure $\mathbf{p}$ with respect to $\mathbf{P}$ is a stationary distribution if it has total mass one (meaning its components sum to one).*

Our goal now is to show that a stationary distribution $\mathbf{p}$ of $\{X(t); t \geq 0\}$ can be calculated by solving the linear system $\mathbf{pQ} = \mathbf{0}$, but justifying this procedure requires several steps, as well as the introduction of a CTMC that behaves in many ways like the time-reversal of an ergodic CTMC.

Definition 3.4 *Given a CTMC $\{X(t); t \geq 0\}$ with state space $S \cup \{\partial\}$, generator $\mathbf{Q}$, and positive subinvariant measure $\mathbf{m}$ with respect to $\mathbf{Q}$, let $\overline{X}$ denote the $\mathbf{m}$-dual $\{\overline{X}(t); t \geq 0\}$ whose state space $S \cup \{\partial\}$ and generator $\overline{\mathbf{Q}}$ satisfy*

$$\overline{q}(x, y) := \frac{m(y)q(y, x)}{m(x)}, \qquad x, y \in S.$$

Note again that this definition determines each value of $\overline{q}(x, \partial)$, for each state $x \in S$, where ∂ is the external absorbing state of $\{\overline{X}(t); t \geq 0\}$. Furthermore, for each $x \in S$,

$$\begin{aligned}
\overline{q}(x, \partial) = -\overline{q}(x, x) - \sum_{y\in S: y\neq x} \overline{q}(x, y) &= -q(x, x) - \frac{1}{m(x)} \sum_{y\in S: y\neq x} m(y)q(y, x) \\
&\geq -q(x, x) + q(x, x) \\
&= 0.
\end{aligned}$$

It is instructive to understand how the distributions of the recurrence times $\{\overline{\tau}_x\}_{x\in S}$ associated with $\{\overline{X}(t); t \geq 0\}$ relate to the distributions of the recurrence times $\{\tau_x\}_{x\in S}$ associated with $\{X(t); t \geq 0\}$. This issue is studied in Lemma 3.3.

Lemma 3.3 *The following equalities hold for each $x, y \in S$, and each $\alpha \in \mathbb{C}_+$:*

$$\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)] = \frac{m(y)}{m(x)}\left(\frac{q(y)+\alpha}{q(x)+\alpha}\right)\mathbb{E}_y[e^{-\alpha\tau_x}\mathbf{1}(\tau_x < \tau_y)]. \tag{64}$$

$$\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)] = \mathbb{E}_x[e^{-\alpha\tau_x}\mathbf{1}(\tau_x < \tau_y)] \tag{65}$$

$$\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}] = \frac{\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}{1-\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]} \tag{66}$$

$$\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}] = \mathbb{E}_x[e^{-\alpha\tau_x}]. \tag{67}$$

$$\overline{\pi}_{x,x}(\alpha) = \pi_{x,x}(\alpha) \tag{68}$$

$$1-\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}] = \left[\frac{1-\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]}{1-\mathbb{E}_y[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}\right](1-\mathbb{E}_y[e^{-\alpha\overline{\tau}_y}]). \tag{69}$$

$$\overline{\pi}_{x,x}(\alpha) = \left[\frac{q(y)+\alpha}{q(x)+\alpha}\right]\left[\frac{1-\mathbb{E}_y[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}{1-\mathbb{E}_x[e^{-\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]}\right]\overline{\pi}_{y,y}(\alpha). \tag{70}$$

Proof We begin by proving (64): observe that

$$\begin{aligned}
\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)] &= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n:x_0=x,x_n=y,x_1,\ldots,x_{n-1}\neq x,y}\prod_{\ell=1}^{n}\frac{\overline{q}(x_{\ell-1},x_\ell)}{\overline{q}(x_{\ell-1})+\alpha} \\
&= \frac{m(y)}{m(x)}\left[\frac{q(y)+\alpha}{q(x)+\alpha}\right]\sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n:x_0=x,x_n=y,x_1,\ldots,x_{n-1}\neq x,y}\prod_{\ell=1}^{n}\frac{q(x_\ell,x_{\ell-1})}{q(x_{\ell-1})+\alpha} \\
&= \frac{m(y)}{m(x)}\left[\frac{q(y)+\alpha}{q(x)+\alpha}\right]\mathbb{E}_y[e^{-\alpha\tau_x}\mathbf{1}(\tau_x < \tau_y)].
\end{aligned}$$

A similar path-summation argument can be used to establish (65).

Identity (66) follows from an application of the strong Markov property:

$$\begin{aligned}
\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}] &= \mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)] + \mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)] \\
&= \mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)] + \mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}]
\end{aligned}$$

which implies (66).

Identity (67) can be derived with another path-summation argument: here

$$\begin{aligned}
\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}] &= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n:x_0=x_n=x,x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{\overline{q}(x_{\ell-1},x_\ell)}{\overline{q}(x_{\ell-1})+\alpha} \\
&= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n:x_0=x_n=x,x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{q(x_\ell,x_{\ell-1})}{q(x_\ell)+\alpha} \\
&= \mathbb{E}_x[e^{-\alpha\tau_x}].
\end{aligned}$$

Identity (68) follows as an immediate consequence of Identity (67).

We next establish Identity (69): here

$$\begin{aligned}
\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}] &= \mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)] + \mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)] \\
&= \mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)] + \mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]\mathbb{E}_y[e^{-\alpha\overline{\tau}_x}] \\
&= \mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)] + \frac{\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]\mathbb{E}_y[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]}{1 - \mathbb{E}_y[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}
\end{aligned}$$

and after a little algebra, we arrive at (69). Finally, Identity (70) is a simple consequence of Identity (69).$\diamondsuit$

One interesting consequence of Lemma 3.3 is that if $\{X(t); t \geq 0\}$ is irreducible and recurrent, with subinvariant measure $\mathbf{m}$, then its $\mathbf{m}$-dual $\{\overline{X}(t); t \geq 0\}$ must also be irreducible and recurrent. This is shown in Lemma 3.4.

Lemma 3.4 *Suppose $\{X(t); t \geq 0\}$ is irreducible and recurrent, with positive subinvariant measure $\mathbf{m}$. Then $\{\overline{X}(t); t \geq 0\}$ must also be irreducible and recurrent, and $\mathbf{m}$ is an invariant measure of $\{X(t); t \geq 0\}$ with respect to $\mathbf{Q}$.*

Proof Irreducibility of $\{\overline{X}(t); t \geq 0\}$ follows immediately from irreducibility of $\{X(t); t \geq 0\}$, since any feasible path $x = x_0, x_1, \ldots, x_n = y$ under $\{X(t); t \geq 0\}$ corresponds to a feasible path $y = x_n, x_{n-1}, \ldots, x_1, x_0 = x$ under $\{\overline{X}(t); t \geq 0\}$. Recurrence of $\{\overline{X}(t); t \geq 0\}$ follows from (67), which shows that for each state $x \in S$, the distribution of $\overline{\tau}_x$ under $\mathbb{P}(\cdot \mid \overline{X}(0) = x)$ is the same as the distribution of τ_x under the probability measure $\mathbb{P}(\cdot \mid X(0) = x)$. Finally, observe that because all states of $\{\overline{X}(t); t \geq 0\}$ are recurrent, it must be the case that for each $x \in S$, $\overline{q}(x, \partial) = 0$, which in turn means

$$\begin{aligned}
-q(x,x) = -\overline{q}(x,x) &= \overline{q}(x,\partial) + \sum_{y \in S: y \neq x} \overline{q}(x,y) \\
&= \frac{1}{m(x)} \sum_{y \neq x} m(y)q(y,x)
\end{aligned}$$

which in turn implies

$$-m(x)q(x,x) = \sum_{y \in S: y \neq x} m(y)q(y,x)$$

or, equivalently,

$$0 = \sum_{y \in S} m(y)q(y,x)$$

i.e. $\mathbf{m}$ is an invariant measure of $\{X(t); t \geq 0\}$ with respect to $\mathbf{Q}$. $\diamondsuit$

Theorem 3.6 *Suppose $\mathbf{m}$ is a positive subinvariant measure with respect to $\mathbf{Q}$, and let $\{\tilde{X}(t); t \geq 0\}$ be another CTMC with state space $S \cup \{\partial\}$. Then the following statements are equivalent:*

(i) $\tilde{\mathbf{Q}}$ *satisfies $m(x)\tilde{q}(x,y) = m(y)q(y,x)$;*

(ii) the transition functions $\{\tilde{p}_{x,y}\}_{x,y \in S}$ associated with $\{\tilde{X}(t); t \geq 0\}$ and $\{X(t); t \geq 0\}$ satisfy $m(x)\tilde{p}_{x,y}(t) = m(y)p_{y,x}(t)$ for each real $t \geq 0$.

Proof Clearly statement (ii) implies statement (i), since for each $x, y \in S$ satisfying $x \neq y$,

$$m(x)\tilde{q}(x,y) = \lim_{h \downarrow 0} m(x)\frac{\tilde{p}_{x,y}(h)}{h} = \lim_{h \downarrow 0} m(y)\frac{p_{y,x}(h)}{h} = m(y)q(y,x).$$

Next, assume statement (i) holds, which of course means that $\{\tilde{X}(t); t \geq 0\}$ is actually the $\mathbf{m}$-dual $\{\overline{X}(t); t \geq 0\}$ of $\{X(t); t \geq 0\}$. In order to prove (ii) for $x \neq y$, it suffices, due to the uniqueness property of Laplace transforms, to show that for each $\alpha \in \mathbb{C}_+$,

$$m(x)\overline{\pi}_{x,y}(\alpha) = m(y)\pi_{y,x}(\alpha).$$

Through the application of various facts contained in Lemma 3.3, we find that for each $x, y \in S$ satisfying $x \neq y$,

$$\begin{aligned}
m(x)\overline{\pi}_{x,y}(\alpha) &= m(x)\overline{\pi}_{x,x}(\alpha)\mathbb{E}_y\left[\mathbf{1}(\sigma_x < \infty)e^{-\alpha\theta_x}\prod_{\ell=1}^{\sigma_x}\frac{\overline{q}(X_\ell, X_{\ell-1})}{q(X_{\ell-1}, X_\ell)}\right] \\
&= m(y)\overline{\pi}_{x,x}(\alpha)\mathbb{E}_y[e^{-\alpha\sigma_x}] \\
&= m(y)\overline{\pi}_{y,y}(\alpha)\left[\frac{\overline{q}(y)+\alpha}{\overline{q}(x)+\alpha}\right]\left[\frac{1-\mathbb{E}_y[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}{1-\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]}\right]\frac{\mathbb{E}_y[e^{-\alpha\tau_x}\mathbf{1}(\tau_x < \tau_y)]}{1-\mathbb{E}_y[e^{-\alpha\tau_y}\mathbf{1}(\tau_y < \tau_x)]} \\
&= m(x)\overline{\pi}_{y,y}(\alpha)\frac{\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}\mathbf{1}(\overline{\tau}_y < \overline{\tau}_x)]}{1-\mathbb{E}_x[e^{-\alpha\overline{\tau}_x}\mathbf{1}(\overline{\tau}_x < \overline{\tau}_y)]} \\
&= m(x)\pi_{y,y}(\alpha)\mathbb{E}_x[e^{-\alpha\overline{\tau}_y}] \\
&= m(y)\pi_{y,x}(\alpha)
\end{aligned}$$

proving (ii). $\diamondsuit$

Theorem 3.6 can now be used to give a quick proof of Theorem 3.7, which first appeared in Kelly [17]. Note that half of our proof of Theorem 3.7 makes use of the random-product technique and Lemma 3.2, and the other half makes use of Lemma 3.3: unlike the arguments given in [17], our approach does not require the use of forward and backward integral recursions.

Theorem 3.7 *Suppose* $\mathbf{p} = [p(x)]_{x\in S}$ *is a positive row vector. Then the following statements are equivalent:*
(a) $\mathbf{p}$ *is an invariant measure for* $\{X(t); t \geq 0\}$.
(b) $\mathbf{p}$ *satisfies* $\mathbf{pQ} = \mathbf{0}$, *and the* $\mathbf{p}$*-dual* $\{\overline{X}(t); t \geq 0\}$ *is regular.*

Proof Assume first that statement (a) is true. Using Fatou's Lemma, we see that for each state $x \in S$,

$$\sum_{y\in S} p(y)q(y,x) \leq 0$$

which proves $\mathbf{p}$ is a positive subinvariant measure with respect to $\mathbf{Q}$. Letting $\{\overline{X}(t); t \geq 0\}$ denote the $\mathbf{p}$-dual of $\{X(t); t \geq 0\}$, we find through an application of Theorem 3.6 that for each $x, y \in S$,

$$p(x)\overline{p}_{x,y}(t) = p(y)p_{y,x}(t) \tag{71}$$

and summing (71) over all $y \in S$ reveals

$$p(x)\sum_{y\in S}\overline{p}_{x,y}(t) = \sum_{y\in S} p(y)p_{y,x}(t) = p(x)$$

meaning

$$\sum_{y\in S}\overline{p}_{x,y}(t) = 1$$

proving $\{\overline{X}(t); t \geq 0\}$ with state space S is regular, and $\overline{\mathbf{Q}}$ is conservative, i.e. $\mathbf{pQ} = \mathbf{0}$.

Assume now that statement (b) is true. Then by Lemma 3.2, we get

$$\begin{aligned}\sum_{y\in S} p(y)\alpha\pi_{y,x}(\alpha) &= \sum_{y\in S} p(y)\alpha\pi_{y,y}(\alpha)\mathbb{E}_x\left[e^{-\alpha\overline{\theta}_y}\prod_{\ell=1}^{\overline{\sigma}_y}\frac{q(\overline{X}_\ell,\overline{X}_{\ell-1})}{\overline{q}(\overline{X}_{\ell-1},\overline{X}_\ell)}\right]\\ &= \sum_{y\in S} p(y)\alpha\overline{\pi}_{y,y}(\alpha)\frac{p(x)}{p(y)}\mathbb{E}_x[e^{-\alpha\overline{\theta}_y}]\\ &= p(x)\sum_{y\in S}\alpha\overline{\pi}_{y,y}(\alpha)\mathbb{E}_x[e^{-\alpha\overline{\theta}_y}]\\ &= p(x)\end{aligned}$$

which proves (a). $\diamondsuit$

Our next result is a corollary of Theorem 3.7, as it shows that when $\{X(t); t \geq 0\}$ is further assumed to be both irreducible and recurrent, the set of all positive solutions to $\mathbf{p} = \mathbf{p}\mathbf{P}(t)$ for all $t \geq 0$ coincides with the set of all positive probability measures $\mathbf{p}$ on S that satisfy $\mathbf{p}\mathbf{Q} = \mathbf{0}$.

Corollary 3.1 *Suppose $\{X(t); t \geq 0\}$ is irreducible and recurrent, and let $\mathbf{p}$ is a positive probability measure on S. Then $\mathbf{p}$ satisfies $\mathbf{p}\mathbf{Q} = \mathbf{0}$ if and only if for each real $t \geq 0$, $\mathbf{p} = \mathbf{p}\mathbf{P}(t)$.*

Proof Assume first that $\mathbf{p}$ is a positive invariant measure with respect to $\mathbf{P}$. Then by Theorem 3.7, $\mathbf{p}$ must satisfy $\mathbf{p}\mathbf{Q} = 0$. Next, observe that if $\mathbf{p}$ is a positive invariant measure with respect to $\mathbf{Q}$, we can use Lemma 3.4 to conclude that the $\mathbf{p}$-dual $\{\overline{X}(t); t \geq 0\}$ is also irreducible and recurrent, which further implies that $\{\overline{X}(t); t \geq 0\}$ is regular. Thus, by Theorem 3.7, $\mathbf{p}$ is a positive invariant measure with respect to $\mathbf{P}$. $\diamondsuit$

Now that we fully understand the relationship between the set of all positive invariant measures with respect to $\mathbf{Q}$, and the set of all positive invariant measures with respect to $\mathbf{P}$, we can turn to the problem of constructing invariant measures of $\{X(t); t \geq 0\}$ with recurrent states. Our next result, Theorem 3.8, is well-known, as it shows how a recurrent state can be used to construct an invariant measure with respect to $\mathbf{Q}$, and that if $\{X(t); t \geq 0\}$ is irreducible and recurrent, then all positive invariant measures with respect to $\mathbf{Q}$ (and with respect to $\mathbf{P}$) are scalar multiples of each other. Again, the novelty here involves observing how the random-product technique yields an extremely simple derivation of this known result.

Theorem 3.8 *Fix a recurrent state $x \in S$. Then the row vector $\boldsymbol{\nu}^{(x)} := [\nu^{(x)}(y)]_{y\in S}$, defined as*

$$\nu^{(x)}(y) := w_{x,y}(0) = q(x)\mathbb{E}_x\left[\int_0^{\tau_x}\mathbf{1}(X(t) = y)dt\right], \quad y \in S$$

is an invariant measure with respect to $\mathbf{Q}$. Furthermore, if $\{X(t); t \geq 0\}$ is both irreducible and positive recurrent, each positive invariant measure $\boldsymbol{\nu}$ with respect to $\mathbf{Q}$ is a positive scalar multiple of $\nu^{(x)}$.

Proof Fix a recurrent state $x \in S$. Recall that the communicating class C_x containing state x is closed, so without loss of generality we can assume throughout that $S = C_x$, which in turn makes $\{X(t); t \geq 0\}$ both irreducible and recurrent. Our goal, then, is to show that for each state $y \in C_x$,

$$\nu^{(x)}(y)q(y) = \sum_{z\in C_x: z\neq y}\nu^{(x)}(z)q(z,y). \tag{72}$$

The key towards establishing (72) is to first show that for each state $y \in C_x$ (where possibly $y = x$),

$$w_{x,y}(0) = \mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)}\right]. \tag{73}$$

This identity is trivial for the case where $y \neq x$, because under the measure $\mathbb{P}_y$, $\tilde{\eta}_x = \tilde{\sigma}_x$. Next, observe that

$$\begin{aligned}
\mathbb{E}_x\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)}\right] &= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n: x_0=x_n=x, x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{q(x_\ell,x_{\ell-1})}{q(x_{\ell-1})} \\
&= \sum_{n=1}^{\infty}\sum_{x_0,\ldots,x_n: x_0=x_n=x, x_1,\ldots,x_{n-1}\neq x}\prod_{\ell=1}^{n}\frac{q(x_\ell,x_{\ell-1})}{q(x_\ell)} \\
&= \mathbb{P}_x(\eta_x < \infty) \\
&= 1 \\
&= w_{x,x}(0).
\end{aligned}$$

We are now equipped to quickly establish (72): for each state $y \in S$, we can use first-step analysis to show that

$$\begin{aligned}
\nu^{(x)}(y) &= \sum_{z\neq y}\mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=1}^{\tilde{\eta}_x}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)} \mid \tilde{X}_1 = z\right]\frac{\tilde{q}(y,z)}{q(y)} \\
&= \sum_{z\neq y}\mathbb{E}_y\left[\mathbf{1}(\tilde{\eta}_x < \infty)\prod_{\ell=2}^{\tilde{\eta}_x}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)} \mid \tilde{X}_1 = z\right]\frac{q(z,y)}{\tilde{q}(y,z)}\frac{\tilde{q}(y,z)}{q(y)} \\
&= \sum_{z\neq y}\nu^{(x)}(z)\frac{q(z,y)}{q(y)}
\end{aligned}$$

where the last statement follows from the Markov property. Hence,

$$\nu^{(x)}(y) = \sum_{z\neq y}\nu^{(x)}(z)\frac{q(z,y)}{q(y)}$$

or, equivalently,

$$\nu^{(x)}(y)q(y) = \sum_{z\neq y}\nu^{(x)}(z)q(z,y)$$

which establishes (72) for each state $y \in S$.

Assuming now that $\{X(t); t \geq 0\}$ is both irreducible and recurrent, our goal is to show that each positive invariant measure $\boldsymbol{\nu}$ is a positive scalar multiple of the positive invariant measure $\boldsymbol{\nu}^{(x)}$, where x is an arbitrary chosen recurrent state within S. If we apply the random-product technique while choosing $\{\tilde{X}(t); t \geq 0\}$ to be the $\boldsymbol{\nu}$-dual $\{\overline{X}(t); t \geq 0\}$, we find that for each $y \in S$,

$$\nu^{(x)}(y) = \frac{\nu(y)}{\nu(x)}\mathbb{P}_y(\overline{\sigma}_x < \infty) = \frac{\nu(y)}{\nu(x)}$$

proving ν is a positive scalar multiple of $\nu^{(x)}$. $\diamondsuit$

Our next result, Theorem 3.9, shows that if $\{X(t); t \geq 0\}$ is both irreducible and regular, and has a positive invariant measure $\mathbf{p}$ with respect to $\mathbf{Q}$, then each state must be positive recurrent. Asmussen discusses in Chapter 2 of [1] how regularity must be assumed in the hypothesis of this result, and we feel that our derivation of this result (which, again, makes use of the random-product technique) further clarifies why this condition is needed.

Theorem 3.9 *Suppose $\{X(t); t \geq 0\}$ is irreducible and nonexplosive, and there exists a probability measure $\mathbf{p} := [p(y)]_{y\in S}$ satisfying $p(y) > 0$ for each $y \in S$, and $\mathbf{pQ} = \mathbf{0}$. Then $\{X(t); t \geq 0\}$ is positive recurrent.*

Proof Let $\{\overline{X}(t); t \geq 0\}$ denote the **p**-dual of $\{X(t); t \geq 0\}$. Fix a state $x \in S$, and observe that for each state $y \neq x$,

$$\begin{aligned}\nu^{(x)}(y) &= \mathbb{E}_y\left[\mathbf{1}(\overline{\sigma}_x < \infty)\prod_{\ell=1}^{\overline{\sigma}_x}\frac{q(\overline{X}_\ell, \overline{X}_{\ell-1})}{\overline{q}(\overline{X}_{\ell-1}, \overline{X}_\ell)}\right] \\ &= \mathbb{E}_y\left[\mathbf{1}(\overline{\sigma}_x < \infty)\prod_{\ell=1}^{\overline{\sigma}_x}\frac{p(\overline{X}_{\ell-1})}{p(\overline{X}_\ell)}\right] \\ &= \frac{p(y)}{p(x)}\mathbb{P}_y(\overline{\sigma}_x < \infty).\end{aligned}$$

Next, observe that since $\{X(t); t \geq 0\}$ is nonexplosive,

$$\sum_{y \in S}\nu^{(x)}(y) = q(x)\mathbb{E}_x[\tau_x]$$

which yields

$$\begin{aligned}q(x)\mathbb{E}_x[\tau_x] = \nu^{(x)}(x) + \sum_{y \neq x}\nu^{(x)}(y) &= 1 + \sum_{y \neq x}\frac{p(y)}{p(x)}\mathbb{P}_y(\tilde{\eta}_x < \infty) \\ &\leq \frac{p(x)}{p(x)} + \sum_{y \neq x}\frac{p(y)}{p(x)} = \frac{1}{p(x)} < \infty\end{aligned}$$

thus proving state x is positive recurrent, as $p(x) > 0$ by assumption. Finally, since x was an arbitrarily chosen state, we conclude that all states are positive recurrent, meaning $\{X(t); t \geq 0\}$ is positive recurrent. $\diamondsuit$
Readers should observe that if we do not assume $\{X(t); t \geq 0\}$ is regular, our proof only shows that

$$\mathbb{E}_x[\min(T_\infty, \tau_x)] < \infty$$

which does not imply positive recurrence.

4 Stochastic Networks

In this section, we explain how the ideas behind the random-product technique can be used to provide new insight into the structure of the stationary distribution associated with various types of Markovian queueing networks.

4.1 Whittle Networks and Jackson Networks

Consider an open Markovian queueing network consisting of n nodes, labeled $1, 2, \ldots, n$. We use the set $N := \{0, 1, \ldots, n\}$ to represent the set of all nodes in the network, where the extra 'node' 0 corresponds to the "outside" of the network (usage of this extra node will make our results both more elegant and transparent). For each $t \geq 0$, and each $i \in \{1, 2, \ldots, n\}$, we let $X_i(t)$ denote the number of units present at node i at time t, and we further define $X(t) := (X_1(t), \ldots, X_n(t))$ which keeps track of the number of units present at all nodes within the network. The underlying primitives of the network (unit arrivals, service times, routing scheme) are assumed to behave in a manner so that the stochastic process $\{X(t); t \geq 0\}$ is a CTMC, whose state space is $S = \mathbb{N}_0^n$, where $\mathbb{N}_0$ denotes the set of all nonnegative integers. We will let $e_1, e_2, \ldots, e_n \in S$ denote the n unit basis vectors, meaning that for each $i \in \{1, 2, \ldots, N\}$, e_i is a vector whose ith component is equal to one, and all other components of e_i are equal to zero. The vector e_0 will be used to denote the zero vector. Finally, we define the collection of transformations T_{ij}, $0 \leq i, j \neq N$, where for each $x \in S$,

$$T_{ij}x := x - e_i + e_j.$$

Usage of these transformations makes it convenient to describe various quantities that capture aspects of unit routing behavior within the network.

In order to properly state the transition rate matrix $\mathbf{Q}$ of $\{X(t); t \geq 0\}$, we first need to introduce some new notation. Let $\{\phi_i\}_{i \in N}$ denote a collection of functions, where each function in this collection maps from S to $[0, \infty)$: these will essentially be used to capture the service behavior of the queueing system found at each node. Next, we define the matrix $\mathbf{\Lambda} := [\lambda_{i,j}]_{i,j \in N}$ which will govern how each unit within the network makes transitions between nodes in the network. This matrix $\mathbf{\Lambda}$ it itself assumed to be a transition rate matrix, meaning (i) $\lambda_{i,j} \geq 0$ for each $i, j \in \{0, 1, 2, \ldots, n\}$ satisfying $j \neq k$, and (ii) for each $i \in N$,

$$\lambda_{i,i} = -\sum_{j \neq i} \lambda_{i,j}.$$

We assume, without loss of generality, that $\mathbf{\Lambda}$ is a irreducible rate matrix: under this assumption, it follows from earlier results in this paper that $\mathbf{\Lambda}$ has associated with it a unique stationary distribution, and each positive invariant measure $\boldsymbol{\alpha}$ of $\mathbf{\Lambda}$ is a positive scalar multiple of the stationary distribution.

Using both the collection of functions $\{\phi_i\}_{i \in N}$ and $\mathbf{\Lambda}$, we assume the transition rate matrix of $\{X(t) : t \geq 0\}$ is given by $\mathbf{Q} = [q(x, y)]_{x,y \in S}$ whose elements are as follows:

$$q(x, y) = \begin{cases} \lambda_{ij}\phi_i(x) & \text{if } y = T_{ij}x \in S \text{ for some } i \neq j \text{ in } N, \\ 0 & \text{otherwise.} \end{cases} \tag{74}$$

Readers should note that we are, for the most part, following the notation scheme found in Chapter 1 of [22]. In fact, Proposition 1.10 on page 13 of [22] shows that when Λ is assumed to be irreducible, it must also be the case that $\{X(t); t \geq 0\}$ is irreducible as well.

Observe that for $i, j \in \{1, 2, \ldots, n\}$, a transition from state x to state $T_{ij}x$ corresponds to the movement of a single unit from node i to node j. Furthermore, a transition from state x to state $T_{i0}x$ corresponds to a unit from node i departing the network, and a transition from state x to state $T_{0j}x$ corresponds to a unit arriving from outside of the network to node j. Each $\phi_i(x)$ term, for $1 \leq i \leq n$, $x \in S$, represents the service rate at node i whenever the system is in state x. For instance, if the queueing system present at node i is a multiserver queue with s_i servers that process work at unit rate, where services at that node are not further influenced by units present at other nodes in the network, we see that $\phi_i(x) = \max(x_i, s_i)\mu_i$, assuming of course that each unit arriving to node i brings with it an exponentially distributed amount of work with rate μ_i for processing. In many applications, for each $i \in \{1, 2, \ldots, n\}$, one can interpret $\lambda_{i,j}$ as the probability that when a single unit departs node i, it next moves to node j, with $\lambda_{i,0}$ denoting the probability that the departing unit leaves the network, but we stress that these terms can be more generally interpreted as routing rates rather than routing probabilities. Furthermore, for this open network, one can say that units arrive from outside of the network to node i according to a (state-dependent) Poisson process with rate $\lambda_{0i}\phi_0(x)$, meaning $(-\lambda_{0,0})\phi_0(x)$ represents the total (state-dependent) arrival rate of units to the network.

Definition 4.1 *We say that $\{X(t); t \geq 0\}$ is a **Whittle Network** if there exists a function $\Phi : S \to (0, \infty)$ such that for each state $x \in S$, and each $i, j \in N$, $i \neq j$, with $T_{ij}x \in S$,*

$$\Phi(x)\phi_i(x) = \Phi(T_{ij}x)\phi_j(T_{ij}x). \tag{75}$$

*We refer to Equation (75) as the Φ-**balancing property** of the Whittle network.*

It is not entirely clear at this point what consequences arise from requiring $\{X(t); t \geq 0\}$ to satisfy the Φ-balancing property, but in the proof of the next result, we will see precisely why this type of assumption helps to establish a closed-form representation for the stationary distribution $\mathbf{p}$ of $\{X(t); t \geq 0\}$.

Theorem 4.1 *Suppose $\boldsymbol{\alpha}$ is the unique positive solution of the system of balance equations $\boldsymbol{\alpha\Lambda} = \mathbf{0}$ that further satisfies $\alpha_0 = 1$. If Φ and $\boldsymbol{\alpha}$ are such that*

$$\sum_{x \in S} \Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell} < \infty$$

then $\{X(t); t \geq 0\}$ is both irreducible and positive recurrent, with a unique stationary distribution $\mathbf{p}$ that satisfies, for each $x \in S$,

$$p(x) = c\Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}.$$

where c satisfies

$$c^{-1} = \sum_{x \in S} \Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}.$$

Theorem 4.1 is essentially the same as Theorem 1.15 on page 15 of [22]: there the result is proven by showing that $\mathbf{p}$ satisfies a system of *partial balance equations*, but we will instead verify the result by making use of the random-product technique. One could argue that our usage of the random-product technique in the proof of this result is really an application of the generalized Kolmogorov criterion, yet technically the random-product technique can be used without knowing in advance whether or not $\{X(t); t \geq 0\}$ is positive recurrent.

Proof In order to apply the random-product technique, we first need to select another feasible CTMC $\{\tilde{X}(t); t \geq 0\}$ whose state space is given by S. It is natural to try to chose $\{\tilde{X}(t); t \geq 0\}$ so that $\{\tilde{X}(t); t \geq 0\}$ is a Markovian queueing network, so we assume first that the transition rate matrix $\tilde{\mathbf{Q}}$ of $\{\tilde{X}(t); t \geq 0\}$ is of the form

$$\tilde{q}(x, y) = \begin{cases} \tilde{\lambda}_{ij}\tilde{\phi}_i(x) & \text{if } y = T_{ij}x \in S \text{ for some } i \neq j \text{ in } N, \\ 0 & \text{otherwise} \end{cases}$$

where $\tilde{\mathbf{\Lambda}} := [\tilde{\lambda}_{i,j}]_{i,j \in N}$ represents the routing matrix associated with $\{\tilde{X}(t); t \geq 0\}$, and $\{\tilde{\phi}_i\}_{i \in N}$ represents its arrival and service rate functions. Given that each function $\tilde{\phi}_i$ is a function of only one state variable, it is worth first further assuming that $\tilde{\phi}_i = \phi_i$ for each $i \in \{0, 1, 2, \ldots, n\}$. Furthermore, given the success we have had in working with dual processes in Section 3, we will also assume that $\tilde{\mathbf{\Lambda}} = \overline{\mathbf{\Lambda}}$, where $\overline{\mathbf{\Lambda}}$ is the $\boldsymbol{\alpha}$-dual with respect to $\mathbf{\Lambda}$, meaning that for each $i, j \in N$ satisfying $i \neq j$,

$$\overline{\lambda}_{i,j} = \frac{\alpha_j \lambda_{j,i}}{\alpha_i}.$$

In other words, we further assume that $\tilde{\mathbf{Q}}$ is of the form

$$q(x, y) = \begin{cases} \overline{\lambda}_{ij}\phi_i(x) & \text{if } y = T_{ij}x \in S \text{ for some } i \neq j \text{ in } N, \\ 0 & \text{otherwise.} \end{cases} \tag{76}$$

Now that we have constructed $\{\tilde{X}(t) : t \geq 0\}$, we can make use the random-product technique. Letting state 0 be the reference state, we construct a measure $\boldsymbol{w}$ on S, where for each $x \neq 0$,

$$w(x) = \mathbb{E}_x \left[\mathbf{1}(\tilde{\sigma}_0 < \infty) \prod_{\ell=1}^{\tilde{\sigma}_0} \frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right].$$

Observe that each transition of $\{\tilde{X}(t); t \geq 0\}$ from state x to state $T_{ij}x$ contributes the following term to the random product:

$$\begin{aligned} \text{contribution of transition from } x \text{ to } T_{ij}x &= \frac{q(T_{ij}x, x)}{\tilde{q}(x, T_{ij}x)} \\ &= \frac{\lambda_{ji}\phi_j(T_{ij}x)}{\overline{\lambda}_{ij}\phi_i(x)} \\ &= \frac{\alpha_i\phi_j(T_{ij}x)}{\alpha_j\phi_i(x)} \\ &= \frac{\alpha_i\Phi(x)}{\alpha_j\Phi(T_{ij}x)}, \end{aligned} \tag{77}$$

where the last equality follows from the Φ-balancing property.

There are many useful things to observe from (77). First, each time a unit moves from node i to node j, that movement contributes to the random product the term α_i/α_j: in fact, when a new unit arrives to node i of the network, it contributes a term α_0/α_i, then it makes some transitions to other nodes in the network, then it finally leaves from some node j, which contributes the term α_j/α_0: thus, if a unit arrives to node i_0, then visits nodes $i_1, i_2, \ldots, i_n$, then leaves the network contributes

$$\frac{\alpha_0}{\alpha_{i_1}}\left[\prod_{\ell=1}^{n-1}\frac{\alpha_{i_\ell}}{\alpha_{i_{\ell+1}}}\right]\frac{\alpha_{i_n}}{\alpha_0} = 1.$$

From this observation, we can see that the only terms from $\boldsymbol{\alpha}$ that do not disappear from $w(x)$ are those that result from transitions made by units present in the system at time zero, meaning

$$w(x) = \left[\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell}\right]\mathbb{E}_x\left[\mathbf{1}(\tilde{\eta}_0 < \infty)\prod_{\ell=1}^{\tilde{\eta}_0}\frac{\Phi(\tilde{X}_{\ell-1})}{\Phi(\tilde{X}_\ell)}\right] = \left[\frac{\Phi(x)}{\Phi(0)}\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell}\right]\mathbb{P}_x(\tilde{\sigma}_0 < \infty).$$

If we further assume that

$$\sum_{x\neq 0}\frac{\Phi(x)}{\Phi(0)}\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell} < \infty$$

we get that $\{X(t); t \geq 0\}$ is positive recurrent, which implies it has a unique stationary distribution $\mathbf{p}$ on S. It is also true that the process $\{\tilde{X}(t); t \geq 0\}$ must be both irreducible and positive recurrent: irreducibility of $\{\tilde{X}(t); t \geq 0\}$ follows from $\{X(t); t \geq 0\}$ being irreducible, and the key to showing positive recurrence involves analyzing the measure $\tilde{\boldsymbol{w}}$ on S, where we again use state 0 as the reference state, and for each $x \neq 0$,

$$\tilde{w}(x) = \mathbb{E}_x\left[\mathbf{1}(\sigma_0 < \infty)\prod_{\ell=1}^{\sigma_0}\frac{\tilde{q}(X_\ell, X_{\ell-1})}{q(X_{\ell-1}, X_\ell)}\right].$$

It is not difficult to show that for each $x \neq 0$,

$$\tilde{w}(x) = \mathbb{P}_x(\sigma_0 < \infty)\frac{\Phi(x)}{\Phi(0)}\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell} = \frac{\Phi(x)}{\Phi(0)}\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell}$$

where the last equality follows from $\{X(t); t \geq 0\}$ being positive recurrent. Finally, since

$$\sum_{x\neq 0}\frac{\Phi(x)}{\Phi(0)}\prod_{\ell=1}^{n}\alpha_\ell^{x_\ell} < \infty$$

we conclude that

$$\sum_{x \in S} \tilde{w}(x) < \infty$$

which in turn implies $\{\tilde{X}(t); t \geq 0\}$ must also be positive recurrent. Hence, for each $x \neq 0$,

$$w(x) = \frac{\Phi(x)}{\Phi(0)} \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}.$$

Multiplying both sides of this equality by $p(0)$ further yields

$$p(x) = p(0) \frac{\Phi(x)}{\Phi(0)} \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}.$$

Finally, since

$$1 = \sum_{x \in S} p(x) = p(0) \left[\frac{\Phi(0)}{\Phi(0)} + \sum_{x \neq 0} \frac{\Phi(x)}{\Phi(0)} \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell} \right]$$

we conclude that

$$p(0) = \frac{\Phi(0)}{\sum_{x \in S} \Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}}$$

which implies

$$p(x) = c\Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell}$$

where

$$c = \left[\sum_{x \in S} \Phi(x) \prod_{\ell=1}^{n} \alpha_\ell^{x_\ell} \right]^{-1}.$$

This proves the result. $\diamondsuit$

We now turn our attention to Jackson networks. A Jackson network is a special type of Whittle network, where the service rate functions are node-dependent instead of system-dependent, meaning that for each $i \in \{1, 2, \ldots, n\}$, ϕ_i is only a function of x_i, not of the entire vector x, and the function ϕ_0 satisfies $\phi_0(x) = 1$ for each $x \in S$. As discussed in [22], one can easily show that for a Jackson network, the service rate function is Φ-balanced by the balancing function

$$\Phi(x) = \prod_{i \in N} \prod_{m=1}^{x_i} \phi_i(m)^{-1}. \tag{78}$$

This observation proves that a Jackson network is a special type of Whittle network, and as a consequence of Theorem 4.1, we get the following result.

Theorem 4.2 *The invariant measure of a Jackson network $\{X(t) : t \geq 0\}$ is given by*

$$w(x) = \prod_{i=1}^{n} \prod_{m=1}^{x_i} \frac{\alpha_i}{\phi_i(m)}, \quad x \in S := \mathbb{N}_0^n,$$

with $\alpha = (\alpha_i)_{i\in N}$ being a positive measure satisfying traffic equations. Furthermore, the process is positive recurrent if and only if the following condition holds for each $i = 1, \dots, n$,

$$c_i^{-1} = \sum_{x_i=0}^{\infty} \prod_{m=1}^{x_i} \frac{\alpha_i}{\phi_i(m)} < \infty.$$

Moreover, if this condition holds, the stationary distribution of the Jackson network is

$$\pi(x) = \prod_{i=1}^{n} \pi_i(x_i),$$

where

$$\pi_i(x_i) = c_i \prod_{m=1}^{x_i} \frac{\alpha_i}{\phi_i(m)}.$$

Proof One could simply say that this result follows immediately from Theorem 4.1, but if we apply our proof technique to the Jackson network, we gain additional insight into what makes the stationary distribution exhibit product-form. Observe that for this process, if we choose $\{\tilde{X}(t); t \geq 0\}$ in precisely the same manner as we did for the Whittle network, we see that for each state x, and each $i, j \in \{1, 2, \dots, n\}$,

$$\begin{aligned}\text{contribution of a transition from } x \text{ to } T_{ij}x &= \frac{q(T_{ij}x, x)}{\tilde{q}(x, T_{ij}x)} \\ &= \frac{\lambda_{ji}\phi_j(x_j+1)}{\overline{\lambda}_{ij}\phi_i(x_i)} \\ &= \frac{\alpha_i\phi_j(x_j+1)}{\alpha_j\phi_i(x_i)} \\ &= \frac{\alpha_i}{\phi_i(x_i)}\frac{\phi_j(x_j+1)}{\alpha_j}.\end{aligned}$$

Now, each time a unit leaves the queue at node i, it contributes to the product the term $\alpha_i/\phi_i(x_i)$ (where x_i includes the departing unit in its count) and each time a unit enters the queue at node i, it contributes to the product the term $\phi_j(x_j+1)/\alpha_j$. Moreover, for each state $i \in \{1, 2, \dots, n\}$,

$$\begin{aligned}\text{contribution of a transition from } x \text{ to } T_{i0}x &= \frac{q(T_{i0}x, x)}{\tilde{q}(x, T_{i0}x)} \\ &= \frac{\lambda_{0i}}{\overline{\lambda}_{i0}\phi_i(x_i)} \\ &= \frac{\alpha_i}{\alpha_0\phi_i(x_i)} \\ &= \frac{\alpha_i}{\phi_i(x_i)}\end{aligned}$$

and for each state $j \in \{1, 2, \dots, n\}$,

$$\begin{aligned}\text{contribution of a transition from } x \text{ to } T_{0j}x &= \frac{q(T_{0j}x, x)}{\tilde{q}(x, T_{0j}x)} \\ &= \frac{\lambda_{j0}\phi_j(x_j+1)}{\overline{\lambda}_{0j}} \\ &= \frac{\alpha_0\phi_j(x_j+1)}{\alpha_j} \\ &= \frac{\phi_j(x_j+1)}{\alpha_j}.\end{aligned}$$

From this observation, we see that from the telescoping property, the only nontrivial terms that appear in the random-product when it reaches state 0 are those that stem from the units present in the system at time zero: in other words, for each state $x \neq 0$,

$$w(x) = \mathbb{E}_x \left[\mathbf{1}(\tilde{\sigma}_0 < \infty) \prod_{\ell=1}^{\tilde{\sigma}_0} \frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right] = \left[\prod_{\ell=1}^{n} \left[\prod_{m=1}^{x_\ell} \frac{\alpha_\ell}{\phi_\ell(m)} \right] \right] \mathbb{P}_x(\tilde{\eta}_0 < \infty) = \prod_{\ell=1}^{n} \left[\prod_{m=1}^{x_\ell} \frac{\alpha_\ell}{\phi_\ell(m)} \right]$$

where one can show that $\mathbb{P}_x(\tilde{\eta}_0 < \infty) = 1$ through reasoning analogous to that used to derive the stationary distribution of a Whittle network. Hence, for each state $x \neq 0$,

$$p(x) = p(0) \prod_{\ell=1}^{n} \left[\prod_{m=1}^{x_\ell} \frac{\alpha_\ell}{\phi_\ell(m)} \right].$$

From here the normalizing constant $p(0)$ can be found in a straightforward manner. $\diamondsuit$

4.2 Stochastic Networks in Random Environments

We close this paper by studying the stationary distributions of two different types of CTMCs, where the CTMC is influenced in some manner by another underlying CTMC referred to as the environment process.

4.2.1 Markovian Queueing Networks in Random Environments

In the work of Zhu [24], the author considers an open Jackson network consisting of n nodes, where both the routing matrix and the ϕ_i functions are further influenced by a random environment. More particularly, consider an underlying irreducible, positive recurrent CTMC $\{E(t); t \geq 0\}$ with a countable state space S_2, generator $\mathbf{R} := [r(u,v)]_{u,v \in S_2}$ and stationary distribution $\boldsymbol{\pi}$: we refer to $\{E(t); t \geq 0\}$ as the underlying environment. While the environment process is state $u \in S_2$, the network behaves as a Jackson network with irreducible routing matrix $\mathbf{\Lambda}(u)$ and functions $\{\phi_i^{(u)}\}_{0 \leq i \leq n}$. Again, letting $\{X(t) : t \geq 0\}$ be a vector-valued process that keeps track of the unit populations at all n nodes, with state space $S_1 := \mathbb{N}_0^N$, while setting $Y(t) := (X(t), E(t))$, we have that the process $\{Y(t); t \geq 0\}$ is a CTMC with state space $S = S_1 \times S_2$ and generator $\mathbf{Q} = [q((x,u),(y,v))]_{(x,u),(y,v) \in S}$, where

$$q((x,u),(y,v)) = \begin{cases} \lambda_{i,j}(u)\phi_i^{(u)}(x) & \text{if } y = T_{ij}x \in S_1 \text{ for some distinct } i,j \in N, \;\; u = v \in S_2, \\ r(u,v) & \text{if } y = x \in S_1, u \neq v \in S_2, \\ 0 & \text{otherwise.} \end{cases} \tag{79}$$

For each state $u \in S_2$, we can associate with $\mathbf{\Lambda}(u)$ the unique invariant measure $\boldsymbol{\alpha}(u) := [\alpha_i(u)]_{i \in N}$ that satisfies $\alpha_0(u) = 1$. Our next result shows that under certain conditions, the stationary distribution of $\{Y(t), t \geq 0\}$ exists, and has a product-form structure.

Theorem 4.3 ([24]) *Suppose that for each integer $n \geq 1$, and each $i \in N$, the product*

$$\frac{\alpha_i(u)}{\phi_i^{(u)}(n)}$$

is independent of u, and furthermore, for each $i \in \{1, 2, \ldots, n\}$,

$$\sum_{y=0}^{\infty} \prod_{m=1}^{y} \frac{\alpha_i(u)}{\phi_i^{(u)}(m)} < \infty.$$

Under these conditions, the CTMC $\{Y(t); t \geq 0\}$ is both irreducible and positive recurrent, and its unique stationary distribution is given by

$$p(x,u) = c\pi(u) \prod_{i=1}^{n} \prod_{m=1}^{x_i} \frac{\alpha_i(u)}{\phi_i^{(u)}(m)}$$

where

$$c = \left(\sum_{u \in S_2} \sum_{x \in S_1} \pi(u) \prod_{i=1}^{n} \prod_{m=1}^{x_i} \frac{\alpha_i(u)}{\phi_i^{(u)}(m)} \right)^{-1}.$$

Proof Just as we did for both Whittle networks and Jackson networks, we define the process $\left\{\tilde{Y}(t); t \geq 0\right\}$ having generator $\tilde{\mathbf{Q}} = \left(\tilde{q}((x,u),(y,v))\right)_{x,y \in S_1, u,v \in S_2}$, which satisfies

$$\tilde{q}((x,u),(y,v)) = \begin{cases} \overline{\lambda}_{ij}(u)\phi_i^{(u)}(x) & \text{if } y = T_{ij}x \in S_1 \text{ for some distinct } i,j \in N, \;\; u = v \in S_2, \\ \overline{r}(u,v) & \text{if } y = x \in S_1, u \neq v \in S_2, \\ 0 & \text{otherwise,} \end{cases} \tag{80}$$

with $\overline{\mathbf{\Lambda}}(u)$ is the generator of the $\boldsymbol{\alpha}(u)$-dual of the CTMC having generator $\mathbf{\Lambda}$, and $\overline{\boldsymbol{R}}$ is the $\boldsymbol{\pi}$-dual of $\boldsymbol{R}$. Using this particular choice for $\{\tilde{Y}(t); t \geq 0\}$, we fix a state $(0, u_0) \in S$ as our reference state, and we focus on calculating, for each state $(x,u) \in S$,

$$w((x,u)) = \mathbb{E}_{(x,u)} \left[\mathbf{1}(\tilde{\eta}_{(0,u_0)} < \infty) \prod_{\ell=1}^{\tilde{\eta}_{(0,u_0)}} \frac{q(\tilde{Y}_\ell, \tilde{Y}_{\ell-1})}{\tilde{q}(\tilde{Y}_{\ell-1}, \tilde{Y}_\ell)} \right].$$

Observe that under the tilde process $\{\tilde{Y}(t); t \geq 0\}$, a transition from state (x,u) to state $(T_{ij}x, u)$ contributes to the random product the term

$$\begin{aligned} \text{contribution of unit transition} &= \frac{q((T_{ij}x,u),(x,u))}{\tilde{q}((x,u),(T_{ij}x,u))} \\ &= \frac{\lambda_{ji}(u)\phi_j^{(u)}(x_j+1)}{\tilde{\lambda}_{ij}(u)\phi_i^{(u)}(x_i)} \\ &= \frac{\alpha_i(u)\phi_j^{(u)}(x_j+1)}{\alpha_j(u)\phi_i^{(u)}(x_i)} \\ &= \frac{\alpha_i(u)}{\phi_i^{(u)}(x_i)} \frac{\phi_j^{(u)}(x_j+1)}{\alpha_j(u)} \\ &= \frac{\alpha_i(u_0)}{\phi_i^{(u_0)}(x_i)} \frac{\phi_j^{(u_0)}(x_j+1)}{\alpha_j(u_0)} \end{aligned} \tag{81}$$

where it is important to notice that, regardless of the state of the environment, each arrival (from either inside or outside of the network) to node i contributes to the product the term $\phi_i^{(u_0)}(x_i + 1)/\alpha_i(u_0)$, and each departure from node i contributes to the product the term $\alpha_i(u_0)/\phi_i^{(u_0)}(x_i)$.

Similarly, a transition from state (x, u) to (x, v) contributes the term

$$\begin{aligned}\text{contribution of environment transition} &= \frac{q((x,v),(x,u))}{\tilde{q}((x,u),(x,v))} \\ &= \frac{r(v,u)}{\tilde{r}(u,v)} \\ &= \frac{\pi(u)}{\pi(v)}.\end{aligned} \tag{82}$$

Reasoning just as did for the ordinary Jackson network, we conclude that

$$w(x,u) = \frac{\pi(u)}{\pi(u_0)} \prod_{i=1}^{N} \prod_{n=1}^{x_i} \frac{\alpha_i(u)}{\phi_i^{(u)}(n)}$$

from which we can derive the stationary distribution of $\{Y(t); t \geq 0\}$ through an obvious normalization method. $\diamondsuit$

4.2.2 General Markov Processes in Random Environments

Our final example stems from the recent work of Belopolskaya and Suhov [2], which further builds on the results of Gannon et al. [12]. In [2], the authors consider a CTMC $\{X(t); t \geq 0\}$ whose state space S is of the form $S = S_1 \times S_2$, where both S_1 and S_2 are countable sets. The transition rate matrix $\mathbf{Q}$ of $\{X(t); t \geq 0\}$ satisfies

$$q((x,u),(y,v)) = \begin{cases} q^{(u)}(x,y) & \text{if } y \neq x, \;\; u = v, \\ \frac{1}{m^{(u)}(x)} \tau^{(x)}(u,v) & \text{if } y = x, u \neq v, \\ 0 & \text{otherwise,} \end{cases} \tag{83}$$

where for each fixed $u \in S_2$, $\boldsymbol{Q}^{(u)} := [q^{(u)}(x,y)]_{x,y \in S_1}$ represents the transition rate matrix associated with some other CTMC $\{X^{(u)}(t); t \geq 0\}$ having state space S_1: we assume each CTMC $\{X^{(u)}(t); t \geq 0\}$ is irreducible with positive invariant measure $\boldsymbol{m}^{(u)}$. Moreover, for each $x \in S_1$, $\boldsymbol{\tau}^{(x)} := [\tau^{(x)}(u,v)]_{u,v \in S_2}$ represents the transition rate matrix associated with another CTMC, which is also assumed to be irreducible with invariant measure $\boldsymbol{\nu}^{(x)}$. In some ways the structure of $\{X(t); t \geq 0\}$ is similar in flavor to the structure of the Jackson network in a random environment considered in [24] (i.e. the model from the previous section, which is described using notation from [22]), in that while the environment process in state u, the chain evolves according to a CTMC with generator $\boldsymbol{Q}^{(u)}$. On the other hand, the environment process associated with $\{X(t); t \geq 0\}$ is also influenced by the other component of the process, yet the structure of $\mathbf{Q}$ yields a very elegant stationary distribution, when it exists.

Our next result provides an expression for the stationary distribution of $\{X(t); t \geq 0\}$, when it exists. This is a special case of Theorem 4.1 from [2], which goes beyond the countable state case.

Theorem 4.4 *Suppose that the invariant measures $\boldsymbol{\nu}^{(x)}$ do not vary with respect to changes in x, meaning there is some positive measure $\boldsymbol{\nu}$ such that $\boldsymbol{\nu}^{(x)} = \boldsymbol{\nu}$ for each $x \in S_1$. If it is also true that*

$$\sum_{y \in S_1} \sum_{v \in S_2} m^{(v)}(y) \nu(v) < \infty$$

then $\{X(t); t \geq 0\}$ is irreducible and positive recurrent, and its stationary distribution is given by

$$p(x,u) = c\, m^{(u)}(x) \nu(u),$$

where $c = \left(\sum_{y \in S_1} \sum_{v \in S_2} m^{(v)}(y) \nu(v)\right)^{-1}$ is a normalizing constant.

Proof Using our previous examples as a guide, we construct the process $\{\tilde{X}(t); t \geq 0\}$ having state space S and transition rate matrix $\tilde{Q} = [\tilde{q}\left((x,u),(y,v)\right)]_{x,y\in S_1,u,v\in S_2}$ defined as

$$\tilde{q}\big((x,u),(y,v)\big) = \begin{cases} \overline{q}^{(u)}(x,y) & \text{if } y \neq x, \quad u = v, \\ \frac{1}{m^{(u)}(x)}\overline{\tau}^{(x)}(u,v) & \text{if } y = x, u \neq v, \\ 0 & \text{otherwise,} \end{cases} \tag{84}$$

where $\overline{\boldsymbol{Q}}^{(u)}$ represents the $\boldsymbol{m}^{(u)}$-dual of $\{X^{(u)}(t); t \geq 0\}$, and $\overline{\boldsymbol{\tau}}^{(x)}$ represents the $\boldsymbol{\nu}^{(x)}$-dual, or rather the $\boldsymbol{\nu}$-dual, of the CTMC having transition rate matrix $\boldsymbol{\tau}^{(x)}$.

Observe that under the tilde process $\{\tilde{X}(t); t \geq 0\}$, a transition from state (x,u) to (y,u) contributes to the random product the term

$$\begin{aligned} \text{contribution of transition} &= \frac{q((y,u),(x,u))}{\tilde{q}((x,u),(y,u))} \\ &= \frac{q^{(u)}(y,x)}{\tilde{q}^{(u)}(x,y)} \\ &= \frac{m^{(u)}(x)}{m^{(u)}(y)}. \end{aligned} \tag{85}$$

On the other hand, each transition from state (x,u) to (x,v) contributes the term

$$\begin{aligned} \text{contribution of transition} &= \frac{q((x,v),(x,u))}{\tilde{q}((x,u),(x,v))} \\ &= \frac{\frac{1}{m^{(v)}(x)}\tau^{(x)}(v,u)}{\frac{1}{m^{(u)}(x)}\tilde{\tau}^{(x)}(u,v)} \\ &= \frac{m^{(u)}(x)}{m^{(v)}(x)}\frac{\nu(u)}{\nu(v)}. \end{aligned} \tag{86}$$

Next, fix a reference point $(x_0,u_0) \in S_1 \times S_2$: one can at this point easily show that from both (85) and (86), combined with a line of reasoning similar to that we used in the derivation of the stationary distribution of a Whittle network, we have

$$w(x,u) = \mathbb{E}_{(x,u)}\left[\mathbf{1}(\tilde{\sigma}_{(x_0,u_0)} < \infty)\prod_{\ell=1}^{\tilde{\sigma}_{(x_0,u_0)}}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)}\right] = \frac{m^{(u)}(x)\nu(u)}{m^{(u_0)}(x_0)\nu(u_0)}.$$

From this point, the rest of the argument follows in a straightforward manner. $\diamondsuit$

References

[1] Asmussen, S. (2003). *Applied Probability and Queues*. Springer-Verlag, New York.

[2] Belopolskaya, Y., and Suhov, Y. (2015). Models of Markov processes with a random transition mechanism. Preprint available on arXiv at https://arxiv.org/abs/1508.05598

[3] Brémaud, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer-Verlag, New York.

[4] Buckingham, P., and Fralix, B. (2015). Some new insights into Kolmogorov's Criterion, with applications to hysteretic queues. *Markov Processes and Related Fields* **21**, 339-368.

[5] Chen, H., and Yao, D.D. (2001). *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*. Springer.

[6] Çinlar, E. (1975). *Introduction to Stochastic Processes*. Academic Press, New York.

[7] Durrett, R. (2016). *Essentials of Stochastic Processes*. Springer International Publishing, Switzerland.

[8] Durrett, R. (2019). *Probability: Theory and Examples*, 5th Edition. Cambridge University Press.

[9] Fralix, B. (2015). When are two Markov chains similar? *Statistics and Probability Letters* **107**, 199-203.

[10] Fralix, B. (2018). A new look at a smart polling model. *Mathematical Methods of Operations Research* **88**, 339-367.

[11] Fralix, B. (2019). On the time-dependent behavior of a Markovian reentrant-line model. *Probability in the Engineering and Informational Sciences* **33**, 367-386.

[12] Gannon, M., Pechersky, E., Suhov, Y., and Yambartsev, A. (2016). Random walks in a dynamic network environment. *Journal of Applied Probability* **53**, 448-462.

[13] Göbel, J., Keeler, H.P., Krezesinski, A., and Taylor, P.G. (2016). Bitcoin blockchain dynamics: the selfish-mine strategy in the presence of propagation delay. *Performance Evaluation* **104**, 23-41.

[14] Javier, K. and Fralix, B. (2019). A further study of some Markovian Bitcoin models from Göbel et al. *Stochastic Models*, to appear.

[15] Joyner, J. and Fralix, B. (2016). A new look at Markov processes of $G/M/1$-type. *Stochastic Models* **32**, 253-274.

[16] Kelly, F. (1979). *Reversibility and Stochastic Networks*. John Wiley and Sons.

[17] Kelly, F. (1983). Invariant measures and the Q-matrix. *Probability, Statistics and Analysis*, J.F.C. Kingman and G.E.H. Reuter, Editors, 143-160. Cambridge University Press.

[18] Last, G., and Brandt, A. (1995). *Marked Point Processes on the Real Line: The Dynamic Approach*. Springer.

[19] Liu, X. and Fralix, B. (2019). On lattice-path counting and the random-product representation, with applications to the E_r/M/1 queue and the M/E_r/1 queue. *Methodology and Computing in Applied Probability* **21**, 1119-1149.

[20] Pollett, P.K. (2000). Similar Markov chains. *Advances in Applied Probability*

[21] Selen, J. and Fralix, B. (2017). Time-dependent analysis of a $M/M/c$ priority system with two priority classes. *Queueing Systems* **87**, 379-415.

[22] Serfozo, R. (1999). *Introduction to Stochastic Networks*. Springer, New York.

[23] Serfozo, R. (2009). *Basics of Applied Stochastic Processes*. Springer, Berlin.

[24] Zhu, Y. (1994). Markovian queueing networks in a random environment. *Operations Research Letters* **15**, 11-17.