Chapter 11

Markov Chains

In the chapter on limit theorems, we saw sequences of random variables indexed by the natural numbers. The underlying experiments generated sequences of independent, identically distributed random variables, from which we constructed derived sequences such as the partial sum or the incremental average of the random variables.

Collections of random variables $\{X_t\}$ indexed by the natural numbers are known as discrete-time stochastic processes, or discrete-time random processes. The index is used to represent time. Such models are often used to represent random time signals that arise in dynamical systems, and have many interesting applications in engineering.

In this chapter, we focus on a class of discrete time random processes known as Markov chains. Markov chains are a special class of discrete time random processes because of two properties. First, the range of the individual variables X_t is discrete, which will will allow us to develop a rich connection between the probability models and concepts from graph theory. Second, the joint probability distribution functions of Markov chains will satisfy the Markov property, which we discuss later in this chapter. Markov chains can also be defined as random processes that are indexed in continuous time, but those extensions are outside the scope of this course.

Markov chains were introduced by Andrey Markov in 1906 to study extensions of the Law of Large Numbers and the Central Limit Theorem to sequences where the random variables were not independent and identically distributed. Such models form the basis for many interesting applications such as speech recognition, communications networks analysis and stochastic automata. Markov chains provide the foundation for many of today's leading technologies. Google's page rank algorithm was based on a Markov chain model of how websites are visited. Viterbi decoding, named after one of Qualcomm's founder, is based on Hidden Markov Model techniques, and is used extensively in modern communications. Markov chain models play fundamental roles in speech and natural language recognition. Markov models are used extensively in mathematical finance to analyze expected returns of different investment mixtures. They also provide the foundation for the analysis and design of network systems for handling random traffic demands.

In the remainder of this chapter, we discuss the foundations of discrete-time Markov Chains and explore their properties. First, we will introduce discrete-time, discrete-space Markov processes, and define the Markov property that characterizes such processes. Following this, we develop tools for computing probabilities in Markov chains. We develop approaches for characterizing how the marginal probability of the Markov chain evolves with the time index, and explore the limiting behavior of such systems. We also introduce tools for analysis of the transient behavior of Markov chains.

11.1 Definition of Markov Chains

Let $\mathcal{R}_X$ be a finite, or countably infinite set of possible values, which we call the **state space**. This set is a subset $\mathcal{R}_X \subset \Re$. Define a probability space $(\Omega, \mathcal{E}, \mathbb{P})$ that generates a countably infinite sequence of random variables $X_0, X_1, X_2, \ldots$, each of which takes values in the state space $\mathcal{R}_X$. An outcome $\omega \in \Omega$ generates a sequence of numbers $X_0(\omega), X_1(\omega), \ldots$ with values in $\mathcal{R}_X$. For each outcome, we refer to this sequence as a trajectory of the Markov chain.

Given a finite subsets of these random variables $X_{t_1}, \ldots, X_{t_n}$, where $t_1, \ldots, t_n \in \{0, 1, 2, \ldots\}$, we can

compute joint probability mass functions of the form $P_{X_{t_1},\ldots,X_{t_m}}(x_{t_1},\ldots,x_{t_m})$. These joint probability mass functions (PMF) can be used to generate conditional probability mass functions as well as marginal probability mass functions. We refer to the indices t as times, so we think of $X_t(\omega), t = 0, 1, \ldots$ as a trajectory over time.

Without loss of generality, assume the indices $t_1 < t_2 < \ldots < t_m$ are ordered linearly in time. Using conditional probabilities and the product rule for probability mass functions, we can write the joint PMF of the random variables with those indices as

$$P_{X_{t_1},\ldots,X_{t_m}}(x_{t_1},\ldots,x_{t_m}) = P_{X_{t_m}|X_{t_1},\ldots,X_{t_{m-1}}}(x_{t_m}|x_{t_1},\ldots,x_{t_{m-1}})$$
$$P_{X_{t_{m-1}}|X_{t_1},\ldots,X_{t_{m-2}}}(x_{t_{m-1}}|x_{t_1},\ldots,x_{t_{m-2}})\cdots P_{X_{t_2}|X_{t_1}}(x_{t_2}|x_{t_1})P_{X_{t_1}}(x_{t_1})$$

We say that the sequence of random variables $X_0, X_1, X_2 \ldots$, satisfies the **Markov Property** if and only if, for any set of times $t > t_m > \ldots > t_1$, we have

$$P_{X_t|X_{t_1},\ldots,X_{t_m}}(x_t|x_{t_1},\ldots,x_{t_m}) = P_{X_t|X_{t_m}}(x_t|x_{t_m}).$$

That is, the conditional probability mass function of the random variable at time t, X_t, given values of random variables at different previous times $t_1, t_2, \ldots, t_m$, depends only on the value of the most recent random variable in its past. This simplifies how we write the joint probability mass function, as

$$P_{X_{t_1},\ldots,X_{t_m}}(x_{t_1},\ldots,x_{t_m}) = P_{X_{t_m}|X_{t_{m-1}}}(x_{t_m}|x_{t_{m-1}})P_{X_{t_{m-1}}|X_{t_{m-2}}}(x_{t_{m-1}}|x_{t_{m-2}})\cdots P_{X_{t_2}|X_{t_1}}(x_{t_2}|x_{t_1})P_{X_{t_1}}(x_{t_1})$$

Thus, we can specify the joint probability mass function (PMF) of a collection of random variables in terms of a product of pairwise conditional PMFs times the marginal PMF of the random variable with the earliest time index. This economical description is very useful in obtaining an economical probabilistic description of the Markov chain.

Of particular interest is the one-step conditional probability $P_{X_{t+1}|X_t}(x_{t+1}|xt)$. In general, this conditional probability depends on time. Assume that the state space is given as $\mathcal{R}_X = \{a_1, a_2, \ldots, a_n, \ldots\}$. Then, $P_{X_{t+1}|X_t}(x_{t+1} = a_k|x_t = a_j)$ depends on a_k, a_j, and t. When this conditional probability does not depend on t, we say the Markov chain is **homogeneous** or **time-invariant**. Homogeneous Markov chains have the nice property that the conditional probability mass function $P_{X_{t+1}|X_t}(x_{t+1} = a_k|x_t = a_j)$ is the same for all $t = 0, 1, 2, \ldots$. Hence, the full probability description of the Markov chain can be obtained from the marginal PMF $P_{X_0}(x_0)$ and the one-step conditional probability $P_{X_{t+1}|X_t}(x_{t+1} = a_k|x_t = a_j)$. As shorthand notation, we define the **transition probability kernel** of the Markov chain as a matrix $\mathbf{P}$ with elements defined as:

$$\mathbf{P}_{jk} = P_{X_{t+1}|X_t}(x_{t+1} = a_k|x_t = a_j),\ j, k \in \{1, 2, \ldots\}.$$

Thus, $\mathbf{P}_{jk}$ is the probability that, if the random variable X_t has value a_j, then the random variable X_{t+1} will take value a_k. The transition probability kernel has the following properties:

- $1 \geq \mathbf{P}_{jk} \geq 0$ for all $j, k \in \{1, 2, \ldots\}$. This follows because it was defined as a conditional probability, which is a probability.
- $\sum_k \mathbf{P}_{jk} = 1$. This property is the normalization property for conditional PMFs.

We refer to the random variable X_t as the **state** at time t. The Markov chain provides a probabilistic description of how the state X_t evolves over time.

Example 11.1
Consider the following Markov chain, where the state space is $\mathcal{R}_X = \{1, 2, 3, 4, 5\}$. We assume initially that $X_0 = 3$; that is, uniform; that is, $P_{X_0}(3) = 1, P_{X_0}(x) = 0$ if $x \neq 3$. Thus, we have defined the marginal PDF at time 0. We now describe the transition probability kernel, as follows: if $j \neq 1, j \neq 5$, then

$$P_{X_{t+1}|X_t}(x_{t+1} = k|x_t = j) = \begin{cases} 0.5 & k = j + 1 \\ 0.5 & k = j - 1 \\ 0 & \text{elsewhere.} \end{cases}$$

For $j = 1$, the transition probability kernel is

$$P_{X_{t+1}|X_t}(x_{t+1} = k|x_t = 1) = \begin{cases} 0.5 & k = 2 \\ 0.5 & k = 1 \\ 0 & \text{elsewhere.} \end{cases}$$

For $j = 5$, the transition probability kernel is

$$P_{X_{t+1}|X_t}(x_{t+1} = k|x_t = 5) = \begin{cases} 0.5 & k = 5 \\ 0.5 & k = 4 \\ 0 & \text{elsewhere.} \end{cases}$$

Note that we can represent this transition probability kernel as a matrix $\mathbf{P}$, where

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.5 & 0 & 0 & 0 \\ 0.5 & 0 & 0.5 & 0 & 0 \\ 0 & 0.5 & 0 & 0.5 & 0 \\ 0 & 0 & 0.5 & 0 & 0.5 \\ 0 & 0 & 0 & 0.5 & 0.5 \end{bmatrix}$$

We now have a complete description of the probabilistic structure of the Markov chain. We can answer questions such as: What is the probability that $X_0 = 3, X_1 = 2, X_2 = 1$? Note that this will be $P_{X_0}(3)\mathbf{P}_{32}\mathbf{P}_{21} = 0.25$. Another question might be what is the probability that $X_3 = 3$? Although we don't have an easy way of computing this yet, we see that there are two ways that $X_3 = 3$, which is with $X_2 = 2$ and $X_2 = 4$. Each of those two paths will have probability 0.25, so the probability that $X_3 = 3$ will be 0.5.

More rigorously, we would compute the joint probability of $X_1 = 3, X_2 = k, X_3 = 3$ as $\mathbf{P}_{3k}\mathbf{P}_{k3}$. To get the probability that $X_3 = 3$, we would sum over k this joint probability, thereby marginalizing the intermediate random variable $X_2 = k$. It so happens that this product is nonzero only for $k = 2$ and $k = 4$, so the sum is again 0.5.

In the special case that the state space $\mathcal{R}_X$ is finite, the set of possible states is $\{a_1, a_2, \ldots, a_K\}$, and the transition probability kernel $\mathbf{P}_{jk}$ can be represented as a $K \times K$ matrix $\mathbf{P}$ with $(j,k)-$element $\mathbf{P}_{jk}$. In this case we denote $\mathbf{P}$ as the **state transition matrix** or the **transition probability matrix** of the Markov chain. We study the special case of homogeneous, finite state Markov chains next.

11.2 Finite State Markov Chains

11.2.1 Graphical representation of the Markov chain

Consider a finite state Markov chain, with state space $\mathcal{R}_X = \{a_1, a_2, \ldots, a_K\}$. To simplify notation, we assume $\mathcal{R}_X = \{1, 2, \ldots, K\}$. For a homogeneous, finite state Markov chain, the transition probability kernel is represented by a state transition matrix $\mathbf{P}$, with properties

- $\mathbf{P}_{jk} \in [0, 1],\ j, k \in \{1, \ldots, K\}$.
- $\sum_{k=1}^{n} \mathbf{P}_{jk} = 1$ for $j = 1, \ldots, K$.

That is, all of the elements of $\mathbf{P}$ are nonnegative numbers less than or equal to 1, and the sum of every row equals one. Matrices that satisfy these two properties are known as **stochastic matrices.** Later in this section, we will describe some useful properties of stochastic matrices that help us understand the behavior of Markov chains.

The state transition matrix $\mathbf{P}$ is often sparse, containing many zeros. In Example 11.1, over half the matrix was composed of zeros. We can represent the contents of the matrix $\mathbf{P}$ in graphical form, where nodes indicate possible values of the state, and directed arcs between nodes represent transition probabilities. Thus,

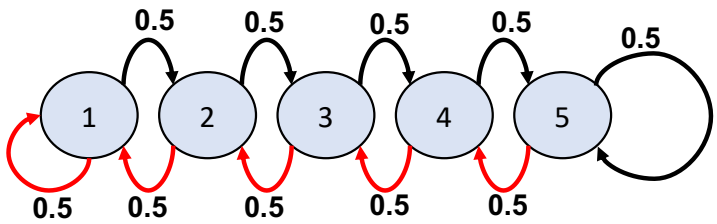

Figure 11.1: Graph of Markov chain state transition matrix for Example 11.1.

the graph contains K nodes (the cardinality of $\mathcal{R}_X$) and a number of directed arcs equal to the number of positive elements in $\mathbf{P}$. Figure 11.1 shows the graph that represents the state transition matrix $\mathbf{P}$ in Example 11.1. Note the directed arcs, with weights that correspond to the non-zero entries of $\mathbf{P}$. The condition that the rows of $\mathbf{P}$ must each add up to 1 implies that the sum of the probabilities of the arcs that leave each node must equal 1. This includes self-loop arcs where the transition is from a particular state to itself.

Example 11.2
Consider a four state Markov chain, with state transition matrix shows the graph for a four state Markov matrix

$$\mathbf{P} = \begin{bmatrix} P_{11} & P_{12} & 0 & 0 \\ 0 & 0 & P_{23} & P_{24} \\ P_{31} & 0 & 0 & 0 \\ 0 & 0 & P_{43} & P_{44} \end{bmatrix}$$

What is the graph of the Markov chain?

The graph is shown in the figure below. The graph has 7 directed arcs, corresponding to the 7 non-zero elements of $\mathbf{P}$.

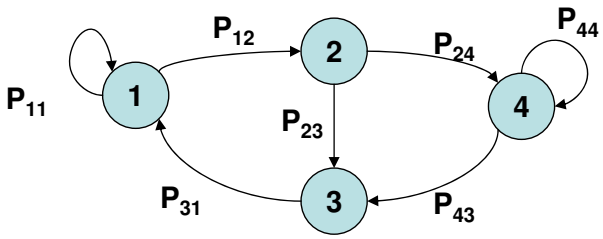

Figure 11.2: Graph of Markov Chain transition probabilities.

Example 11.3
One of the simplest Markov chain models has two states, $\mathcal{R}_X = \{1, 2\}$, corresponding to an on-off system. This model is often used for failure-repair processes. When the model is in state 1, the "on" state, there is a probability of failure p at each time. Eventually, a failure happens, and the state of the system transitions to state 2, the "off" state. In this state, there is a probability of repair q at each time. Eventually, the state transitions back to the "on" state 1. The state transition diagram is shown in Figure 11.3.

The graph representation of the state transition matrix helps us understand how the Markov chain behaves as a function of time. One view of the Markov chain is that it is a collection $\{X_t, t = 0, 1, \ldots\}$ of random variables with joint probability mass functions that satisfy the Markov property. A different view is to consider the sequence of values $\{X_0(s), X_1(s), \ldots\}$ that would occur from a single realization s of the experiment that generated the chain. We refer to such a sequence as a **trajectory** of the Markov chain. A trajectory is a time sequence of state values, and can be viewed as a trajectory on the graph, where transitions between states that are adjacent in time can only happen if there is a directed arc from the previous state to the next state. With this perspective, the Markov chain generates a probability distribution

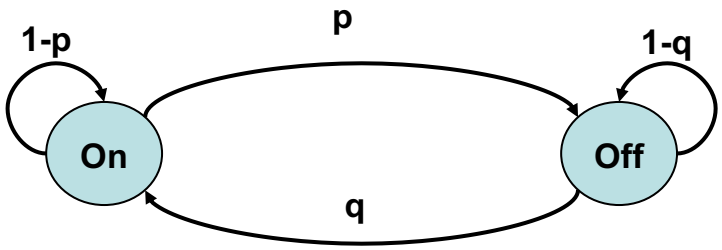

Figure 11.3: Graph of Markov Chain transition probabilities.

over possible state trajectories on the Markov chain graph. The Markov property establishes that, given knowledge that the chain is in state $X_t = k$ at time t, the probability distribution on the future trajectory of the state depends only on $X_t = k$, and not on any values $X_s, s < t$. Thus, $X_t = k$ has all the information needed to predict the future state values $X_\tau, \tau > t$.

Example 11.3 illustrates an important property of the state trajectories of Markov chains. We know the system remains in the same state for a random number of time steps before transitioning to another state. A possible state trajectory for the first 30 steps is 000000011000000000011110000011, where we see the trajectory start at state 0, stay there for 7 times before transitioning to state 1 in time 8. The next time the system visits state 0, it transitions to state 1 in 11 times. Because of the Markov property, the amount of time it takes to transition out of state 0 has the same distribution for every visit in the trajectory. For each state k, define the random variable $H_k(\omega)$ as follows:

$$H_k(\omega) = \min_{t>0}\{t : X_0(\omega) = k, X_t(\omega) \neq k\}.$$

Note we have included the explicit dependence on the realization of the trajectory ω. $H_k(\omega)$ is the first exit time that the Markov chain trajectory would leave state k, given that it started at time 0 in state k. .

H_k is a discrete random variable, with values in $\{1, 2, \ldots\}$. It can even take an arbitrarily large value, albeit with decreasing probability. The following result characterizes the PMF of H_k.

Lemma 11.1
For a homogeneous Markov chain with state transition matrix $\mathbf{P}$, the first exit time from state k, H_k is a geometric random variable with success probability $1 - \mathbf{P}_{kk}$.

To show this, note that, if the Markov chain is in state $X_0 = k$ at time 0, the probability that it exits at the next time is $1 - \mathbf{P}_{kk}$. Thus, $\mathbb{P}[H_k = 1] = 1 - \mathbf{P}_{kk}$. If it does not exit, then the chain remains in state $X_1 = k$ with probability $\mathbf{P}_{kk}$. The event that the chain exits at time 2 is independent of the prior history of the Markov chain, because of the Markov property, and has probability q_k of occurring. Hence, $\mathbb{P}[H_k = 2] = (1-\mathbf{P}_{kk})\mathbf{P}_{kk}$, and the probability that $X_2 = k$ is P^2_{kk}. Continuing by induction, we can establish that $\mathbb{P}[H_k = \ell] = (1 - \mathbf{P}_{kk})(\mathbf{P}_{kk})^{\ell-1}$, which is the PMF of a geometric random variable.

11.2.2 Evolution of marginal probabilities

Let $X_t, t = 0, 1, \ldots$ be a discrete time, finite-valued Markov chain with values in $\mathcal{R}_X = \{1, 2, \ldots, K\}$. The Markov chain has a marginal distribution at $t = 0$ as $P_{X_0}(x_0)$. We can represent this distribution as a vector, as illustrated below.

$$\underline{p}(0) = \begin{bmatrix} P_{X_0}(1) \\ P_{X_0}(2) \\ \vdots \\ P_{X_0}(K) \end{bmatrix}$$

Similarly, we denote the marginal PMF of X_t as a vector $\underline{p}(t)$, defined as

$$\underline{p}(t) = \begin{bmatrix} P_{X_t}(1) \\ P_{X_t}(2) \\ \vdots \\ P_{X_t}(K) \end{bmatrix}$$

The state transition matrix $\mathbf{P}$ can be used to compute the evolution of the marginal probability vectors $\underline{p}(t)$ over time, as follows: Note that, at time 1,

$$P_{X_0,X_1}(j,k) = P_{X_1|X_0}(k|j)P_{X_0}(j) = \mathbf{P}_{jk}P_{X_0}(j)$$

Hence, the marginal probability at time 1 is given by summing over the possible values j of X_0, as

$$P_{X_1}(k) = \sum_{j=1}^{K} \mathbf{P}_{jk}P_{X_0}(j)$$

which can be written in terms of matrix operations as

$$\underline{p}(1) = \mathbf{P}^T\underline{p}(0)$$

Extending the above argument inductively yields the following recursion:

$$\underline{p}(t) = \mathbf{P}(t)^T\underline{p}(0),$$

where $\mathbf{P}(m) \equiv \mathbf{P}^m$ for $m \geq 0$ is the m-step transition probability matrix. The multistep transition matrix satisfies the Chapman-Kolmogorov equation

$$\mathbf{P}(n+m) = \mathbf{P}(m)\mathbf{P}(n) = \mathbf{P}(n)\mathbf{P}(m) \quad \text{for } n, m \geq 0.$$

Note that $\mathbf{P}(0)$ is the K-dimensional identity matrix $\mathbf{I}_K$.

Note that the state transition matrix $\mathbf{P}$ and the multi-step transition matrix $\mathbf{P}(m)$ must satisfy the laws of conservation of probability. That is, for any row k, we must have

$$\sum_{j=1}^{\infty} \mathbf{P}_{kj} = 1; \quad \sum_{j=1}^{\infty} \mathbf{P}(m)_{kj} = 1;$$

Example 11.4

Assume a person starts in the middle of a room. At each time, with probability $p = 0.5$, they take a step to the right. With probability 0.5, they take a step to the left. However, if they are at the wall, and they try to take a step into the wall, they stay in place. Assume the walls on the left and right are five steps away from the center of the room. What is the probability that the person will be next to the right wall at time 10?

The figure below illustrates the Markov chain for this problem, under the assumption that $p = 0.5$. The starting position is in state 6, so that $P_{X_0}(6) = 1$. The state transition matrix is given by

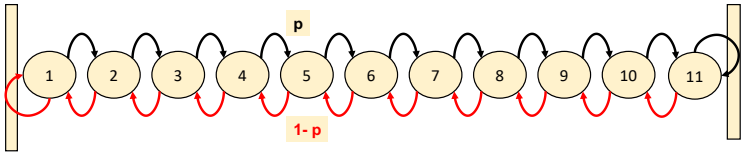

Figure 11.4: Random walk in a closed room.

$$\mathbf{P} = \begin{bmatrix} 1-p & p & 0 & 0 & \cdots & 0 \\ 1-p & 0 & p & 0 & \cdots & 0 \\ 0 & 1-p & 0 & p & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & 1-p & 0 & p \\ 0 & 0 & \cdots & 0 & 1-p & p \end{bmatrix}$$

We are interested in computing $\mathbf{P}(10)_{6,11}$, the probability that, starting at state 6 at time 0, we are in state 11 at time 10. By direct computation, we get $\mathbf{P}(10)_{6,11} = 0.0439$. How would this change if we increased the time to 20? The probability of being next to the wall increases to 0.0741. If we consider the same question at time 100, the probability increases to 0.0905. After 200 steps, the marginal probability vector is

$$\underline{p}(200) = \begin{bmatrix} 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \\ 0.0909 \end{bmatrix},$$

having reached a steady state.

Example 11.5
In Example 11.4, we saw the marginal probability vector $\underline{p}(t)$ approach a limit as $t \to \infty$. Do we see similar behavior in other examples? Consider the on-off system of Example 11.3. Let $p = 0.1, q = 0.2$. In this case, the state transition matrix is

$$\mathbf{P} = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix} = \begin{bmatrix} 0.9 & 0.1 \\ 0.2 & 0.8 \end{bmatrix}.$$

Assuming we start in the "on" state 1, we compute the marginal probability vector after 5, 10, 20, and 40 times. The results are shown below:

$$\underline{p}(5) = \begin{bmatrix} 0.7227 \\ 0.2773 \end{bmatrix};\ \underline{p}(10) = \begin{bmatrix} 0.6761 \\ 0.3239 \end{bmatrix};\ \underline{p}(20) = \begin{bmatrix} 0.6669 \\ 0.3331 \end{bmatrix};\ \underline{p}(40) = \begin{bmatrix} 0.6667 \\ 0.3333 \end{bmatrix}.$$

Again, we see the marginal probability vector approach a steady state with increasing t.

Assume that the marginal distribution vectors converge to a steady state marginal distribution $\underline{\pi}$. In this case, this steady state distribution must satisfy $\mathbf{P}^T\underline{\pi} = \underline{\pi}$. That implies that $\underline{\pi}$ is an eigenvector of the matrix $\mathbf{P}^T$, corresponding to an eigenvalue of 1. We know that $\mathbf{P}$ has an eigenvalue of 1, with eigenvector corresponding to the K-dimensional vector of all ones, because the sum of every row of $\mathbf{P}$ equals one. That is,

$$\mathbf{P}\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \begin{bmatrix} \sum_{k=1}^{K} P_{1k} \\ \sum_{k=2}^{K} P_{1k} \\ \vdots \\ \sum_{k=1}^{K} P_{Kk} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix}.$$

Since the eigenvalues of $\mathbf{P}$ and $\mathbf{P}^T$ are the same, $\mathbf{P}^T$ also has an eigenvalue of 1, with corresponding eigenvector. Note also that $\underline{\pi}$ is the limit of a sequence of marginal probability mass functions, and hence the limit will also be a valid probability mass function: $\pi_k \in [0,1], \sum_{k=1}^{K} \pi_k = 1$.

To better understand the limit behavior of Markov chains, we discuss the properties of stochastic matrices that control the evolution of the marginal distributions.

11.2.3 Stochastic matrices

When the number of states is finite and equal to K, the state transition matrix will be an $K \times K$ matrix $\mathbf{P}$, where $\mathbf{P}$ is such that all of its entries are nonnegative and the rows sum up to 1. Nonnegative matrices with the property that the rows sum up to 1 are known as **stochastic matrices**.

We first quote a theorem for linear algebra that relates the locations of the eigenvalues of matrices to the elements of its rows.

Theorem 11.1 (Gershgoren's Theorem)
Consider a square matrix $\mathbf{A}$ of dimension $K \times K$. Define distances $d_i = \sum_{j=1, j\neq i}^{n} |\mathbf{A}_{ij}|$. Define the set of complex numbers

$$L = \{\lambda \in \mathbb{C} : |\lambda - \mathbf{A}_{ii}| \leq d_i \text{ for some } i \in \{1, \ldots, K\}\}.$$

Then, all of the eigenvalues of $\mathbf{A}$ are contained in the set L.

The distance d_i is the sum of the magnitude of the off-diagonal elements in row i. The set L consists of the union of circles of radius d_i centered around each of the diagonal elements $\mathbf{A}_{ii}$. Figure 11.5 illustrates the implications of Gershgoren's theorem for the matrix $A = \begin{bmatrix} 3 & 2 \\ 1 & 1 \end{bmatrix}$. The eigenvalues must lie in the union of two circles in the complex plane, centered at the diagonal elements (3,0) and (1,0), with radii 2 and 1, respectively. By direct computation, the eigenvalues are 3.7321 and 0.2679, which are in the union of the two circles.

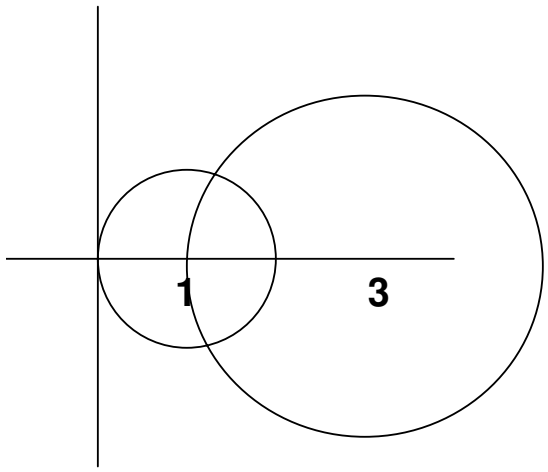

Figure 11.5: Illustration of Gershgoren's Theorem.

For stochastic matrices $\mathbf{A}$, the rows add up to 1, and all the elements are non-negative. This means that $d_i + \mathbf{A}_{ii} = 1$, and the center of the circle is on the non-negative real line. Hence, all of the eigenvalues of a stochastic matrix must be on or inside the unit circle of radius 1, centered at 0. Furthermore, since every row adds to 1, we know that the vector $\underline{1} = [1, 1, \ldots, 1]^T$ satisfies $\mathbf{A}\underline{1} = \underline{1}$, and is thus an eigenvector of the matrix $\mathbf{A}$ with eigenvalue equal to 1.

Figure 11.6 illustrates Gershgoren's theorem for stochastic matrices. Note that all of the eigenvalues λ of $\mathbf{A}$ must satisfy $|\lambda| \leq 1$.

The other theorem from mathematics that relates to the eigenvalues and eigenvectors of stochastic matrices is the Perron-Frobenius theorem, stated below:

Theorem 11.2 (Perron-Frobenius Theorem)
Consider a square matrix $\mathbf{A}$ of dimension $K \times K$ with non-negative elements. Then, there exists a non-negative real eigenvalue λ_{PF} with associated non-negative eigenvector, such that $|\lambda| \leq \lambda_{PF}$ for any other eigenvalue λ of $\mathbf{A}$. Furthermore, if $\mathbf{A}$ is such that $\mathbf{A}^k$ is strictly positive for some k, then $|\lambda| < \lambda_{PF}$ and the associated eigenvector with λ_{PF} can be chosen as strictly positive.

The Perron-Frobenius theorem establishes that $\lambda_{PF} = 1$ and that the associated eigenvector $\underline{\pi}$ can be chosen so that $\underline{\pi}$ is non-negative. Furthermore, it establishes the condition that is needed to ensure that $\underline{\pi} > 0$ and is a unique stationary density: if there exists k such that every element of $\mathbf{P}^k = \mathbf{P}(k)$ is positive. We will provide graphical conditions that are necessary and sufficient for this to be true.

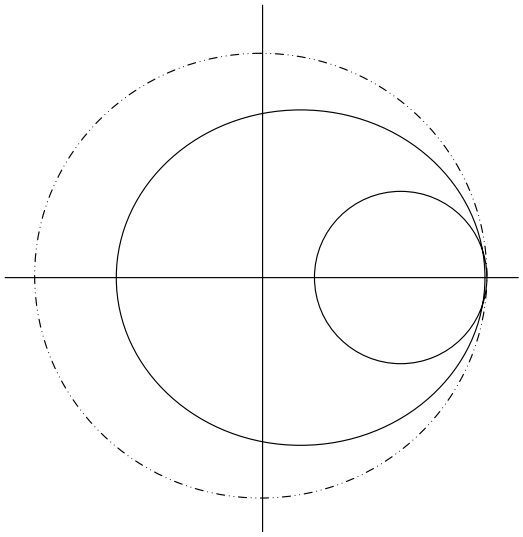

Figure 11.6: Illustration of Gershgoren's Theorem for stochastic matrices.

Example 11.6
Consider the "on"-"off" example in Example 11.5, with state transition matrix

$$\mathbf{P} = \begin{bmatrix} 1-p & p \\ q & 1-q \end{bmatrix}.$$

The eigenvalues of this matrix are the solution of the quadratic equation

$$(s-p)(s-q) - (1-p)(1-q) = s^2 - (p+q)s + pq - 1 + (p+q) - pq = (s-1)(s-(p+q-1)) = 0$$

which are $1, p+q-1$. The magnitude of the second eigenvalue is strictly less than 1, unless both p, q are either 0 or 1. Note that, if $p, q \in (0,1)$, then $\mathbf{P} > 0$ and, by the Perron-Frobenius Theorem, there is at most one eigenvalue with magnitude 1, and the limit eigenvector can be chosen to be strictly positive. The eigenvector of $\mathbf{P}^T$ corresponding to the eigenvalue 1 satisfies:

$$\mathbf{P} = \begin{bmatrix} 1-p & q \\ p & 1-q \end{bmatrix} \underline{\pi} = \underline{\pi}.$$

This results in the equations

$$\begin{aligned} (1-p)\pi_1 + q\pi_2 = \pi_1 &\iff -p\pi_1 + q\pi_2 = 0 \\ p\pi_1 + (1-q)\pi_2 = \pi_2 &\iff p\pi_1 - q\pi_2 = 0 \end{aligned}$$

which reduce to $\pi_2 = \frac{p}{q}\pi_1$. To find π_1, we use the normalization property of PMFs, which says that $\pi_1 + \pi_2 = \pi_1(1+\frac{p}{q}) = 1$. This implies that $\pi_1 = \frac{q}{p+q}, \pi_2 = \frac{p}{p+q}$.

11.2.4 Steady-state behavior of Markov chains

As discussed previously, the marginal probability mass function $\underline{p}(t)$ evolves according to a linear system:

$$\underline{p}(t+1) = \mathbf{P}^T \underline{p}(t)$$

For homogeneous Markov chains in discrete time, this equation may have a limit as $t \to \infty$, as all the eigenvalues of $\mathbf{P}$ will have magnitude less than or equal to 1. We are interested in providing conditions where

$$\lim_{t\to\infty} \mathbf{P}^t = \mathbf{P}_\infty$$

and

$$\lim_{t\to\infty} \underline{p}(t) = \lim_{t\to\infty} (\mathbf{P}^t)^T \underline{p}(0) = \mathbf{P}_\infty^T \underline{p}(0) = \underline{\pi}$$

To illustrate issues that can arise, consider the two graphs illustrated in Figures 11.7(a) and 11.7(b). The first graph shows that, after starting in state 2, one can either go to state 1 or to states 3 and 4. Depending on which transition is used, the limit will be different. It is clear that this Markov chain may have multiple limiting distributions. The second figure illustrates a more complex case. If one starts in state 1 at time 0, note that one can only be in an odd-valued state at even times! This Markov chain will not approach a limit, but rather will oscillate between two limits!

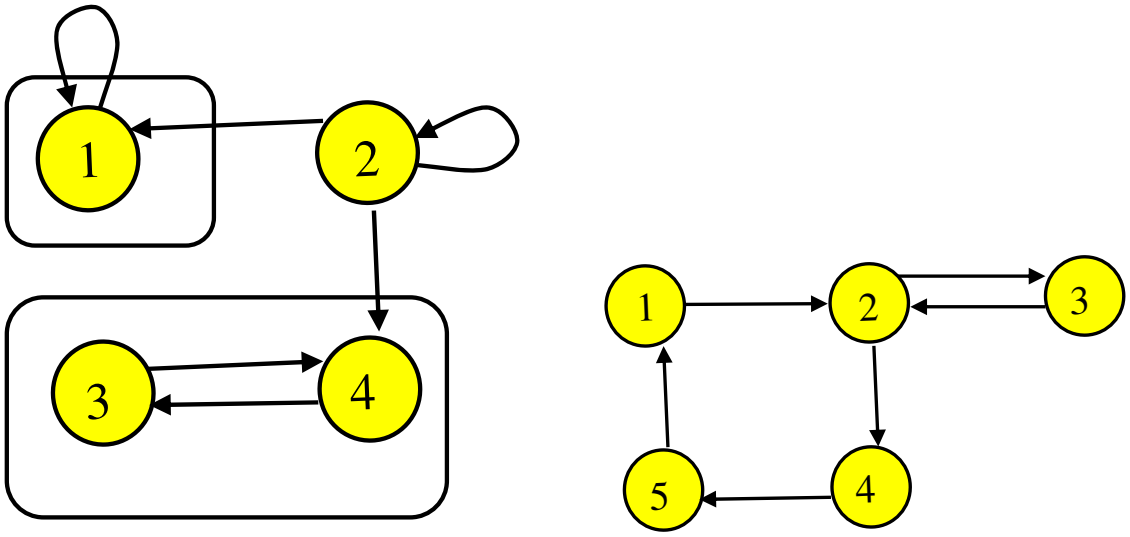

(a) Markov chain with two limits (b) Period 2 Markov chain

Figure 11.7: Illustration of Markov chains with difficult limit behavior.

For finite state Markov chains, one can define regularity conditions that guarantee that there is a unique eigenvalue of P with magnitude 1, so that there are unique limits. Furthermore, these conditions can be established from the transition diagram of the Markov chain! We discuss these next.

Consider two states i, j of the Markov chain. State j is said to be **accessible** from state i if there exists a time n such that $(P^n)_{ij} > 0$. An equivalent graphical condition is that there exists a *directed path* with positive probability arcs from node i to node j in the Markov chain graph. In the reflected random walk diagram of Figure 11.4 in Example 11.4, every state is accessible from every other state. However, consider the minor variation shown in Figure 11.8, where one of the feasible arcs has been removed. In this case, state 7 is accessible from state 6, but state 6 is not accessible from state 7.

Two states i, j are said to **communicate** if i is accessible from j and j is accessible from i; by convention, every state is said to communicate with itself. Communication is a transitive, symmetric and reflexive binary relationship, hence it is an equivalence relationship. A **communicating class** is a non-empty set of states that communicate with each other, and no state in the class communicates with any state outside the class. The set of possible states of a finite-valued Markov Chain can be partitioned into disjoint communicating classes. For instance, the Markov Chain illustrated in Figure 11.8 has 2 communicating classes: $\{1, 2, 3, 4, 5, 6\}$ and $\{7, 8, 9, 10\}$.

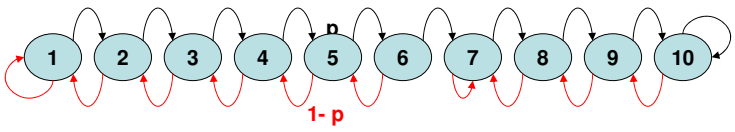

Figure 11.8: Example of Markov Chain with inaccessible states

When a Markov Chain has only one communicating class, it is said to be **irreducible**. In irreducible Markov Chains, every state communicates with every other state, as in Fig. 11.4.

A state i in a homogeneous Markov Chain is said to be **transient** if, given that the Markov Chain starts at state i, there is a non-zero probability that the state never returns to state i. Formally, assume $X_0 = i$, and define the random time $T = \min[t > 0 : X_t = i]$. Then, $\mathbb{P}[\{T = \infty\}] > 0$. Thus, there is positive probability that, when the trajectory of the Markov chain leaves a transient state, it will never return to it.

For finite-state Markov Chains, there is graphical way of identifying a transient state: A state i is transient if and only if there is a second state j such that j is accessible from i, but i is not accessible from j. In Figure 11.8, states 1, 2, 3, 4, 5 and 6 are transient states, and they can each access state 7, but cannot be accessed from state 7. Note that, if a state i is transient, every other state k in its communicating class is also transient, because that state k can communicates with state i and therefore can access a state j not in its communicating class.

When a state is not transient, it is called **recurrent**: recurrent states have the property that the expected time to return to the state, given that the Markov Chain starts in that state, is finite. In terms of the random time T defined previously, $\mathbb{E}[T] < \infty$ for recurrent states. In Fig. 11.8, states 7, 8, 9 and 10 are recurrent states. Note that, for finite state Markov Chains, we can label each communicating class as either recurrent or transient.

The meaning of transient states is that, as time grows, the probability of being in a transient state decays to zero. If there is a limiting probability distribution $\underline{\pi}$ and state i is transient, then $\underline{\pi}_i = 0$.

Note the following: If a finite state Markov chain has more than one recurrent communicating class, there will be more than one limiting distribution for $\underline{p}(t)$, and the limit will depend on the initial distribution $\underline{p}(0)$. The matrix P will have more than one eigenvalue equal to 1. This is the case in the Markov Chain in Fig. 11.7, where state 1 is one recurrent communicating class, and states $3, 4$ are the other recurrent communicating class.

When there is only one recurrent communicating class, there is a unique stationary probability distribution $\underline{\pi}$ such that

$$\mathbf{P}^T \underline{\pi} = \underline{\pi} \tag{11.1}$$

Specifically, the matrix P will have a single eigenvalue with value 1. However, this condition is insufficient to guarantee that this stationary probability distribution will be the limit distribution for arbitrary initial probability distributions.

Specifically, consider Fig. 11.7(b). It is easy to verify that all states belong to a single communicating class, which is recurrent. However, we have already established that, starting from the initial condition $X_0 = 1$, the probabilities $\underline{p}(t)$ do not approach a limit! Indeed, they will approach a limit cycle where they will shift among two different limits for odd and even values of n. In this case, there is a second eigenvalue of $\mathbf{P}$ on the unit circle, with value -1.

For a finite state Markov Chain, we define the **period** of state j as the greatest common divisor of the lengths of all the cycles from state j to itself in the graph of the Markov Chain. A more mathematical definition is that the period d is the largest integer d such that $(\mathbf{P}^n)_{jj} = 0$ unless n is divisible by d. A state with period 1 is said to be *aperiodic*.

Note that the *period of all the states in the same communicating class must be the same*. This follows because of the cycles for a state k in this communicating class must consist of states in that communicating class. The proof of this is a bit involved but straightforward from the definition.

A communicating class is periodic with period d if every state has period d greater than 1. There is a simple condition to recognize whether a communicating class is aperiodic: As long as one of the states in the communicating class has a self-loop (e.g. $\mathbf{P}_{ii} > 0$ for some i), the period of that state is 1, and the communicating class must be aperiodic.

We can now give conditions for a finite state Markov Chain to have a unique limiting probability distribution $\underline{\pi}$, which is approached from any initial probability distribution $\underline{p}(0)$. We state this below as a theorem.

Theorem 11.3
Assume that X_t is a finite state homogeneous Markov chain with state transition matrix $\mathbf{P}$. If the Markov chain has a single recurrent communicating class, and the class is *aperiodic*, then there exists a unique limit distribution $\underline{\pi}$.

Note that a Markov chain with transient states can approach a unique limit distribution $\underline{\pi}$ as long as there is only one recurrent, aperiodic communicating class. This limit distribution will have $\pi_k = 0$ for all transient states k.

There is a stronger result for the special case of irreducible Markov chains which have a single communicating class.

Theorem 11.4
Assume that X_t is a finite state homogeneous Markov chain with state transition matrix $\mathbf{P}$. If the Markov chain is *irreducible* and *aperiodic*, then there exists a unique limit distribution $\underline{\pi}$. Furthermore, this limit has the property that $\underline{\pi}_j > 0$ for all states j. Such a Markov Chain is called **ergodic**.

The combination of the irreducible and aperiodic conditions imply that there exists $k > 0$ such that $\mathbf{P}^k > 0$, that is, a matrix with strictly positive entries. In this case, the Perron-Frobenius theorem establishes the existence of a unique eigenvector of $\mathbf{P}^T$ for the eigenvalue 1 with strictly positive elements. The limit distribution $\underline{\pi}$ is this unique positive eigenvector of the matrix $\mathbf{P}^T$ corresponding to the eigenvalue 1, normalized so that its entries that sum up to 1.

11.2.5 Computing stationary probability distributions

An important problem in the analysis of Markov chains is computing the stationary probability distribution $\underline{\pi}$. The algebraic characterization is $\mathbf{P}^T\underline{\pi} = \underline{\pi}$, where $\mathbf{P}$ is the state transition matrix. This can be a cumbersome set of equations to solve. There is another set of equations based on the graphical representation of the Markov chain transitions that can be easier to analyze. A **cut C** of a directed graph is a set of arcs such that, when the arcs are removed from the graph, the graph is divided into two disjoint set of nodes with no arcs between them.

The useful property of cuts is that, given any cut of the Markov chain graph, the probability flow across that cut must equal zero once the system reaches the stationary distribution. A cut C specifies a subset $A \subset \mathcal{R}_X$ and its complement A^c in $\mathcal{R}_X$, and consists of the arcs going from A to A^c, and from A^c to A. Given a distribution $\underline{\pi}$ on the states of the Markov Chain, the net probability flow on a cut C is defined as

$$F(A, A^c) = \sum_{i\in A}\sum_{j\in A^c} \mathbf{P}_{ij}\pi_i - \sum_{j\in A^c}\sum_{i\in A} \mathbf{P}_{ji}\pi_j$$

The main result is that, if $\underline{\pi}$ is a stationary distribution of a Markov chain, then the net probability flow along any cut must be zero! This is summarized in the theorem below:

Theorem 11.5
$\underline{\pi}$ is a stationary distribution of a Markov chain if and only if $\sum_i \pi_i = 1$ and the net probability flow on any cut in the Markov chain graph is zero. That is, for any $A \subset \mathcal{R}_X$, we have

$$\sum_{i\in A}\sum_{j\in A^c} \mathbf{P}_{ij}\pi_i - \sum_{j\in A^c}\sum_{i\in A} \mathbf{P}_{ji}\pi_j = 0$$

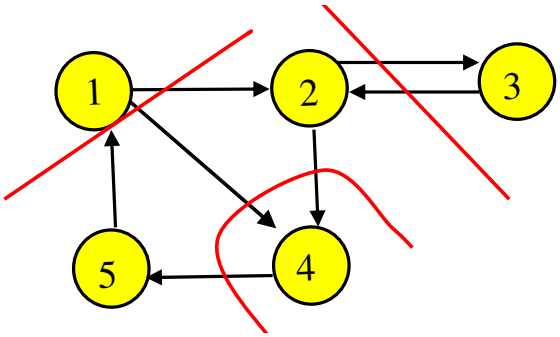

Figure 11.9: Illustration of probability balance

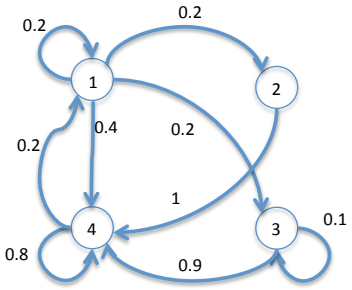

Figure 11.10: Diagram of the Markov Chain for the example

This property is referred to as **probability balance**.

To see that the theorem is equivalent to stationarity, note that if we select $A = \{i\}$, you get exactly the balance equations for the eigenvector:

$$\sum_{j \neq i} \mathbf{P}_{ij}\pi_i = \sum_{j \in \mathcal{R}_X, j \neq i} \mathbf{P}_{ji}\pi_j.$$

If we add $\mathbf{P}_{ii}\pi_i$ to both sides, we have

$$\left(\sum_{j=1}^{K} \mathbf{P}_{ij}\right)\pi_i = \pi_i = \sum_{j=1}^{K} \mathbf{P}_{ji}\pi_j.$$

This is the i-th equation of $\mathbf{P}^T\underline{\pi} = \underline{\pi}$. It is also easy to show the converse, so that starting from balance equations, one can show flow in and out of any group of states is zero for stationary distributions.

Why is this useful? Sometimes, it is easy to identify cuts that yield equations that are simpler than the eigenvector equations. To illustrate how to use probability balance to compute stationary distributions, consider the example in Figure 11.9. The example shows three different cuts, that separate the graph into two disconnected sets of nodes with no arcs across them. Applying fflow balance to each of these cuts yields the equations:

$$\begin{aligned} \mathbf{P}_{14}\pi_1 + \mathbf{P}_{12}\pi_1 - \mathbf{P}_{51}\pi_5 &= 0 \\ \mathbf{P}_{23}\pi_2 - \mathbf{P}_{32}\pi_3 &= 0 \\ \mathbf{P}_{24}\pi_2 + \mathbf{P}_{14}\pi_1 - \mathbf{P}_{45}\pi_4 &= 0. \end{aligned}$$

The above yields three equations in five unknowns, so it is insufficient to find a solution. We can add another cut, isolating state 5, to obtain the following equation: $\mathbf{P}_{51}\pi_5 = \mathbf{P}_{45}\pi_4$. Other cuts are possible, but will be redundant with these equations. Notice that none of those equations include a constant, so the solution $\pi_i = 0$ satisfies the equations. Just like we had to do in the eigenvector method, we must add a normalization equation:

$$\sum_{j=1}^{5} \pi_j = 1.$$

With that as a fifth equation, we now have a unique solution which will yield a positive, normalized $\underline{\pi}$.

Example 11.7
Consider a 4-state discrete time Markov chain, with transition probability matrix described below:

$$\mathbf{P} = \begin{pmatrix} 0.2 & 0.2 & 0.2 & 0.4 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0.1 & 0.9 \\ 0.2 & 0 & 0 & 0.8 \end{pmatrix}$$

The graph illustrating the transitions of this Markov chain is shown in Fig. 11.10:

Looking at the diagram, it is easy to see that all 4 states are recurrent, as there are directed paths from any one state to any other state. Thus, the chain has a single recurrent communicating class, and thus is irreducible. One can also determine that the Markov chain is aperiodic, because there are some self-loops of length 1. Thus, the Markov chan has a unique steady state distribution, which can be computed as follows: To compute the steady state distribution, we need 4 equations. One of them is:

$$\pi_1 + \pi_2 + \pi_3 + \pi_4 = 1$$

To find 3 others, cut node 2 away from the graph. The flow on that cut yields:

$$0.2\pi_1 = \pi_2$$

Cut node 3 away from the graph, to get:

$$0.2\pi_1 = 0.9\pi_3$$

To get the last equation, we can cut around node 1 to get:

$$0.8\pi_1 = 0.2\pi_4$$

Using the last 3 equations, we get:

$$\pi_2 = \pi_1/5;\ \pi_3 = 2\pi_1/9; \pi_4 = 4\pi_1$$

Substituting into the first equation yields:

$$\pi_1(1 + 1/5 + 2/9 + 4) = 1 \Rightarrow \pi_1 = \frac{45}{244}$$

$$\pi_2 = \frac{9}{244}; \quad \pi_3 = \frac{10}{244}; \quad \pi_r = \frac{180}{244}$$

Example 11.8
We want to model a counter that behaves as follows: The counter has three states: $\mathcal{R}_X = \{1, 2, 3\}$. When the counter is in state 3, it shifts to state 2 at the next time. When it is in state 2, it shifts to state 1 at the next time. When it is in state 1, it shifts to states 1, 2, or 3 at the next time, each with probability $\frac{1}{3}$.

The state transition matrix of this Markov chain is $\mathbf{P} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$. The state transition diagram for the Markov chain is show in Figure 11.11.

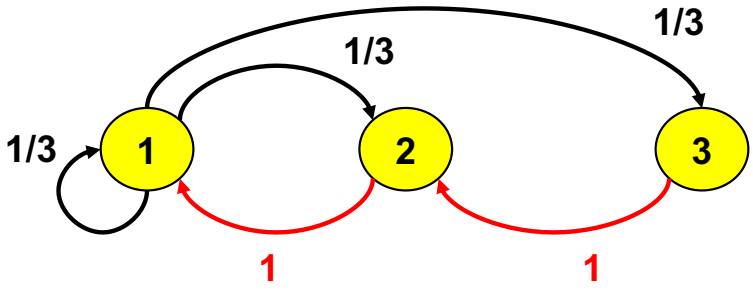

Figure 11.11: Diagram of the Markov chain for Example 11.8

A quick analysis of the graph shows that there is a single recurrent class, and that there are no transient states. Cuts around nodes 3 and 1 plus the normalization equations yields the following equations:

$$\frac{\pi_1}{3} = \pi_3$$

$$\frac{2\pi_1}{3} = \pi_2$$

$$\pi_1 + \pi_2 + \pi_3 = 1$$

Solving this yields the stationary distribution: $\underline{\pi} = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{3} \\ \frac{1}{6} \end{bmatrix}$.

Example 11.9
Consider a model of a reflected random walk in a contained space. The state space is $\mathcal{R}_X = \{1, 2, \ldots, 10\}$. At each time t, if the state k is in $\{2, \ldots, 9\}$, the next state is $k+1$ with probability p and $k-1$ with probability $1-p$. If the current state is $k = 1$, then the next state is 1 with probability $1-p$ and 2 with probability p. If the current state is $k = 10$, the next state is 10 with probability p and 9 with probability $1-p$. The diagram of the Markov chain is displayed in Figure 11.12.

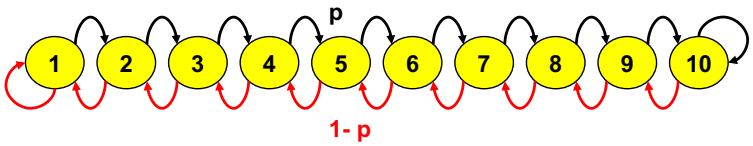

Figure 11.12: Diagram of the Markov chain for Example 11.9.

It is clear that the graph of this Markov chain is irreducible, and the presence of two self-loops makes it aperiodic. Hence, there is a unique stationary distribution. The linear structure of the Markov chain graph makes it easy to find 9 cuts, separating states $k, k+1$, for $k = 1, 2, \ldots, 9$. These cuts yield the following equations:

$$\begin{aligned} &p\pi_1 = (1-p)\pi_2; \quad && p\pi_2 = (1-p)\pi_3; \quad && p\pi_3 = (1-p)\pi_4; \\ &p\pi_4 = (1-p)\pi_5; \quad && p\pi_5 = (1-p)\pi_6; \quad && p\pi_6 = (1-p)\pi_7; \\ &p\pi_7 = (1-p)\pi_8; \quad && p\pi_8 = (1-p)\pi_9; \quad && p\pi_9 = (1-p)\pi_{10}; \end{aligned}$$

Solving, we get the following relationships:

$$\pi_k = \left(\frac{p}{1-p}\right)^{k-1} \pi_1, \ k = 2, \ldots, 10.$$

The tenth equation needed is the normalization equation:

$$\sum_{k=1}^{10} \pi_k = 1 \iff \sum_{k=1}^{10} \left(\frac{p}{1-p}\right)^{k-1} \pi_1 = 1.$$

Fortunately, we can sum this term:

$$\sum_{k=1}^{10} \left(\frac{p}{1-p}\right)^{k-1} = \frac{1 - \left(\frac{p}{1-p}\right)^{10}}{1 - \frac{p}{1-p}}.$$

Hence, $\pi_1 = \frac{1-\frac{p}{1-p}}{1-\left(\frac{p}{1-p}\right)^{10}}$, and $\pi_k = \left(\frac{p}{1-p}\right)^{k-1} \pi_1, \ k = 2, \ldots, 10$. This expression is valid as long as $p \neq 1-p$. Thus, if $p = 0.4$, we obtain $\pi_1 = 0.3392$, and $\pi_{10} = 0.0088$.

If we have symmetry, and $p = 1-p = 0.5$, the balance equations indicate that $\pi_j = \pi_k$ for all $j, k \in 1, \ldots, 10$ so the steady-state distribution is $\pi_k = 0.1$.

Although we have focused our analysis on ergodic Markov chains so far, it is often possible to analyze the limiting behavior of non-ergodic Markov chains. We illustrate this with two different examples.

Example 11.10
Consider the Markov chain with state transition diagram shown in Figure 11.13. The Markov chain has a single recurrent class, but has period 2. The state transition matrix is

$$\mathbf{P} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0.5 & 0.5 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The state transition matrix still has an eigenvalue of 1, and there is a stationary distribution $\underline{\pi}$, which we can find using probability balance, as:

$$\pi_5 = \pi_1; \quad \pi_4 = \pi_5; \quad \pi_3 = 0.5\pi_2; \quad \pi_4 = 0.5\pi_2$$

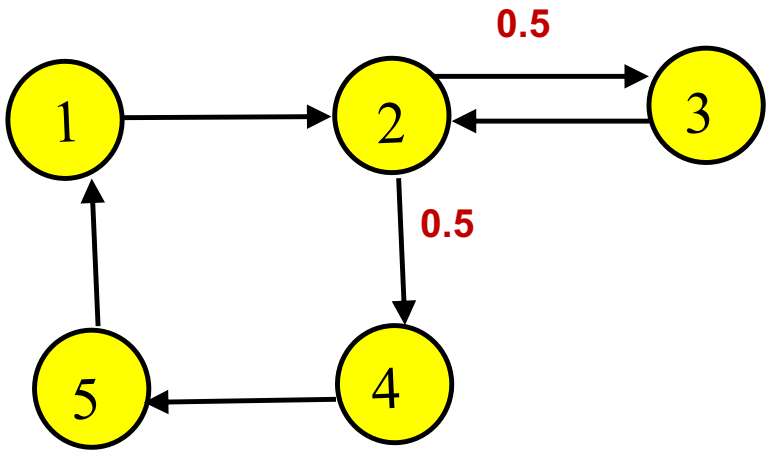

Figure 11.13: Diagram of the Markov chain for Example 11.10.

This means $\pi_1 = \pi_3 = \pi_4 = \pi_5$, and $\pi_2 = 2\pi_1$. Coupled with the normalization constraint, this yields $6\pi_1 = 1$, so $\pi_1 = \pi_3 = \pi_4 = \pi_5 = \frac{1}{6}, \pi_2 = \frac{1}{3}$. If the Markov chain starts with this distribution, it will stay in this distribution.

However, for different initial conditions, the limiting behavior will oscillate between two distributions, depending on the initial condition, and it won't converge to the stationary distribution. For instance, if $X_0 = 1$, the two distributions in the limit are $\begin{bmatrix} 0 \\ \frac{2}{3} \\ 0 \\ 0 \\ \frac{1}{3} \end{bmatrix}$ and $\begin{bmatrix} \frac{1}{3} \\ 0 \\ \frac{1}{3} \\ \frac{1}{3} \\ 0 \end{bmatrix}$.

Example 11.11
Consider the Markov chain with state transition diagram shown in Figure 11.13. The Markov chain has two communicating classes (states 1, 2, 3, and states 4, 5), but it has a single recurrent class (1, 2, 3). The state transition matrix is

$$\mathbf{P} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & p & 0 & 1-p \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}.$$

Since we know there is no steady state probability in the two transient states (4, 5), we can simply restrict our analysis

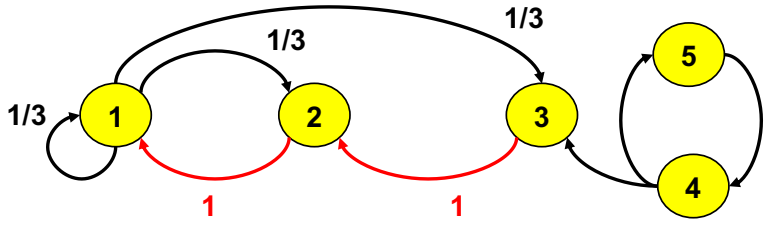

Figure 11.14: Diagram of the Markov chain for Example 11.11.

to the recurrent class, and analyze the steady state behavior of a 3 state model, with transition probability matrix

$$\mathbf{P}_r = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

This is the same Markov chain we analyzed in Example 11.6. Thus, the steady state probability in the original Markov

chain is

$$\underline{\pi} = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{3} \\ \frac{1}{6} \\ 0 \\ 0 \end{bmatrix}$$

11.3 Markov chains with infinite state spaces

The above discussion focused on finite state Markov chains, where the state space $\mathcal{R}_X$ has a finite number of states. What changes when the state space is infinite? We can no longer use linear algebra to establish our results, as the transition probability function $\mathbf{P}_{ij}$ does not have a convenient representation as a finite matrix. We highlight some of the key issues and differences below.

Example 11.12
Consider a random walk with probability 0.5 of going forward or back at each time. For this Markov chain, the state space is the space of integers: $\mathcal{R}_X = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$. It is easy to see that every state communicates with every other state. This Markov chain has period 2, and has a single communicating class. However, there cannot be an equilibrium distribution: we are no longer guaranteed that there is a positive "eigenvector" with eigenvalue 1. If there were, note that, by symmetry, every equilibrium state should have the same probability. However, since there are an infinite number of states, there is no way to select such a probability to satisfy the normalization condition $\sum_{k \in \mathcal{R}_X} \pi_k = 1$.

Example 11.13
Consider a Markov chain defined on the non-negative numbers as follows: $\mathbf{P}_{00} = 1/2, \mathbf{P}_{01} = 1/2$. For $k > 0$, $\mathbf{P}_{(k-1)k} = \mathbf{P}_{k(k+1)} = 1/2$. All other $\mathbf{P}_{ij} = 0, |i - j| \geq 2$. This chain is aperiodic (state 0 has a self-transition, so it has period 1) and has a single communicating class. However, this chain will not have an equilibrium distribution. Looking at balance equations, cutting between states i and j, we the relation:

$$\pi_k = \pi_{k+1}, k = 0, 1, \ldots$$

Hence, every state would have the same steady state probability, but with an infinite number of states, they would all be zero, a contradiction!

One way of seeing this is to look at the expected time to reach state 0 from state n. As we will show later with our transient analysis, no matter what state you start in, the expected number of steps it takes to reach state 0 is infinite!

The first important difference when the Markov chain has an infinite number of states is in the concept of recurrence. When the state transition graph was irreducible and the state space was finite, we could guarantee that $\mathbf{P}^n_{ij} > 0$ for every pair of states i, j; thus, with probability 1, we would visit state j when we start in state i in finite expected time. When the state space is infinite, this condition of irreducibility is no longer sufficient.

Let X_t be a time-homogeneous Markov chain with transition probability $\mathbf{P}$. Note that the state space $\mathcal{R}_X$ may be infinite. Define the following quantities:

$$T_i = \inf\{t \geq 1 : X_t = i\} = \textbf{ first passage time } \text{for state } \ i$$

When $X_0 = i$, then T_i is the revisit time for state i. We can now define some useful quantities relating how X_t visits a particular state i. Let $\mathcal{I}\{X_t = i\}$ denote the indicator function which is 1 when the event $X_t = i$ is true, and zero otherwise. Then,

$$V_i = \sum_{t=0}^{\infty} \mathcal{I}\{X_t = i\} \ \text{ is the number of visits to state } \ i$$
$$f_i = \mathbb{P}[T_i < \infty | X_0 = i] \text{ is the probability that the chain revisits state } i$$
$$m_i = \mathbb{E}[T_i | X_0 = i] \text{ is the expected return time to state } i$$

Consider the case of a finite-state aperiodic Markov chain with a single recurrent communicating class, but with some transient states. Let i be a transient state. Then, V_i is finite, and $f_i < 1$. However, if i is a recurrent state, we get that $V_i = \infty$ with probability 1, $f_i = 1$ and $m_i < \infty$, so that the chain continues to revisit state i. We use these concepts to extend the definition of recurrence to infinite state Markov chains:

Definition 11.1
A state i of a homogeneous Markov chain $\{X_t, t = 0, 1, \ldots\}$ is recurrent if

$$\mathbb{P}[V_i = \infty | X_0 = i] = 1.$$

A recurrent state is one that you return to an infinite number of times. Indeed, we can characterize a recurrent state as one for which $f_i = 1$, and a transient state as one for which $f_i < 1$. When the state space is infinite, we don't have simple graphical characterizations of what recurrent and transient states are. However, we can use the transition probabilities to get equivalent definitions:

Theorem 11.6
State i in a homogeneous Markov chain is recurrent if and only if

$$\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} = \infty$$

To show this, note that for recurrent i, one has $\mathbb{P}[V_i = \infty | X_0 = i] = 1$. Note also the following interpretation:

$$(\mathbf{P}^n)_{ii} = \mathbb{P}[X_n = i | X_0 = i]$$

where $\mathbf{P}^n$ is the n-step transition probability kernel $\mathbf{P}(X_n = j | X_0 = i)$, which can be obtained through direct application of the one-step kernel n times. Thus,

$$\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} = \sum_{n=0}^{\infty} \mathbb{E}[\mathcal{I}\{X_n = i\} | X_0 = i] = \mathbb{E}[\sum_{n=0}^{\infty} \mathcal{I}\{X_n = i\} | X_0 = i] = E[V_i | X_0 = i] = \infty$$

If i is a transient state, then $f_i < 1$. We can view the return process as a geometric random variable because of the Markov nature of the process X_t. The first return occurs with probability f_i, the second return with probability f_i^2, etc. Thus, the expected number of returns is $\frac{1}{1-f_i}$, which is finite. By the above argument, for transient states i, $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} < \infty$.

We can now use the same definitions we had previously for communicating classes. State i communicates with state j if $(\mathbf{P}^n)_{ij} > 0$ for some $n \geq 1$ and $(\mathbf{P}^m)_{ji} > 0$ for some $m \geq 1$. A communicating class C is a set of states such that, if $i, j \in C$, then i communicates with j. Furthermore, there are no states $k \notin C$ such that a state $j \in C$ communicates with state k.

Theorem 11.7
Let C be a communicating class in the homogeneous Markov chain X_t. Then, either all states in C are recurrent or all states in C are transient.

To see this, take any pair of states $i, j \in C$ and suppose that i is a transient state. Since i, j communicate, there exists $n, m \geq 0$ with $(\mathbf{P}^n)_{ij} > 0, (\mathbf{P}^m)_{ji} > 0$. Then, for any $r \geq 0$,

$$(\mathbf{P}^{n+m+r})_{ii} \geq (\mathbf{P}^n)_{ij} (\mathbf{P}^r)_{jj} (\mathbf{P}^m)_{ji}$$

So,

$$(\mathbf{P}^r)_{jj} \leq \frac{1}{(\mathbf{P}^n)_{ij} (\mathbf{P}^m)_{ji}} (\mathbf{P}^{n+m+r})_{ii}$$

Summing over all $r \geq 0$ yields

$$\sum_{r=0}^{\infty} (\mathbf{P}^r)_{jj} \leq \frac{1}{(\mathbf{P}^n)_{ij} (\mathbf{P}^m)_{ji}} \sum_{r=0}^{\infty} (\mathbf{P}^{n+m+r})_{ii}$$

The last sum is finite since i is transient, so the left hand side is also finite, indicating that j is also transient.

As was the case for finite state Markov chains, every recurrent communicating class will be closed: once a Markov chain enters a state in a recurrent class, the future states in the chain must belong to the same recurrent class. Otherwise, there would be a state i in the recurrent class that communicates with a transient state j (so $(\mathbf{P}^n)_{ij} > 0$ for some $n \geq 1$) but j does not communicate with i. We can thus show that this contradicts $\mathbb{P}[V_i = \infty] = 1$, so that i won't get revisited infinitely.

However, the converse is not true. If we have a closed communicating class, it may not be recurrent! We do have the following result: if a closed communicating class has a finite number of states, it must be recurrent. However, there will be examples of closed communicating classes that won't be recurrent. Examples 11.12 and 11.13 show closed communicating classes that are not recurrent.

Recurrence is the key property for extending our previous results to infinite Markov chains. The implications of recurrence are summarized below:

Theorem 11.8
Suppose $\mathbf{P}$ has a single communicating class C, which is recurrent. Then, for every state $j \in C$, $\mathbb{P}[T_j < \infty] = 1$.

We now focus on the steady state behavior. Does a steady state distribution exist? Can there be more than one? How can one calculate it? We define a couple of useful variables to help understand this behavior. Remember that T_k is the first return time for state k. Let

$$\begin{aligned} V_i^k &= \sum_{n=0}^{T_k} \mathcal{I}\{X_t = i\} = \text{ number of visits to state } i \text{ before visiting state } k. \\ \gamma_i^k &= \mathbb{E}[V_i^k | X_0 = k] \text{ expected number of visits to } i \text{ before revisiting } k \\ V_i(n) &= \sum_{k=0}^{n} \mathcal{I}\{X_t = i\} \text{ number of visits to state } i \text{ before time } n \end{aligned}$$

If there were an invariant distribution $\pi_i, i \in \mathcal{R}_X$, then one would like to show

$$\mathbb{E}[T_i | X_0 = i] = \frac{1}{\pi_i}, \quad \gamma_i^k = \frac{\pi_i}{\pi_k}$$

and

$$\lim_{n\to\infty} \frac{V_i(n)}{n} = \pi_i$$

almost everywhere.

The main result for existence and uniqueness of steady state distributions for general Markov chains requires two items: First, one must have recurrent states. Second, one must have the property that, for a recurrent state, the expected return time is finite. We call a state i *positive recurrent* if it is recurrent and $m_i = \mathbb{E}[T_i | X_0 = i] < \infty$. When a recurrent state has infinite expected return time, we call it *null recurrent*.

Theorem 11.9
Let $\mathbf{P}$ be the state transition kernel of an irreducible Markov chain. Then, the Markov chain has a positive recurrent state i if and only if it has an invariant distribution π. Furthermore, if it has an invariant distribution, then all states are positive recurrent, and $\mathbb{E}[T_i | X_0 = i] = \frac{1}{\pi_i}$ for all states i.

Note that this does not guarantee that all initial distributions approach the invariant distribution π. The problem is that we can still have periodic chains! Here is the final extension that we need:

Theorem 11.10
Let $\mathbf{P}$ be the transition probability kernel of an irreducible, aperiodic, positive recurrent Markov chain (also called ergodic), with invariant distribution π. Then, for any initial distribution, the marginal probabilities converge: $P_{X_n}(j) \to \pi_j$ as $n \to \infty$ for all j. In particular,

$$\lim_{n\to\infty} (\mathbf{P}^n)_{ij} = \pi_j$$

The difficulty in applying this theorem is that computing whether the Markov chain is positive recurrent is equivalent to finding the stationary probability distribution. In practice, we simply try to compute the stationary distribution using the properties of probability balance, and either we can find it, or we find a contradiction that shows such a stationary probability distribution cannot exist.

Computing the stationary distribution of ergodic Markov chains when the state space is infinite can be done using the balance equations $\underline{\pi} = \mathbf{P}^T \underline{\pi}$, where the vector notation is extended to infinite dimensions. This will now require solution of an infinite number of linear equations. The use of cuts is helpful in getting these equations into simple form, as illustrated below.

Example 11.14
Consider a Markov chain defined on the non-negative numbers, which is a model for a single-server queue. The state value represents the number of elements in a queue. The transition probabilities are $\mathbf{P}_{00} = 1-p; \mathbf{P}_{k(k+1)} = p, k = 0, 1, 2, \ldots; \mathbf{P}_{k(k-1)} = 1-p, k = 1, 2, 3, \ldots$. A state transition diagram of this Markov chain is shown in Figure 11.15. It

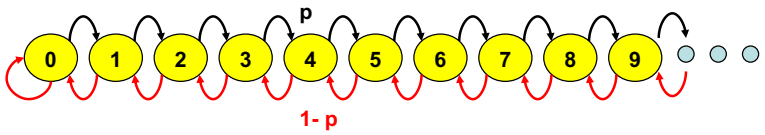

Figure 11.15: Diagram of the Markov chain for Example 11.14.

is clear that there is a single communicating class in this chain, and that the chain is irreducible, as there are no transient states. Furthermore, the chain is aperiodic because of the self-transition present in state 0, that makes the period of the chain equal to 1. Note that, if $p = 1 - p = 0.5$, this is the Markov chain we discussed in Example 11.13.

Assume $p < q$. Since this chain is linear, we can find cuts between any pair of consecutive states. For a cut between states k and $k+1$, probability balance yields the following equation:

$$p\pi_k = (1-p)\pi_{k+1} \iff \pi_{k+1} = \frac{p}{1-p}\pi_k = \left(\frac{p}{1-p}\right)^{k+1} \pi_0, \ k = 0, 1, \ldots$$

Define the utilization factor $\alpha = \frac{p}{1-p}$. Then, we have $\pi_{k+1} = \alpha^{k+1}\pi_0$. Substituting this into the normalization equation yields

$$\sum_{t=0}^{\infty} \alpha^t \pi_0 = 1 \iff \frac{\pi_0}{1-\alpha} = 1 \iff \pi_0 = 1 - \alpha,$$

where we have used the formula for summing a geometric series. Note that this sum exists only for $\alpha < 0$, which means $p < q$.

Thus, the steady state probability distribution is $\pi_k = (1-\alpha)\alpha^k$. This means the Markov chain is positive recurrent when $\alpha < 1$, and is ergodic.

This chain is aperiodic (state 0 has a self-transition, so it has period 1) and has a single communicating class. It is also easy to see that the mean revisit time for state 0 is finite, so the states are positive recurrent, and the chain will be ergodic.

Note that the probability balance equations are the same when $p = q$. However, in this case, we have

$$\pi_k = \pi_0, \ k = 1, 2, \ldots.$$

For this case, the normalization property yields

$$\sum_{k=0}^{\infty} \pi_0 = 1.$$

This equation has no solution, and thus the Markov chain is not positive recurrent and is not ergodic. Similar contradictions can be found for $p > 1 - p$.

Assume $\alpha < 1$. Then, the ergodic distribution is $\pi_k = (1-\alpha)\alpha^k, \ k = 0, 1, \ldots$. Can we compute $\mathbb{E}[X_\infty]$, the expected value of the state of the Markov chain in the limit?

Since X_∞ is a discrete random variable with PMF $P_{X_\infty}(k) = (1-\alpha)\alpha^k$, the expectation is

$$\begin{aligned}\mathbb{E}[X_\infty] &= \sum_{k=0}^{\infty} k(1-\alpha)\alpha^k = \alpha(1-\alpha)\sum_{k=1}^{\infty} k\alpha^{k-1} \\ &= \alpha(1-\alpha)\sum_{k=1}^{\infty}\frac{d}{d\alpha}\alpha^k = \alpha(1-\alpha)\frac{d}{d\alpha}\left(\sum_{k=1}^{\infty}\alpha^k\right) \\ &= \alpha(1-\alpha)\frac{d}{d\alpha}\left(\frac{1}{1-\alpha}-1\right) \\ &= \frac{\alpha}{1-\alpha}\end{aligned}$$

where we have interchanged differentiation and summation because of the convergence of the geometric series when $\alpha < 1$. This implies that, as $\alpha \to 1$, the expected value of the state (the length of the queue) blows up and approaches ∞.

11.4 Ergodicity and the Strong Law of Large Numbers

Markov chains were introduced by Andrey Markov and were named after him. He developed Markov chains to create correlated sequences of random variables, to study extensions of the strong Law of Large Numbers and the Central Limit Theorem for such sequences. In his first paper, in 1906, he proved that, for a Markov chain with positive transition probabilities, the average of the state values along a trajectory converges to the expected value of the limiting distribution (the fixed vector). This was an extension of the weak Law of Large Numbers. In later papers, he proved the Central Limit Theorem for such chains. Subsequently, he established that ergodic Markov chains have properties that generalize the Strong Law of Large Numbers.

Assuming a Markov chain $\{X_t\}$ is ergodic, the marginal distribution $P_{X_t}(x)$ converges to a limit distribution $\underline{\pi}$, where $\pi_i = \mathbb{P}[X_\infty = i]$. Then, for any bounded real-valued function $f : \mathcal{R}_X \to \Re$, we have

$$\lim_{n\to\infty}\frac{1}{n+1}\sum_{j=0}^{n} f(X_j) = \sum_{k\in\mathcal{R}_X} f(k)\pi_k = \mathbb{E}[f(X_\infty)]$$

almost surely. If we choose the function $f(k) = 1, f(j) = 0$ if $j \neq k$, we get the following statement:

$$\lim_{n\to\infty}\frac{1}{n+1}\sum_{j=0}^{n}\mathcal{I}\{X_j = k\} = \pi_k.$$

Hence, π_k is the fraction of time, on average, that the Markov chain spends in state k. If we choose the function $f(k) = k$, we get exactly the strong Law of Large Numbers, although we have to show this using a limiting argument when the number of states is infinite.

What is the key insight behind Markov's results? The Markov property of Markov chains established that the evolution of the process starting from a particular state k was independent of the past trajectory of the process. If state k was positive recurrent, the trajectories of states visited between visits to state k represented an independent sample of such possible trajectories. Defining as a random variable the sum of the function $f(X_j)$ over the number of states visited starting from state k before the next return to k (including the state k), every revisit provided independent, identically distributed random samples for $f(X_j)$. There is a subtle argument needed to handle the fact that each of those restarts might take different times in returning to state k, but again those random times are identically distributed. The results then follow from the strong Law of Large Numbers.

11.5 Transient Analysis of Markov Chains

Let $\{X_t\}$ be a homogeneous, discrete-time Markov chain with transition probability kernel $\mathbf{P}$, taking values in a discrete state space $\mathcal{R}_X$. Suppose we have a subset of states $A \subset \mathcal{R}_X$. Denote the trajectory of the

Markov chain for a specific outcome as $\{X_t(\omega)\}$. The *first hitting time* of the subset A starting from a state $X_0(s) = i$ is a random variable defined as:

$$H_i^A(\omega) = \inf\{n \geq 0 : X_n(\omega) \in A | X_0(\omega) = i\}.$$

H_i^A is a random variable, although we must allow for the possibility that it takes on an infinite value. Thus, it is a random variable taking values in $\Re \cup \{\infty\}$, a generalization of our earlier definitions. If $H_i^A(\omega) = \infty$, it means the process trajectory, for the experiment outcome ω, never reaches any of the states in A. The probability that the process hits A at all when it starts at state $X_0(\omega) = i$ is given by:

$$h_i^A = \mathbb{P}[\{H_i^A(\omega) < \infty\}].$$

In many problems of interest, we want to compute expected hitting times and hitting probabilities given a particular initial state X_0. Such hitting times can indicate successful completion of events and reaching of milestones. What is surprising is that we will be able to do these computations using simple linear algebra techniques, as described below.

Example 11.15
Let's first consider an example. Suppose we have a four state Markov chain, with transition probability matrix P given by:

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1/2 & 0 & 1/2 & 0 \\ 0 & 1/2 & 0 & 1/2 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Note that this system has three communicating classes: 1, 4 and $\{2, 3\}$. However, only 1 and 4 are recurrent classes. Once the state reaches states 1 or 4, the state trajectory stays in those states for all future times.

Suppose we start in state 2. We would like to compute the expected number of steps required to reach states 1 or 4. We can compute this as follows: Let k_i denote the expected time to reach states 1 or 4 starting from from state i. Then, observe the following relationships:

$$k_1 = 0; k_4 = 0$$

What about k_2 and k_3? By the Markov nature of the process, the expected time to reach from state 2 is 1 plus the expected time to reach from whatever the next state is, weighted by the probability of transitioning to that state. In mathematical terms, this yields

$$k_2 = 1 + 0.5k_1 + 0.5k_3; \quad k_3 = 1 + 0.5k_2 + 0.5k_4$$

Basically, any trajectory that starts at i and hits the set $A = \{0, 4\}$ has to take the first step to a state that is connected to i. From that next state, by time invariance, the expected hitting time is the same as that of trajectories that start at that state.

These last two equations are easily solved once we substitute $k_1 = 0, k_4 = 0$ to get $k_2 = k_3 = 2$.

What about a hitting probability? Let the set $A = \{4\}$. Then, reasoning along the same lines, the probability of hitting A from a particular state k is the weightednsum of the probabilities of hitting A from whatever states k transitions, 2, weighted by the transition probabilities. In mathematical terms,

$$h_4^A = 1; \; h_3^A = 0.5h_2^A + 0.5h_4^A; \; h_2^A = 0.5h_1^A + 0.5h_4^A; \; h_1^A = 0$$

Solving these, we get $h_1^A = 0, h_2^A = 1/3, h_3^A = 2/3, h_4^A = 1$.

Can we generalize the insights from this example to arbitrary Markov chains? Let's first focus on Markov chains with fininte state space $\mathcal{R}_X$. The result below characterizes the general solution:

Theorem 11.11
Let $\underline{h}^A$ denote the vector of hitting probabilities for a subset A of the finite state space $\mathcal{R}_X$. Then, $\underline{h}^A$ is the smallest non-negative solution of the following set of linear equations:

$$\begin{cases} h_i^A = 1 & i \in A \\ h_i^A = \sum_j \mathbf{P}_{ij} h_j^A & i \notin A \end{cases}$$

In vector form,

$$\underline{h}^A = \hat{\mathbf{P}}\underline{h}^A; \quad h_i^A = 1, i \in A,$$

where $\hat{\mathbf{P}}$ is a reduced version of matrix $\mathbf{P}$ with the rows corresponding to $i \in A$ deleted. By smallest solution we mean that, if $\underline{x}$ is another non-negative solution, then $x_i \geq h_i^A$.

Note that we have the same number of equations and unknowns, as there is one equation for each $i \notin A$.

Let's prove the above. First, let's show that $\underline{h}^A$ satisfies the equations. Assume $x(0) = i \in A$. Then, the hitting time $H_i^A = 0$, and the hitting probability $h_i^A = 1$, which the theorem guarantees by construction. Now, assume that $x(0) = i \notin A$. Then, $H_i^A > 0$, as it will take at least one step to reach a state in A. By the Markov property of the process,

$$h_i^A = \mathbb{P}[H_i^A < \infty | X_0 = i] = \sum_j \mathbb{P}[H_i^A < \infty, X_1 = j | X_0 = i] = \sum_j \mathbb{P}[H_i^A < \infty | X_1 = j]\mathbf{P}_i j = \sum_j h_j^A P_{ij}$$

which shows that $\underline{h}^A$ satisfies theorem 11.11.

Now, suppose we have a non-negative solution $\underline{g}$ to the equations in theorem 11.11. We want to show that these must be greater than or equal to the expected hitting times. We know that $h_i^A = g_i$ for $i \in A$, as they are set to 1. Suppose $i \notin A$. Then,

$$g_i = \sum_j \mathbf{P}_{ij} g_j = \sum_{j \in A} \mathbf{P}_{ij} g_j + \sum_{j \notin A} \mathbf{P}_{ij} g_j = \sum_{j \in A} \mathbf{P}_{ij} + \sum_{j \notin A} \mathbf{P}_{ij} g_j$$

Now, substitute for g_j in the last term, to get:

$$g_i = \sum_{j \in A} \mathbf{P}_{ij} + \sum_{j \notin A} \mathbf{P}_{ij} (\sum_{k \in A} \mathbf{P}_{jk} + \sum_{k \notin A} \mathbf{P}_{jk} g_k) = \mathbb{P}[X_1 \in A] + \mathbb{P}[X_1 \notin A, X_2 \in A] + \sum_{j,k \notin A} \mathbf{P}_{ij}\mathbf{P}_{jk} g_k.$$

By repeated substitution, we get

$$\begin{aligned} g_i = \mathbb{P}[X_1 \in A] + \mathbb{P}[X_1 \notin A, X_2 \in A] + \mathbb{P}[X_1, X_2 \notin A, X_3 \in A] + \ldots + \\ \mathbb{P}[X_1, \ldots, X_{n-1} \notin A, X_n) \in A] + \sum_{j_1, \ldots, j_n \notin A} \mathbf{P}_{ij_1}\mathbf{P}_{j_1 j_2} \cdots \mathbf{P}_{j_{n-1} j_n} g_{j_n} \end{aligned}$$

Note that the sum of all but the last term are $\mathbb{P}[H_i^A \leq n]$. Thus, $g_i \geq \mathbb{P}[H_i^A \leq n]$ for any n, because the last term is non-negative ($g_k \geq 0$ for all k). Thus,

$$g_i \geq \lim_{n \to \infty} \mathbb{P}[H_i^A \leq n] = \mathbb{P}[H_i^A < \infty] = h_i$$

which shows that $\underline{h}$ is the smallest nonnegative solution.

Can there be multiple solutions? Consider the Markov chain with transition probability matrix

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 1/2 & 1/2 \end{bmatrix}$$

and let $A = \{1\}$. Clearly, $h_1^A = 1$, and $h_2^A = h_3^A = 0$, because starting at either state 2 or 3, one cannot reach state 1 at all. Note, however that the equations in theorem 11.11 can be solved by any vector $\underline{g} = (1, k, k)^T$. Of course, the smallest nonnegative solution among this is $(1, 0, 0)^T$, which are the hitting times.

The above theorem is true even if the state space $\mathcal{S}$ is infinite. However, we now have an infinite number of equations to consider, which makes numerical computation harder.

Example 11.16

Consider a random walk on $\{0, 1, 2, \ldots\}$, where $\mathbf{P}_{00} = 1$, and $\mathbf{P}_{i(i+1)} = \mathbf{P}_{i(i-1)} = 1/2$ for $i \geq 1$, $\mathbf{P}_{ij} = 0, |i - j| \geq 2$. This corresponds to an infinite gambler's ruin problem where the gambler never leaves until he is broke. We would like to

compute the hitting probability for the set $A = \{0\}$, corresponding to the gambler leaving broke. Here are the relevant equations for the hitting probability h_i^A:

$$\begin{aligned} h_0^A &= 1 \\ h_1^A &= 0.5h_0^A + 0.5h_2^A \\ &\vdots \\ h_n^A &= 0.5h_{n-1}^A + 0.5h_{n+1}^A \qquad \vdots \end{aligned}$$

We can solve this via z-transforms, as follows: the characteristic equation of the recursion is

$$0.5z^2 - z + 0.5 = 0$$

. By inspection, this has a repeated root at $z = 0$. Thus, this admits solutions of the form $h_n^A = C + Dn$ for some constants C, D. To match the initial condition $h_0^A = 1$, we get $C = 1$. The second equation yields $C + D = 0.5 + 0.5(C + 2D)$, which is true for all D. Thus, any value of $D \geq 0$ will yield a valid nonnegative solution! However, h_n^A is a probability, and as such, it must be less than 1. Indeed, the only solution that will yield a probability is $D = 0$, so $h_n^A = 1$ for all n! This means that you will always go broke, no matter where you start!

What if we change the problem so that $P_{i(i+1)} = 3/4, P_{i(i-1)} = 1/4$? This is a very nice game, with odds in the players' favor. In this case, the main recursion yields

$$h_n^A = 0.25h_{n-1}^A + 0.75h_{n+1}^A$$

with characteristic equation $1 - 4z + 3z^2 = 0$, which yields solution of the form $h_n^A = C + D(1/3)^n$. To fit the initial condition $h_0^A = 1$, we have $C + D = 1$, or $D = 1 - C$. Thus, the general form of the solution is

$$h_n^A = (1 - C)(\frac{1}{3})^n + C = (\frac{1}{3})^n + C(1 - (\frac{1}{3})^n)$$

Note that, for any $C \geq 0$, this remains nonnegative.

Thus, we don't have an easy way to select C. Here is where the choice of smallest non-negative solution gives an answer: the smallest non-negative solution is given by $C = 0$, which is $h_n^A = (\frac{1}{3})^n$. In this case, the probability of going broke decreases exponentially with increasing initial condition.

Theorem 11.11 deals with hitting probabilities. We can develop a similar result for hitting times.

Theorem 11.12
Let $\underline{k}^A$ denote the vector of expected hitting times for a subset A of the state space $\mathcal{R}_X$, where these values could be infinite. Then, $\underline{k}^A$ is the smallest non-negative solution of the following set of linear equations:

$$\begin{cases} k_i^A = 0 & i \in A \\ k_i^A = 1 + \sum_{j \in \mathcal{R}_X} \mathbf{P}_{ij} k_j^A & i \notin A \end{cases}$$

In vector form, $\underline{k}^A = \underline{1} + \hat{\mathbf{P}}\underline{k}^A; k_i^A = 0, i \in A$, where $\hat{\mathbf{P}}$ is the state transition matrix $\mathbf{P}$ with the rows for $i \in A$ removed.

To show this, we proceed as before. We show that $\underline{k}^A$ satisfies the equations in theorem 11.12. If $X_0 = i \in A$, then $H_i^A = 0$, so $k_i^A = 0$. If $X_0 = i \notin A$, then $H_i^A \geq 1$. By the Markov property, when $i \notin A$,

$$\mathbb{P}[H_i^A = n | X_0 = i] = \sum_{j \in \mathcal{R}_X} \mathbb{P}[H_i^A = n, X_1 = j | X_0 = i] = \sum_{j \in \mathcal{R}_X} \mathbb{P}[H_i^A = n | X_1 = j]\mathbf{P}_{ij}$$

Thus,

$$\begin{aligned} k_i^A &= \sum_{n=1}^{\infty} n\mathbb{P}[H_i^A = n] + \infty\mathbb{P}[H_i^A = \infty] = \sum_{n=1}^{\infty} \mathbb{P}[H_i^A \geq n] \\ &= \sum_{n=1}^{\infty} \sum_{j\in\mathcal{S}} \mathbb{P}[H_i^A \geq n, X_1 = j) = \sum_{n=1}^{\infty} \sum_{j\in\mathcal{R}_X} \mathbb{P}[H_i^A \geq n | X_1 = j]\mathbf{P}_{ij} \\ &= \sum_{j\in\mathcal{R}_X} \mathbf{P}_{ij} \sum_{n=1}^{\infty} \mathbb{P}[H_i^A \geq n | x(1) = j] \\ &= \sum_{j\in\mathcal{R}_X} \mathbf{P}_{ij}(1 + \mathbb{E}[H_j^A]) = 1 + \sum_{j\in\mathcal{R}_X} \mathbf{P}_{ij} k_j^A \end{aligned}$$

which shows that the expected hitting times satisfy the equations of theorem 11.12, even when they have infinite value!

Let g be any solution of the linear equations in the Theorem. Then, $g_i = k_i^A = 0$ for $i \in A$. Suppose $i \notin A$. then,

$$\begin{aligned} g_i &= 1 + \sum_{j\notin A} \mathbf{P}_{ij} g_j \\ &= 1 + \sum_{j\notin A} \mathbf{P}_{ij} \Big(1 + \sum_{k\notin A} \mathbf{P}_{jk} g_k\Big) \\ &= \mathbb{P}[H_i^A \geq 1] + \mathbb{P}[H_i^A \geq 2] + \sum_{j,k\notin A} \mathbf{P}_{ij}\mathbf{P}_{jk} g_k \end{aligned}$$

Continuing the substitutions, we get

$$g_i = \mathbb{P}[H_i^A \geq 1] + \mathbb{P}[H_i^A \geq 2] + \cdots + \mathbb{P}[H_i^A \geq n] + \sum_{j_1,\ldots,j_n\notin A} \mathbf{P}_{ij_1}\mathbf{P}_{j_1 j_2}\cdots\mathbf{P}_{j_{n-1}j_n} g_{j_n}$$

Noting that $g_j \geq 0$, we have

$$g_i \geq \lim_{n\to\infty} \big(\mathbb{P}[H_i^A \geq 1] + \mathbb{P}[H_i^A \geq 2] + \cdots + \mathbb{P}[H_i^A \geq n]\big) = \mathbb{E}[H_i^A] = k_i^A$$

which shows that $\underline{k}^A$ is the smallest nonnegative solution.

Example 11.17

Consider the previous example 11.16, where we set $\mathbf{P}_{i(i+1)} = 1/4, \mathbf{P}_{i(i-1)} = 3/4$. Note that, in average, we are headed towards 0. We want to compute the expected time to reach state 0 from any state n. The relevant equations from theorem 11.12 are:

$$\begin{aligned} k_0^0 &= 0; \\ k_1^0 &= 1 + 0.75k_0^0 + 0.25k_2^0 \\ &\vdots \\ k_n^0 &= 1 + 0.75k_{n-1}^0 + 0.25k_{n+1}^0 \\ &\vdots \end{aligned}$$

Note that this set of linear equations has an input which is a constant on the right hand side, corresponding to a pole at $z = 1$. Furthermore, the characteristic equation for this system is $(z-1)(z-3) = 0$, so the pole at $z = 1$ is repeated. This means the solution is of the form

$$k_n^0 = Kn + A + B3^n.$$

Substituting into the above equations yields $K = 2$. The initial condition $k_0^0 = 0$ means $A = -B$. Note that $B \geq 0$ is required for the solution to stay non-negative. The smallest non-negative solution is $B = 0$, which yields $k_n^0 = 2n$.

Example 11.18
Here is a much more complex example. Consider a Tennis game, where the server's probability of winning a point is p, and the receiver's probability of winning a point is $1-p$. We assume that each point is an independent event, and that the probability of winning the point by Player 1 is the same no matter what the score. We can view the evolution of the score of the game as a Markov chain, which eventually ends in either Player 1, the server, winning the game, or Player 2, the receiver, winning the game. A state transition diagram using 17 states is show in Figure 11.16(a), where the tennis score is shown in the circle, and the state number is outside. The red transitions indicate points won by Player 2, and the black transitions indicate points won by Player 1. We have shortened the states somewhat by matching the 30-30 score and Deuce into the same state node, requiring two consecutive points to win by any player.

Analysis of this Markov chain shows that there are only two recurrent communicating classes: state 17 where Player 1 wins, and state

Can we compute the expected duration of the game as a function of p, the probability that Player 1 wins a point? At first, that seems like a daunting task given the size of the network. However, we can solve for this in stages. Conditioned on starting in the Deuce state, corresponding to state number 12, what is the expected number of games? We can solve this by analyzing the much smaller chain in Figure 11.16(b). Indeed, the expected exit time equations for the exit states $16, 17$ are:

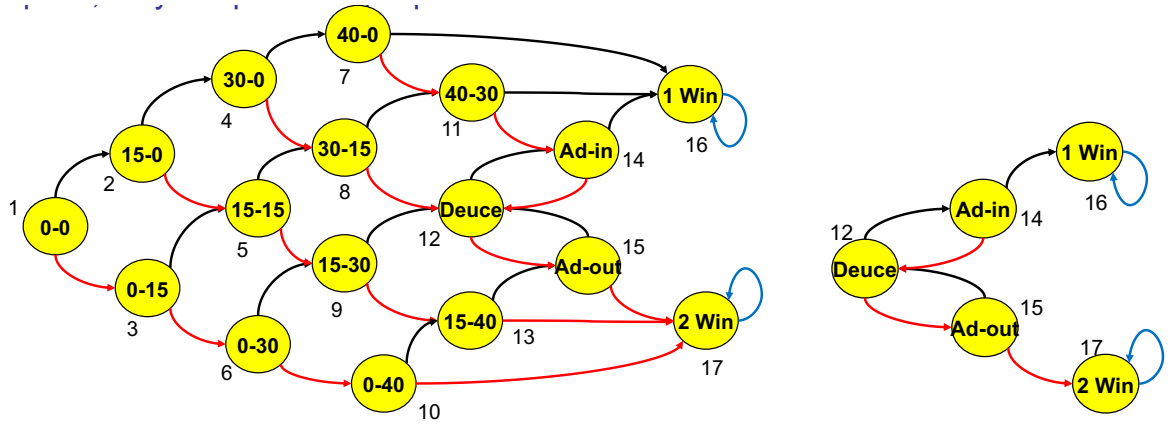

(a) The full Tennis Markov chain. (b) The chain starting at Deuce.

Figure 11.16: Diagram of the Markov chain for Example 11.18.

$$k_{12} = 1 + pk_{14} + (1-p)k_{15}; \quad k_{14} = 1 + (1-p)k_{12}; \quad k_{15} = 1 + pk_{12};$$

Substituting the last two equations into the first one yields the solution:

$$\begin{aligned} k_{12} &= 1 + p(1 + (1-p)k_{12}) + (1-p)(1 + pk_{12}) \\ &= 2 + 2p(1-p)k_{12} \\ \Rightarrow k_{12} &= \frac{2}{1-2p(1-p)}, \; k_{14} = \frac{1+2(1-p)^2}{1-2p(1-p)}, \; k_{15} = \frac{1+2p^2}{1-2p(1-p)} \end{aligned}$$

Let's ask a second question: what is the probability that Player 1 wins, given we have reached Deuce? This is an exit probability question on the same Markov chain, where we want the probability that the Markov chain will reach state 16. The relevant equations are:

$$\begin{aligned} &h_{17} = 0; \; h_{16} = 1; \; h_{14} = ph_{16} + (1-p)h_{12} \\ &h_{12} = ph_{14} + (1-p)h_{15}; \; h_{15} = ph_{12} + (1-p)h_{17} \end{aligned}$$

Solving these yields the following:

$$h_{12} = p(p + (1-p)h_{12}) + (1-p)ph_{12} \Rightarrow h_{12} = \frac{p^2}{1-2p(1-p)}.$$

Now that we have solved for these, note that we can compute the exit times for any of the other states in the full Markov chain by back substitution! For instance, the following equations propagate the solution two layers back.

$$k_{11} = 1 + (1-p)k_{14}; \qquad k_{13} = 1 + pk_{15}$$
$$k_7 = 1 + (1-p)k_{11};\ k_8 = 1 + pk_{11} + (1-p)k_{12};$$
$$k_9 = 1 + pk_{12} + (1-p)k_{13};\ k_{10} = 1 + pk_{13}$$

It is straightforward to write the remaining equations, until we compute $k_1 = pk_2 + (1-p)k_3$, yielding the expected number of games to play.

11.6 Applications

In this section, we discuss two popular applications of the theory of Markov chains.

11.6.1 Google PageRank algorithm

Larry Page and Sergey Brin developed PageRank at Stanford University in 1996 as part of a research project about a new kind of search engine. Sergey Brin had the idea that information on the web could be ordered in a hierarchy by "link popularity": a page ranks higher as there are more links to it. Shortly after, Page and Brin founded Google Inc., the company behind the Google search engine.

The PageRank algorithm outputs a probability distribution used to represent the likelihood that a person randomly clicking on links will arrive at any particular page. Google recalculates PageRank scores each time it crawls the Web and rebuilds its index. The formula uses a model of a random surfer who reaches their target site after several clicks, then switches to a random page. The PageRank value of a page reflects the chance that the random surfer will land on that page by clicking on a link.

The PageRank algorithm can best be modeled as a Markov chain in which the states are pages. Let j denote the state corresponding to a random surfer being in page k. The probability of transitioning to another page k is zero, unless there is a link on page j to page k. Then, the probability is uniform among the number of outgoing links to different pages out of j:

$$\mathbf{P}_{jk} = \begin{cases} 0 & \text{if there is no link to page } k \text{ on page } j, \\ \frac{1}{\text{Number distinct page links on page } j} & \text{elsewhere.} \end{cases}$$

What types of Markov chain does such a construction yield? First of all, it is a large one, with nearly a billion states. Second, the Markov chain is sparse, so that the number of transitions out of every row is a very small fraction of the number of nodes: hence, it is ideally viewed in terms of a graph. However, it is unclear that the resulting chain is irreducible. If a page has no links to other pages, it becomes a sink and there are no transitions out of it. Hence, there can be many transient states that have transitions that lead to such sinks. Thus, the resulting Markov chain is not ergodic.

One idea for making it ergodic is to allow the chain to transition to a random page, uniformly over all pages, when it reaches a page with no links. That would guarantee that the Markov chain would be not have any absorbing states, and that it would be aperiodic, as self-transitions would be possible, and it would even guarantee irreducibility. However, this would yield a hard Markov chain to analyze, as it would lose all the sparsity that was present in the original chain.

What Google's founders did was simpler and more clever: In addition to having probability of transitioning to any of the outgoing links in a page, they added a probability that they would transition from any page to any other page, uniformly. That is, let N be the total number of pages. Let $\alpha \in (0,1)$ be a relaxation factor. Then, the new transition probability was

$$\mathbf{P}_{ij}^{new} = \alpha * \mathbf{P}_{ij} + \frac{1-\alpha}{N}$$

Note that this guarantees that

$$\sum_{j=1}^{N} \mathbf{P}_{ij}^{new} = \alpha \sum_{j=1}^{N} \mathbf{P}_{ij} + \sum_{j=1}^{N} \frac{1-\alpha}{N} = \alpha + 1 - \alpha = 1.$$

Since $\mathbf{P}_{ij}^{new} \in (0,1)$, then $\mathbf{P}_{ij}^{new}$ is a valid stochastic matrix, and a state transition matrix for the new Markov chain. Furthermore, since $\mathbf{P}_{ij}^{new} > 0$, this chain is ergodic.

Google PageRank computes the stationary distribution of this Markov chain $\underline{\pi}$, and ranks pages in order of decreasing π_i. In principle, π_i is equal to the fraction of time that a random web surfer would spend on particular pages. However, solving for the eigenvector of a matrix of size $10^9 \times 10^9$ seems like a daunting task.

In this regard, the idea of adding the uniform transition probability makes this computation easier. Specifically, we can start with $\underline{p}(0) = \begin{bmatrix} \frac{1}{N} \\ \vdots \\ \frac{1}{N} \end{bmatrix}$. Then, we can compute

$$\underline{p}(t+1) = (\mathbf{P}^{new})^T \underline{p}(t).$$

In coordinates, this update is

$$\pi_j(t+1) = \sum_{i=1}^{N} \alpha \mathbf{P}_{ij} \pi_i(t) + \frac{1-\alpha)}{N} \pi_j(t)$$

Note that this is a very sparse update, so that computing an update iteration is of order $O(N)$, linear in the number of nodes. However, how many iterations are required? The rate of convergence of the iteration to steady state depends on the magnitude of the second largest eigenvalue of $\mathbf{P}^{new}$. Fortunately, that magnitude is no larger than α, so by selecting α, one can control the number of iterations. In practice, α is selected to be around 0.85, and the number of iterations required to converge is around 60.

11.6.2 Consensus Algorithms

Consider the following situation: a group of persons in a room generate estimates of a quantity X. Each person generates an estimate X_i. Each person shares their estimate with their immediate neighbors; each person then revised their estimate using a weighted linear combination of their own estimate and the estimate of their neighbors. Following this, another round of communication and averaging takes place. If we repeat this for many rounds, will ever person's estimate converge to the same estimate? Furthermore, if they converge, what estimate will they converge to?

While this problem seems a bit artificial in its description, the problem is at the heart of many applications: distributed training of deep neural networks where each agent only has part of the training data, formation flight of aircraft or birds, distributed control of robots, and similar problems.

Let's formulate this as a Markov chain problem. Assume there are K persons, and each person is represented by a node i. We assume that person i has n_i neighbors, denoted by a set N_i. For every node i and node $j \in N_i$, we assume there is an arc from i to j, and an arc from j to i. We assume the graph is connected, so that there is a path between every pair of nodes.

Let's define the update algorithm for node i. Denote by $X_i(n)$ the estimate of person i after the $n-th$ round of exchanges is complete. $X_i(0)$ is the initial estimate. Then,

$$X_i(n+1) = a_i X_i(n) + \frac{1-a_i}{n_i} \sum_{j \in N_i} X_j(n),$$

where $a_i \in (0,1)$. Each person i can have their own weight for their own estimate relative to that of their neighbors.

Writing this as a vector recursion, this is stated as:

$$\underline{X}(n+1) = \mathbf{P}\underline{X}(n)$$

where $\mathbf{P}_{ij} \in [0,1], \sum_{j=1}^{K} \mathbf{P}_{ij} = 1$. Hence, $\mathbf{P}$ is a stochastic matrix, and thus is the state transition matrix of a Markov chain. Furthermore, since we assume there is a path in the graph between every pair of nodes i, j, the Markov chain is irreducible, and it is aperiodic because there are self-loops, since $a_i > 0$. Thus, the Markov chain is ergodic, and $\mathbf{P}$ has a unique eigenvector corresponding to the eigenvalue 1. Indeed, since all the rows sum up to 1, we know that the eigenvector corresponding to the eigenvalue 1 is the vector of all ones.

This means that the estimates will converge:

$$\lim_{n\to\infty} \underline{X}(n+1) = C \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}$$

for some constant C. Note that the convergence to consensus will happen independent of the numerical choices used to average the neighbors' estimates. Convergence is inevitable because of the ergodicity of the underlying Markov chain.

However, what will be the limit of the estimates that the persons converge to? That depends on the averaging parameters we choose. Denote the stationary distribution of the ergodic Markov chain with state transition matrix $\mathbf{P}$ as $\underline{\pi}$. Then, the consensus algorithm will converge to $\underline{\pi}^t \underline{X}(0)$, the average of the initial estimates using the stationary probability distribution of the Markov chain.

To establish this, define $\underline{1}$ to be the K-dimensional vector of all ones. Then, since $\underline{X}(n)$ converges to the consensus value $C\underline{1}$, we have that

$$\lim_{n\to\infty} \frac{1}{K}\underline{1}^T \underline{X}(n) = \lim_{n\to\infty} \frac{1}{K}\underline{1}^T \mathbf{P}^n \underline{X}(0) = C$$

Note also that, since the Markov chain is ergodic and $\frac{1}{N}\underline{1}$ is a probability distribution,

$$l \lim_{n\to\infty} \frac{1}{K}\underline{1}^T \mathbf{P}^n = \underline{\pi}^T.$$

This establishes that the average value $C = \underline{\pi}^T \underline{X}(0)$.