Part VIII

Discrete-Time Markov Chains

This final part of the book is devoted to the topic of Markov chains. Markov chains are an extremely powerful tool used to model problems in computer science, statistics, physics, biology, and business – you name it! They are used extensively in AI/machine learning, computer science theory, and in all areas of computer system modeling (analysis of networking protocols, memory management protocols, server performance, capacity provisioning, disk protocols, etc.). Markov chains are also very common in operations research, including supply chain, call center, and inventory management.

Our goal in discussing Markov chains is two-fold. On the one hand, as always, we are interested in applications and particularly applications to computing. On the other hand, Markov chains are a core area of probability theory and thus we have chosen to cover the theory of Markov chains in some depth here.

In Chapter 24, we introduce finite-state Markov chains, limiting distributions, and stationary distributions.

In Chapter 25, we delve into the theory of finite-state Markov chains, discussing whether the limiting distribution exists and whether the stationary distribution is unique. We also introduce time reversibility, time averages, and mean passage times. A more elementary class might choose to skip this chapter, but it is my experience that undergraduates are fully capable of understanding this material if they proceed slowly and focus on examples to help illustrate the concepts.

In Chapter 26, we turn to infinite-state Markov chains. These are great for modeling the number of packets queued at a router, or the number of jobs at a data center. Although we skip the hardest proofs here, there is still a lot of intuition to be gained just in understanding definitions like transient and positive-recurrent.

All these chapters are full of examples of the application of Markov chains for modeling and solving problems. However, it is the final chapter, Chapter 27 on queueing theory, which really ties it all together. Through queueing theory, we see a real-world application of all the abstract concepts introduced in the Markov chain chapters.

24 Discrete-Time Markov Chains: Finite-State

This chapter begins our study of Markov chains, specifically discrete-time Markov chains. In this chapter and the next, we limit our discussion to Markov chains with a finite number of states. Our focus in this chapter will be on understanding how to obtain the limiting distribution for a Markov chain.

Markov chains come up in almost every field. As we study Markov chains, be on the lookout for Markov chains in your own work and the world around you. They are everywhere!

24.1 Our First Discrete-Time Markov Chain

Love is complicated. Figure 24.1 depicts the day-by-day relationship status of CMU students.

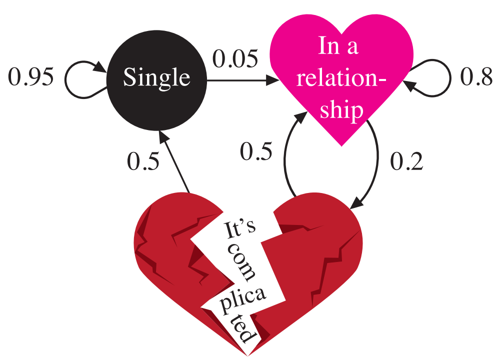

Figure 24.1 *The states of love, according to Facebook.*

There are three possible states for the relationship status. We assume that the relationship status can change only at the end of each day, according to the probabilities shown. For example, if we're "single" today, with probability 0.95 we will still be single tomorrow. When entering the "relationship" state, we stay there on average for five days (note the Geometric distribution), after which we move into the "it's complicated" state. From the "it's complicated" state, we're equally likely to return to the single state or the relationship state.

For such a Markov chain, we will ask questions like: *What fraction of time does one spend in the "relationship" state, as opposed to the "single" state?*

24.2 Formal Definition of a DTMC

Definition 24.1 *A* **discrete-time Markov chain (DTMC)** *is a stochastic process* $\{X_n, n = 0, 1, 2, \ldots\}$*, where* X_n *denotes the state at (discrete) time step* n *and such that* $\forall n \geq 0$*,* $\forall i, j$*, and* $\forall i_0, \ldots, i_{n-1} \in \mathbb{Z}$*,*

$$\begin{aligned} &\mathbf{P}\left\{X_{n+1} = j \;\middle|\; X_n = i, X_{n-1} = i_{n-1}, \ldots, X_0 = i_0\right\} \\ &= \mathbf{P}\left\{X_{n+1} = j \;\middle|\; X_n = i\right\} \qquad \textit{(Markovian property)} \\ &= P_{ij} \qquad \textit{(stationary property)}, \end{aligned}$$

where P_{ij} *is independent of the time step and of past history.*

Let's try to understand this definition line-by-line.

Question: First, what is a "stochastic process"?

Answer: A **stochastic process** is simply a sequence of random variables. In the case of Markov chain, this is a sequence of the states at each time step.

Question: What is being stated in the equality in marked "Markovian property" in the definition?

Answer: In a nutshell, past states don't matter. Only the current state matters.

Definition 24.2 *The* **Markovian property** *states that the conditional distribution of any future state* X_{n+1}*, given past states* $X_0, X_1, \ldots, X_{n-1}$*, and the present state* X_n*, is independent of past states and depends only on the present state* X_n*.*

Question: What is being stated in the equality marked "stationary property" in the definition?

Answer: The **stationary property** indicates that the transition probability, P_{ij}, is independent of the time step, n.

Definition 24.3 *The* **transition probability matrix** *associated with any DTMC is a matrix,* **P***, whose* (i, j)*th entry,* P_{ij}*, represents the probability of moving to state* j *on the next transition, given that the current state is* i*.*

Observe that, by definition, $\sum_j P_{ij} = 1, \ \forall i$, because, given that the DTMC is in state i, it must next transition to some state j.

Finite state versus infinite state: This chapter and the next will focus on DTMCs with a finite number of states, M. In Chapter 26, we will generalize to DTMCs with an infinite (but still countable) number of states.

DTMCs versus CTMCs: In a DTMC, the state can only change at synchronized (discrete) time steps. This book focuses on DTMCs. In a continuous-time Markov chain (CTMC) the state can change at any moment of time. CTMCs are outside the scope of this book, but we refer the interested reader to [35].

Ergodicity issues: In working with Markov chains, we will often be trying to understand the "limiting probability" of being in one state as opposed to another (limiting probabilities will be defined very soon). In this chapter, we will *not* dwell on the question of whether such limiting probabilities exist (called *ergodicity* issues). Instead we simply assume that there exists some limiting probability of being in each state of the chain. We defer all discussion of ergodicity to Chapter 25.

The three Ms: Solving Markov chains typically requires solving large systems of simultaneous equations. We therefore recommend taking the time to familiarize yourself with tools like Matlab [52], Mathematica [80], or Maple [50].

24.3 Examples of Finite-State DTMCs

We start with a few examples of simple Markov chains to illustrate the key concepts.

24.3.1 Repair Facility Problem

A machine is either working or is in the repair center. If it is working today, then there is a 95% chance that it will be working tomorrow. If it is in the repair center today, then there is a 40% chance that it will be working tomorrow. We are interested in questions like, "What fraction of time does my machine spend in the repair shop?"

Question: Describe the DTMC for the repair facility problem.

Answer: There are two states, "Working" and "Broken," where "Broken" denotes that the machine is in repair. The transition probability matrix is

$$\mathbf{P} = \begin{array}{c c} & \begin{array}{cc} W & B \end{array} \\ \begin{array}{c} W \\ B \end{array} & \begin{bmatrix} 0.95 & 0.05 \\ 0.40 & 0.60 \end{bmatrix} \end{array}.$$

The Markov chain diagram is shown in Figure 24.2.

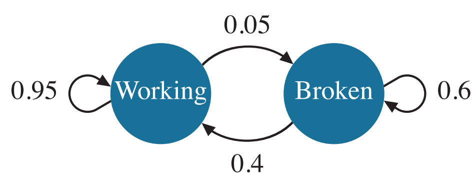

Figure 24.2 *Markov chain for the repair facility problem.*

Question: Now suppose that after the machine remains broken for four days, the machine is replaced with a new machine. How does the DTMC diagram change?

Answer: The revised DTMC is shown in Figure 24.3.

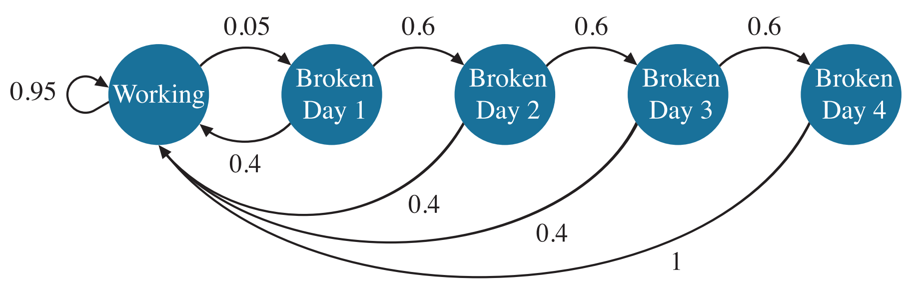

Figure 24.3 *Markov chain for the repair facility problem with a four-day limit.*

24.3.2 Umbrella Problem

An absent-minded professor has two umbrellas that she uses when commuting from home to office and back. If it rains and an umbrella is available in her location, she takes it. If it is not raining, she always forgets to take an umbrella. Suppose that it rains with probability p each time she commutes, independently of prior commutes. Our goal is to determine the fraction of commutes during which the professor gets wet.

Question: What is the state space?

Hint: Try to use as few states as possible!

Answer: We only need three states. The states track the number of umbrellas available at the current location, regardless of what this current location is. The DTMC is shown in Figure 24.4.

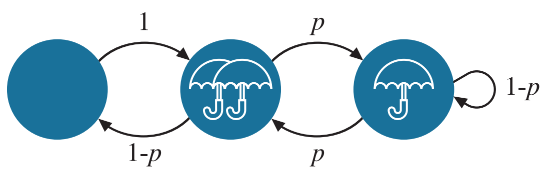

Figure 24.4 *DTMC for the umbrella problem.*

The transition probability matrix is $\mathbf{P} = \begin{array}{c} \\ 0 \\ 1 \\ 2 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} 0 & \quad 1 & \quad 2 \end{array} \\ \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1-p & p \\ 1-p & p & 0 \end{bmatrix} \end{array}.$

The probability of getting wet is the probability that it rains during a commute from a location with zero umbrellas.

24.3.3 Program Analysis Problem

A program has three types of instructions: CPU (C), memory (M), and user interaction (U). In analyzing the program, we note that a C instruction with probability 0.7 is followed by another C instruction, with probability 0.2 is followed by an M instruction and with probability 0.1 is followed by a U instruction. An M instruction with probability 0.1 is followed by another M instruction, with probability 0.8 is followed by a C instruction, and with probability 0.1 is followed by a U instruction. Finally, a U instruction with probability 0.9 is followed by a C instruction, and with probability 0.1 is followed by an M instruction.

In the exercises for this chapter and the next, we answer questions like, "What is the fraction of C instructions?" and "How many instructions are there on average between consecutive M instructions?" For now, we simply note that the program can be represented as a Markov chain with the transition probability matrix:

$$\mathbf{P} = \begin{array}{c} \\ C \\ M \\ U \end{array}\!\!\begin{array}{c} \begin{array}{ccc} C & \quad M & \quad U \end{array} \\ \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.8 & 0.1 & 0.1 \\ 0.9 & 0.1 & 0 \end{bmatrix} \end{array}.$$

24.4 Powers of $\mathbf{P}$: *n*-Step Transition Probabilities

Definition 24.4 *Let* $\mathbf{P}^n = \mathbf{P} \cdot \mathbf{P} \cdots \mathbf{P}$*, multiplied n times. Then* $(\mathbf{P}^n)_{ij}$ *denotes the* (i, j)*th entry of matrix* $\mathbf{P}^n$*. Occasionally, we will use the shorthand:*

$$P_{ij}^n \equiv (\mathbf{P}^n)_{ij}\,.$$

Question: What does $(\mathbf{P}^n)_{ij}$ represent?

Answer: To answer this, we first consider two examples.

Example 24.5 (Back to the umbrellas)

Consider the umbrella problem from before, where the chance of rain on any given day is $p = 0.4$. We then have:

$$\mathbf{P} = \begin{bmatrix} 0 & 0 & 1 \\ 0 & .6 & .4 \\ .6 & .4 & 0 \end{bmatrix} \qquad \mathbf{P}^5 = \begin{bmatrix} .06 & .30 & .64 \\ .18 & .38 & .44 \\ .38 & .44 & .18 \end{bmatrix} \qquad \mathbf{P}^{30} = \begin{bmatrix} .230 & .385 & .385 \\ .230 & .385 & .385 \\ .230 & .385 & .385 \end{bmatrix}.$$

Observe that all the rows become the *same*! Note also that, for all the above powers, each row sums to 1.

Example 24.6 (Back to the repair facility)

Now, consider again the simple repair facility problem, with general transition probability matrix $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} 1-a & a \\ b & 1-b \end{bmatrix}, \qquad 0 < a < 1,\ 0 < b < 1.$$

You should be able to prove by induction that

$$\mathbf{P}^n = \begin{bmatrix} \frac{b+a(1-a-b)^n}{a+b} & \frac{a-a(1-a-b)^n}{a+b} \\ \frac{b-b(1-a-b)^n}{a+b} & \frac{a+b(1-a-b)^n}{a+b} \end{bmatrix}$$

$$\lim_{n\to\infty} \mathbf{P}^n = \begin{bmatrix} \frac{b}{a+b} & \frac{a}{a+b} \\ \frac{b}{a+b} & \frac{a}{a+b} \end{bmatrix}.$$

Question: Again, all rows are the same. Why? What is the meaning of the row?

Hint: Consider a DTMC in state i. Suppose we want to know the probability that it will be in state j two steps from now. To go from state i to state j in two steps, the DTMC must have passed through some state k after the first step. Below we condition on this intermediate state k.

For an M-state DTMC, as shown in Figure 24.5,

$$\begin{aligned}\left(\mathbf{P}^2\right)_{ij} &= \sum_{k=0}^{M-1} P_{ik} \cdot P_{kj} \\ &= \text{Probability of being in state } j \text{ in two steps, given we're in state } i \text{ now.}\end{aligned}$$

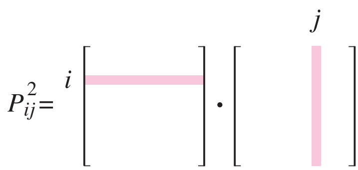

Figure 24.5 $\left(\mathbf{P}^2\right)_{ij}$.

Likewise, the n-wise product can be viewed by conditioning on the state k after $n-1$ time steps:

$$\begin{aligned}(\mathbf{P}^n)_{ij} &= \sum_{k=0}^{M-1} \left(\mathbf{P}^{n-1}\right)_{ik} P_{kj} \\ &= \text{Probability of being in state } j \text{ in } n \text{ steps, given we are in state } i \text{ now.}\end{aligned}$$

24.5 Limiting Probabilities

We now move on to looking at the limit. Consider the (i, j)th entry of the power matrix $\mathbf{P}^n$ for large n:

$$\lim_{n\to\infty} (\mathbf{P}^n)_{ij} \equiv \left(\lim_{n\to\infty} \mathbf{P}^n\right)_{ij}.$$

This quantity represents the limiting probability of being in state j infinitely far into the future, given that we started in state i.

Question: So what is the limiting probability of having zero umbrellas?

Answer: According to $\mathbf{P}^{30}$, it is 0.23.

Question: The fact that the rows of $\lim_{n\to\infty} \mathbf{P}^n$ are all the same is interesting because it says what?

Answer: The fact that $(\mathbf{P}^n)_{ij}$ is the same for all values of i says that the starting state, i, does not matter.

Definition 24.7 *Let*

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

π_j *represents the* **limiting probability** *that the chain is in state* j*, independent of the starting state i. For an* M*-state DTMC, with states* $0, 1, \ldots, M-1$*,*

$$\vec{\pi} = (\pi_0, \pi_1, \ldots, \pi_{M-1}), \quad \textit{where} \sum_{i=0}^{M-1} \pi_i = 1,$$

represents the **limiting distribution** *of being in each state.*

Important note: As defined, π_j is a limit. Yet it is not at all obvious that the limit π_j exists! It is also not obvious that $\vec{\pi}$ represents a distribution (that is, $\sum_i \pi_i = 1$), although this latter part turns out to be easy to see (Exercise 24.2). For the rest of this chapter, we will assume that the limiting probabilities exist. In Chapter 25 we look at the existence question in detail.

Question: So what is the limiting probability that the professor gets wet?

Answer: The professor gets wet if both (1) the state is 0, that is, there are zero umbrellas in the current location (π_0); and (2) it is raining ($p = 0.4$). So the limiting probability that the professor gets wet is $\pi_0 \cdot p = (0.23)(0.4) = 0.092$.

Question: Can you see why the limiting probability of having one umbrella is equal to the limiting probability of having two umbrellas?

Answer: Let's go back to Figure 24.4. Suppose now that we're only trying to determine the fraction of time that we're in a location with one umbrella versus the fraction of time that we're in a location with two umbrellas. In that case, all that matters is the number of visits to state 1 versus the number of visits to state 2. But, over a long period of time, the number of visits to state 1 and the number to state 2 are equal. To see this, if one considers only those two options of 1 and 2, then the chain from Figure 24.4 collapses to that shown in Figure 24.6. But the chain in Figure 24.6 is symmetric, hence the equal limiting probabilities.

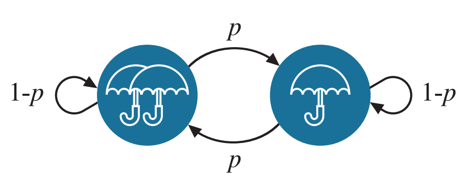

Figure 24.6 *Compressed umbrella problem.*

24.6 Stationary Equations

Question: Based only on what we have learned so far, how do we determine $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$?

Answer: We take the transition probability matrix $\mathbf{P}$ and raise it to the nth power for some large n and look at the jth column, any row.

Question: Multiplying $\mathbf{P}$ by itself many times sounds quite onerous. Also, it seems one might need to perform a very large number of multiplications if the Markov chain is large. Is there a more efficient way?

Answer: Yes, by solving stationary equations, given in Definition 24.8.

Definition 24.8 *A probability distribution* $\vec{\pi} = (\pi_0, \pi_1, \ldots, \pi_{M-1})$ *is said to be* **stationary** *for the Markov chain with transition matrix* $\mathbf{P}$ *if*

$$\vec{\pi} \cdot \mathbf{P} = \vec{\pi} \quad \textit{and} \quad \sum_{i=0}^{M-1} \pi_i = 1.$$

Figure 24.7 provides an illustration of $\vec{\pi} \cdot \mathbf{P} = \vec{\pi}$.

$$\begin{bmatrix} \pi_0 & \pi_1 & \pi_2 \end{bmatrix} \begin{bmatrix} P_{00} & P_{01} & P_{02} \\ P_{10} & P_{11} & P_{12} \\ P_{20} & P_{21} & P_{22} \end{bmatrix} = \begin{bmatrix} \pi_0 & \pi_1 & \pi_2 \end{bmatrix}$$

Figure 24.7 *Visualization of* $\vec{\pi} \cdot \mathbf{P} = \vec{\pi}$ *for the case of* $M = 3$ *states.*

Doing the row-by-column multiplication in Figure 24.7 results in the following **stationary equations**:

$$\begin{aligned} \pi_0 \cdot P_{00} + \pi_1 \cdot P_{10} + \pi_2 \cdot P_{20} &= \pi_0 \\ \pi_0 \cdot P_{01} + \pi_1 \cdot P_{11} + \pi_2 \cdot P_{21} &= \pi_1 \\ \pi_0 \cdot P_{02} + \pi_1 \cdot P_{12} + \pi_2 \cdot P_{22} &= \pi_2 \\ \pi_0 + \pi_1 + \pi_2 &= 1. \end{aligned}$$

These stationary equations can be written more compactly as follows:

$$\sum_{i=0}^{M-1} \pi_i P_{ij} = \pi_j, \ \forall j \quad \text{and} \quad \sum_{i=0}^{M-1} \pi_i = 1. \tag{24.1}$$

Question: What does the left-hand side of the first equation in (24.1) represent?

Answer: The left-hand side represents the probability of being in state j one transition from now, given that the current probability distribution on the states is $\vec{\pi}$. So (24.1) says that if we start out distributed according to $\vec{\pi}$, then one step later our probability of being in each state will still follow distribution $\vec{\pi}$. Thus, from then on we will always have the same probability distribution on the states. Hence, we call the distribution "stationary," which connotes the fact that we stay there forever.

24.7 The Stationary Distribution Equals the Limiting Distribution

The following theorem relates the *limiting distribution* to the *stationary distribution* for a finite-state DTMC. Specifically, the theorem says that for a finite-state DTMC, the stationary distribution obtained by solving (24.1) is unique and represents the limiting probabilities of being in each state, assuming these limiting probabilities exist.

Theorem 24.9 (Stationary distribution = limiting distribution) *In a finite-state DTMC with M states, let*

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

be the limiting probability of being in state j (independent of the starting state i) and let

$$\vec{\pi} = (\pi_0, \pi_1, \ldots, \pi_{M-1}), \quad \textit{where} \sum_{i=0}^{M-1} \pi_i = 1,$$

be the limiting distribution. Assuming that $\vec{\pi}$ exists, then $\vec{\pi}$ is also a stationary distribution and no other stationary distribution exists.

Question: What's the intuition behind Theorem 24.9?

Answer: Intuitively, given that the limiting distribution, $\vec{\pi}$, exists, it makes sense that this limiting distribution should be stationary, because we're not leaving the limit once we get there. It's not as immediately obvious that this limiting distribution should be the only stationary distribution.

Question: What's the impact of Theorem 24.9?

Answer: Assuming that the limiting distribution exists, Theorem 24.9 tells us that to get the limiting distribution we don't need to raise the transition matrix to a high power, but rather we can just solve the stationary equations.

Proof: [Theorem 24.9] We prove two things about the limiting distribution $\vec{\pi}$:

1. We will prove that $\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots, \pi_{M-1})$ is a stationary distribution. Hence, at least one stationary distribution exists.
2. We will prove that any stationary distribution must be equal to the limiting distribution.

Important: Throughout the proof, $\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots, \pi_{M-1})$ is used to refer to the *limiting distribution*.

Part 1: Proof that $\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots, \pi_{M-1})$ is a stationary distribution:

Intuitively, this should make a lot of sense. If we have some limiting distribution, then once you get there, you should stay there forever.

$$\begin{aligned}
\pi_j = \lim_{n\to\infty} \left(\mathbf{P}^{n+1}\right)_{ij} &= \lim_{n\to\infty} \sum_{k=0}^{M-1} (\mathbf{P}^n)_{ik} \cdot P_{kj} \\
&= \sum_{k=0}^{M-1} \lim_{n\to\infty} (\mathbf{P}^n)_{ik} P_{kj} \\
&= \sum_{k=0}^{M-1} \pi_k P_{kj}.
\end{aligned}$$

Hence $\vec{\pi}$ satisfies the stationary equations, so it's also a stationary distribution.

Part 2: Proof that any stationary distribution, $\vec{\pi}'$, must equal the limiting distribution, $\vec{\pi}$:

Let $\vec{\pi}'$ be any stationary probability distribution. As usual, $\vec{\pi}$ represents the limiting probability distribution. We will prove that $\vec{\pi}' = \vec{\pi}$, and specifically that $\pi'_j = \pi_j, \forall j$.

Suppose we start at time 0 with stationary distribution $\vec{\pi}' = (\pi'_0, \pi'_1, \ldots, \pi'_{M-1})$. After one step, we will still be in distribution $\vec{\pi}'$:

$$\vec{\pi}' \cdot \mathbf{P} = \vec{\pi}'$$

But this implies that after n steps, we will still be in distribution $\vec{\pi}'$:

$$\vec{\pi}' \cdot \mathbf{P}^n = \vec{\pi}'. \tag{24.2}$$

Looking at the jth entry of $\vec{\pi}'$ in (24.2), we have:

$$\sum_{k=0}^{M-1} \pi'_k \left(\mathbf{P}^n\right)_{kj} = \pi'_j.$$

Taking the limit as n goes to infinity of both sides, we have:

$$\lim_{n\to\infty} \sum_{k=0}^{M-1} \pi'_k \left(\mathbf{P}^n\right)_{kj} = \lim_{n\to\infty} \pi'_j = \pi'_j.$$

We are now ready to prove that $\pi'_j = \pi_j, \forall j$:

$$\begin{aligned} \pi'_j = \lim_{n\to\infty} \sum_{k=0}^{M-1} \pi'_k \left(\mathbf{P}^n\right)_{kj} &= \sum_{k=0}^{M-1} \pi'_k \lim_{n\to\infty} \left(\mathbf{P}^n\right)_{kj} \\ &= \sum_{k=0}^{M-1} \pi'_k \pi_j \\ &= \pi_j \sum_{k=0}^{M-1} \pi'_k = \pi_j. \end{aligned}$$ ■

Note that we were allowed to pull the limit into the summation sign in both parts because we had finite sums (M is finite).

One more thing: In the literature you often see the phrase "consider a stationary Markov chain," or "consider the following Markov chain in steady state ..."

Definition 24.10 *A Markov chain for which the limiting probabilities exist is said to be* **stationary** *or in* **steady state** *if the initial state is chosen according to the stationary probabilities.*

Summary: Finding the limiting probabilities in a finite-state DTMC:

By Theorem 24.9, provided the limiting distribution $\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots, \pi_{M-1})$ exists, we can obtain it by solving the stationary equations:

$$\vec{\pi} \cdot \mathbf{P} = \vec{\pi} \quad \text{and} \quad \sum_{i=0}^{M-1} \pi_i = 1,$$

where $\vec{\pi} = (\pi_0, \pi_1, \ldots, \pi_{M-1})$.

24.8 Examples of Solving Stationary Equations

Example 24.11 (Repair facility problem with cost)

Consider again the repair facility problem represented by the finite-state DTMC shown again in Figure 24.8.

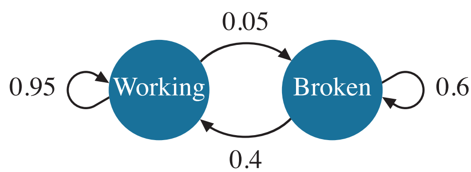

Figure 24.8 *Markov chain for the repair facility problem.*

We are interested in the following type of question.

Question: The help desk is trying to figure out how much to charge me for maintaining my machine. They figure that it costs them \$300 every day that my machine is in repair. What will be my annual repair bill?

To answer this question, we first derive the limiting distribution $\vec{\pi} = (\pi_W, \pi_B)$ for this chain. We solve the stationary equations to get $\vec{\pi}$ as follows:

$$\vec{\pi} = \vec{\pi} \cdot \mathbf{P}, \text{ where } \mathbf{P} = \begin{pmatrix} 0.95 & 0.05 \\ 0.4 & 0.6 \end{pmatrix}$$

$$\pi_W + \pi_B = 1.$$

This translates to the following equations:

$$\begin{aligned} \pi_W &= \pi_W \cdot 0.95 + \pi_B \cdot 0.4 \\ \pi_B &= \pi_W \cdot 0.05 + \pi_B \cdot 0.6 \\ \pi_W + \pi_B &= 1. \end{aligned}$$

Question: What do you notice about the first two equations above?

Answer: They are identical! In general, if $\vec{\pi} = \vec{\pi} \cdot \mathbf{P}$ results in M equations, only $M - 1$ of these will be linearly independent (this is because the rows of $\mathbf{P}$ all sum to 1). Fortunately, the last equation above (the normalization condition) is there to help us out. Solving, we get $\pi_W = \frac{8}{9}$ and $\pi_B = \frac{1}{9}$.

By Theorem 24.9, the stationary distribution also represents the limiting probability distribution. Thus my machine is broken one out of every nine days on average. The expected daily cost is $\frac{1}{9} \cdot 300 = \$33.33$ (with an annual cost of more than \$12,000).

Example 24.12 (Umbrella problem)

Consider again the umbrella problem depicted in Figure 24.9.

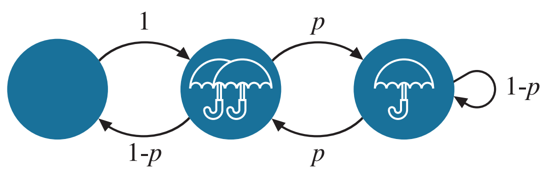

Figure 24.9 *DTMC for the umbrella problem.*

Rather than raising the transition matrix **P** to a high power, this time we use the stationary equations to obtain the limiting probabilities for general p:

$$\begin{aligned}
\pi_0 &= \pi_2 \cdot (1-p) \\
\pi_1 &= \pi_1 \cdot (1-p) + \pi_2 \cdot p \\
\pi_2 &= \pi_0 \cdot 1 + \pi_1 \cdot p \\
\pi_0 + \pi_1 + \pi_2 &= 1.
\end{aligned}$$

Their solution is

$$\pi_0 = \frac{1-p}{3-p} \qquad \pi_1 = \frac{1}{3-p} \qquad \pi_2 = \frac{1}{3-p}.$$

Question: Suppose the professor lives in Pittsburgh, where the daily probability of rain is $p = 0.6$. What fraction of days does the professor get soaked?

Answer: The professor gets soaked if she has zero umbrellas and it is raining: $\pi_0 \cdot p = \frac{0.4}{2.4} \cdot 0.6 = 0.1$. Not too bad. No wonder I never learn!

24.9 Exercises

24.1 **Solving for limiting distribution**
For the program analysis problem from Section 24.3.3, solve the stationary equations to determine the limiting distribution, (π_C, π_M, π_U).

24.2 **Powers of transition matrix**
Given any finite-state transition matrix, $\mathbf{P}$, prove that for any positive integer n, $\mathbf{P}^n$ maintains the property that each row sums to 1.

24.3 **Random walk on clique**
You are given a clique on $n > 1$ nodes (a clique is a graph where there is an edge between every pair of nodes). At every time step, you move to a uniformly random node *other* than the node you're in. You start at node v. Let T denote the time (number of hops) until you first return to v.
(a) What is $\mathbf{E}[T]$?
(b) What is $\mathbf{Var}(T)$?

24.4 **Card shuffling**
You have n distinct cards, arranged in an ordered list: $1, 2, 3, \ldots, n$. Every minute, you pick a card at random and move it to the front of the ordered list. We can model this process as a DTMC, where the state is the ordered list. Derive a stationary distribution for the DTMC. [Hint: Make a guess and then prove it.]

24.5 **Doubly stochastic matrix**
A doubly stochastic matrix is one in which the entries in each row sum up to 1, and the entries in each column sum up to 1. Suppose you have a finite-state Markov chain whose limiting probabilities exist and whose transition matrix is doubly stochastic. What can you prove about the stationary distribution of this Markov chain? [Hint: Start by writing some examples of doubly stochastic transition matrices.]

24.6 **Randomized chess**
In chess, a rook can move either horizontally within its row (left or right) or vertically within its column (up or down) any number of squares. Imagine a rook that starts at the lower left corner of an 8×8 chess board. At each move, a bored child decides to move the rook to a random legal location (assume that the "move" cannot involve staying still). Let T denote the time until the rook first lands in the upper right corner of the board. Compute $\mathbf{E}[T]$ and $\mathbf{Var}(T)$.

24.7 **Tennis match**
[Proposed by William Liu] Abinaya and Misha are playing tennis. They're currently tied at deuce, meaning that the next person to lead by two points wins the game. Suppose that Misha wins each point independently with probability $\frac{2}{3}$ (where Abinaya wins with probability $\frac{1}{3}$).
(a) What is the probability that Misha wins the game?
(b) What is the expected number of remaining points played until someone wins?

24.8 **Markovopoly**
[Proposed by Tai Yasuda] Suppose you are playing a board game where the board has 28 locations arranged as shown in Figure 24.10. You start at the "Go" square, and, at each turn, you roll a six-sided die and move forward in the clockwise direction whatever number you roll. However, the dark squares in the corners are jail states, and once you land there, you must sit out for the next three turns (for the next three turns, you stay in jail instead of rolling a die and moving). On the fourth turn, you can roll the die again and move. Your goal is to figure out the fraction of the turns that you are in jail. (You are "in jail" if you are in a jail square at the end of your turn.) Write stationary equations to determine this fraction.

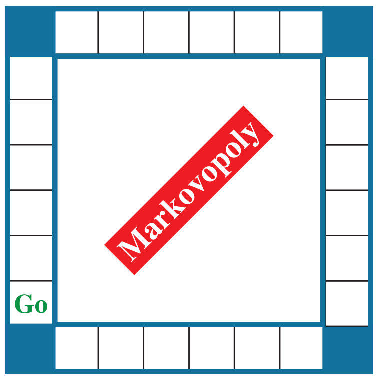

Figure 24.10 *Markovopoly for Exercise 24.8.*

24.9 **Axis & Allies**
In the game Axis & Allies, the outcome of a two-sided naval battle is decided by repeated rolling of dice. Until all ships on at least one side are destroyed, each side rolls one six-sided die for *each* of its existing ships. The die rolls determine casualties inflicted on the opponent; these casualties are removed from play and cannot fire (roll) in subsequent rounds.
There are two types of ships: battleships and destroyers. For a battleship, a die roll of four or lower is scored as a "hit" on the opponent. For a destroyer, a die roll of three or lower is scored as a "hit" on the opponent. It takes two hits (not necessarily in the same round) to destroy a battleship and only one hit to destroy a destroyer. (Note: Battleships are twice as expensive as destroyers.)
For example: Suppose side A has two destroyers and one battleship. Suppose side B has one destroyer and three battleships. Side A rolls two dice for its destroyers (rolling, say, 3 and 6) and one die for its battleship (rolling, say, 5). This means that side A generates one hit against side B.

At the same time, side B rolls one die for its destroyer (rolling, say 5) and three dice for its battleships (rolling, say, 1, 4, and 6). This means that side B generates two hits against side A.
The defender gets to decide to which ship to allocate the hit; we assume that the defender chooses intelligently. In the above example, side A will choose to be left with one destroyer and one weakened battleship. Side B will choose to be left with one destroyer, one weakened battleship and two undamaged battleships.
If two destroyers (side A) engage a battleship (side B) in a battle, what is the probability that the destroyers win? What is the probability that the battleship wins? [Hint: Raise a matrix to a large power.] [Note: A tie is also possible.]

24.10 **The SIR epidemic model**
The SIR model is commonly used to predict the spread of epidemic diseases. We have a population of n people. The state of the system is (n_S, n_I, n_R), where

- n_S is the number of people who are *susceptible* (healthy/uninfected);
- n_I is the number of people who are *infected*;
- n_R is the number of people who are *recovered*. In the SIR model, "recovered" includes both those recovered and deceased. The point is that "recovered" people are no longer susceptible to the disease.

Clearly $n_S + n_I + n_R = n$.
Each individual of the population independently follows this transmission model:

- If the individual is susceptible, then:
  - with probability $p \cdot \frac{n_I}{n}$, the individual will be infected tomorrow;
  - with probability $1 - p \cdot \frac{n_I}{n}$, the individual will stay susceptible tomorrow.
- If the individual is infected, then:
  - with probability $\frac{1}{21}$, the individual will be recovered;
  - with probability $\frac{20}{21}$, the individual will stay infected.
- If the individual is recovered, then with probability 1 the individual stays recovered.

The goal of the SIR model is to predict what *fraction* of people are in the "susceptible" state when the epidemic ends (that is, $n_I = 0$). These are the people who never got sick and thus have the potential to get sick if the disease resurfaces. You will determine this fraction as a function of the parameter p. You will do this by first determining the appropriate probability transition matrix and then raising this matrix to a very high power. For both steps you'll want to use a computer program like Matlab. For the sake of this problem, please assume $n = 3$ (but feel free to try out higher values of n as well).

(a) How many states are there in this system?
(b) How many absorbing states are there in this system, and what are they? Absorbing states are states that you never leave once you enter them. [Hint: What is n_I for an absorbing state?]
(c) Derive the transition probability from state $(2, 1, 0)$ to $(1, 1, 1)$. Be careful to think about all the ways that this transition can happen. Plug in the values of n_I and n and use $p = 0.5$ so that your final answer is a constant.
(d) Use a computer program to generate the entire transition matrix $\mathbf{P}$. Assume that $p = 0.5$. Print out the row corresponding to state $(2, 1, 0)$. Now raise $\mathbf{P}$ to some very high power and watch what happens to row $(2, 1, 0)$. You'll want a high enough power that most of your entries are smaller than 0.01. What is the meaning of the row corresponding to state $(2, 1, 0)$?
(e) The parameter p can be thought of as a social distancing parameter, where lower p represents better social distancing practices. Consider values of p between 0 and 1. For each value of p, determine the expected fraction of the population who are left in the susceptible state when the outbreak is over (you will do this by conditioning on the probability of ending up in each absorbing state). Assume that you start in state $(2, 1, 0)$. Your final output will be a graph with p on the x-axis, but you can alternatively create a table with values of p spaced out by 0.05.