26 Discrete-Time Markov Chains: Infinite-State

So far we have only talked about *finite*-state discrete-time Markov chains (DTMCs) with M states. Now we move on to infinite-state DTMCs. For a Markov chain with an infinite number of states, one can still imagine a transition probability matrix, $\mathbf{P}$, but the matrix has infinite dimension.

For an infinite-state DTMC, we denote the limiting probability distribution on the states by

$$\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots) \quad \text{where} \quad \pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij} \quad \text{and} \quad \sum_{j=0}^{\infty} \pi_j = 1.$$

We say that distribution $\vec{\pi}$ is stationary if

$$\pi_j = \sum_{k=0}^{\infty} \pi_k P_{kj} \quad \text{and} \quad \sum_{j=0}^{\infty} \pi_j = 1.$$

Infinite-state Markov chains are common in modeling systems where the number of customers or number of jobs is unbounded, and thus the state space is unbounded. The typical example is a queue of jobs or packets, where the queue can grow arbitrarily long.

This chapter will introduce infinite-state DTMCs. We will see that many of the definitions, solution techniques, and theorems from finite-state DTMCs carry over to infinite-state DTMCs. However, there is one crucial difference, which comes up in the definition of ergodicity and the existence of a limiting distribution. This difference will be discussed starting in Section 26.4.

26.1 Stationary = Limiting

We have seen that for a finite-state DTMC, if the limiting distribution exists, then the limiting distribution and stationary distribution are equivalent (Theorem 24.9). The same result holds for infinite-state DTMCs.

Theorem 26.1 (Stationary distribution = limiting distribution) *Given an infinite-state DTMC, let*

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

be the limiting probability of being in state j and let

$$\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots), \quad \textit{where} \quad \sum_{i=0}^{\infty} \pi_i = 1,$$

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

Proof: The proof follows along the lines of the proof of Theorem 24.9; however, it is a little more technical because we can't simply interchange the limit and the summation as we did in that proof, because we have an infinite sum over states. Fortunately, one can get around this difficulty by lower-bounding the infinite sum by a finite sum, which allows us to exchange the limit and the summation. After the exchange, we then consider the limit as the number of items in the finite sum approaches infinity. The details of this trickery are given in [35, section 8.9]. ■

26.2 Solving Stationary Equations in Infinite-State DTMCs

So we can obtain the limiting distribution, $\vec{\pi}$, by solving the stationary equations. Yet there are an infinite number of stationary equations! How do we solve them?

Consider an example of a router that has infinite capacity for packets, called an **unbounded queue** (Figure 26.1). Packets arrive at the router and queue up there. We think of the router as a "server" since it serves packets. The server processes the packet at the head of the queue, and when it finishes processing that packet, it moves on to the next packet.

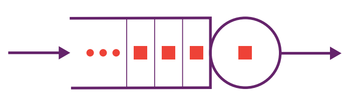

Figure 26.1 *Illustration of a server with unbounded buffer.*

Suppose at every time step, with probability $p = \frac{1}{4}$ one packet arrives, and independently, with probability $q = \frac{1}{3}$ one packet departs. Note that during a time step we might have both an arrival and a transmission, or neither. That is, a packet can "arrive" and "depart" within the same time step, leaving the system in the same state.

We will be interested in answering questions like: *What is the average number of packets in the system?*

To answer this question, we model the problem as a DTMC with an infinite number of states: 0, 1, 2, . . ., representing the number of packets at the router. Let $r = p(1-q) = \frac{1}{6}$ and $s = q(1-p) = \frac{1}{4}$, where $r < s$. Figure 26.2 shows the Markov chain for our problem.

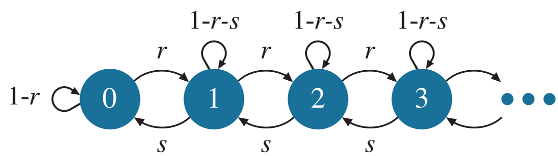

Figure 26.2 *DTMC for a server with unbounded queue.*

Here the transition probability matrix is infinite!

$$\mathbf{P} = \begin{pmatrix} 1-r & r & 0 & 0 & \cdots \\ s & 1-r-s & r & 0 & \cdots \\ 0 & s & 1-r-s & r & \cdots \\ 0 & 0 & s & 1-r-s & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \end{pmatrix}.$$

The stationary equations look like this:

$$\begin{aligned} \pi_0 &= \pi_0(1-r) + \pi_1 s \\ \pi_1 &= \pi_0 r + \pi_1(1-r-s) + \pi_2 s \\ \pi_2 &= \pi_1 r + \pi_2(1-r-s) + \pi_3 s \\ \pi_3 &= \pi_2 r + \pi_3(1-r-s) + \pi_4 s \\ &\vdots \end{aligned}$$

$$\pi_0 + \pi_1 + \pi_2 + \pi_3 + \cdots = 1.$$

Question: How are we going to solve this infinite number of equations?

Answer: It might be easier to write the time-reversibility equations (Theorem 25.29):

$$\begin{aligned} \pi_0 \cdot r &= \pi_1 \cdot s \\ \pi_1 \cdot r &= \pi_2 \cdot s \\ \pi_2 \cdot r &= \pi_3 \cdot s \\ &\vdots \end{aligned}$$

which yield

$$\pi_1 = \left(\frac{r}{s}\right) \cdot \pi_0$$
$$\pi_2 = \left(\frac{r}{s}\right) \cdot \pi_1 = \left(\frac{r}{s}\right)^2 \cdot \pi_0$$
$$\pi_3 = \left(\frac{r}{s}\right) \cdot \pi_2 = \left(\frac{r}{s}\right)^3 \cdot \pi_0$$
$$\vdots$$

We can now make a general "guess":

$$\pi_i = \left(\frac{r}{s}\right)^i \pi_0.$$

Question: How do we verify that this guess is correct?

Answer: To verify your guess, you need to show that it satisfies the stationary equations:

$$\pi_i = \pi_{i-1} r + \pi_i (1 - r - s) + \pi_{i+1} s$$
$$\left(\frac{r}{s}\right)^i \pi_0 = \left(\frac{r}{s}\right)^{i-1} \pi_0 r + \left(\frac{r}{s}\right)^i \pi_0 (1 - r - s) + \left(\frac{r}{s}\right)^{i+1} \pi_0 s. \quad \checkmark$$

Question: Okay, but we still do not know π_0. How can we determine π_0?

Answer: To determine π_0, we make use of the fact that $\sum_i \pi_i = 1$.

This says that

$$\pi_0 \cdot \left(1 + \frac{r}{s} + \left(\frac{r}{s}\right)^2 + \left(\frac{r}{s}\right)^3 + \cdots\right) = 1$$
$$\pi_0 \cdot \left(\frac{1}{1 - \frac{r}{s}}\right) = 1$$
$$\pi_0 = 1 - \frac{r}{s}.$$

So,

$$\pi_i = \left(\frac{r}{s}\right)^i \cdot \left(1 - \frac{r}{s}\right).$$

Question: What is the average number of packets in the system?

Answer: Let N denote the number of packets in the system. Then

$$\mathbf{E}\left[N\right] = \pi_0 \cdot 0 + \pi_1 \cdot 1 + \pi_2 \cdot 2 + \pi_3 \cdot 3 + \cdots$$

Question: Can we get a closed-form expression for $\mathbf{E}[N]$?

Answer: Yes! It will help to define

$$\rho = \frac{r}{s}$$

for shorthand. Then,

$$\pi_i = \rho^i(1-\rho).$$

So,

$$\begin{aligned}
\mathbf{E}[N] &= 1\rho(1-\rho) + 2\rho^2(1-\rho) + 3\rho^3(1-\rho) + \cdots \\
&= (1-\rho)\cdot\rho\cdot\left(1 + 2\rho + 3\rho^2 + 4\rho^3 + \cdots\right) \\
&= (1-\rho)\cdot\rho\cdot\frac{1}{(1-\rho)^2} \qquad \text{by (1.4)} \\
&= \frac{\rho}{1-\rho}. \qquad (26.1)
\end{aligned}$$

Wow! Equation (26.1) is a really simple formula. For our example, $\rho = \frac{\frac{1}{6}}{\frac{1}{4}} = \frac{2}{3}$ and $\mathbf{E}[N] = \frac{\frac{2}{3}}{1-\frac{2}{3}} = 2$. So on average there are two packets in the system.

26.3 A Harder Example of Solving Stationary Equations in Infinite-State DTMCs

Of course not all infinite-state DTMCs are as easy to solve as the one in the previous section. Consider the DTMC shown in Figure 26.3.

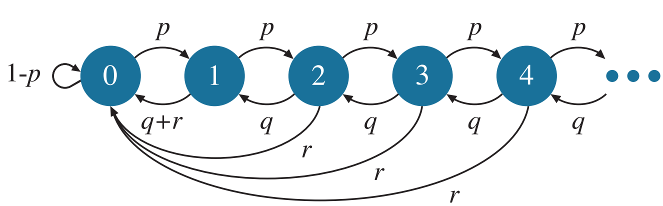

Figure 26.3 *DTMC for processor with failures.*

This kind of chain is often used to model a processor with failures. The chain tracks the number of jobs in the system. At any time step, either the number of jobs increases by 1 (with probability p), or decreases by 1 (with probability q), or a processor failure occurs (with probability r), where $p + q + r = 1$. In the case of a processor failure, all jobs in the system are lost.

To derive the limiting distribution for this chain, simply writing stationary equations will not lead us to the solution. In this case, the z-transform approach (generating functions) from Chapter 6 is very useful. Exercise 26.24 walks you through the steps.

26.4 Ergodicity Questions

We now turn to ergodicity questions.

Recall that in Chapter 25 we asked the following questions for *finite-state* DTMCs:

1. Under what conditions does the limiting distribution exist?
2. How does π_j, the limiting probability of being in state j, compare with p_j, the long-run time-average fraction of time spent in state j?
3. What can we say about m_{jj}, the mean time between visits to state j, and how is this related to π_j?

Recall that in the case of an *ergodic* (aperiodic and irreducible) finite-state DTMC with M states, everything behaves as we would like. Specifically, by Theorem 25.19,

$$0 < \frac{1}{m_{jj}} = \pi_j^{\text{limiting}} = \pi_j^{\text{stationary}} = p_j, \text{ w.p.1.}$$

When the finite-state chain is periodic but irreducible, then the limiting distribution doesn't exist but there's a unique stationary distribution. Specifically, by Theorem 25.20,

$$0 < \frac{1}{m_{jj}} = \pi_j^{\text{stationary}} = p_j, \text{ w.p.1.}$$

For the case of an infinite-state DTMC, we will see that the story is the same when the chain is ergodic. However, the definition of *ergodic* needs to be *strengthened*. The remainder of this chapter is devoted to understanding how to strengthen the definition of *ergodic*. We will figure this out together!

Infinite-state chains are infinitely more complex than finite-state chains. For example, for infinite-state chains, unlike their finite-state counterparts, we will see that when the DTMC is *not* ergodic, even the stationary distribution might not exist. Because of the added complexity inherent in infinite-state chains, we will have to omit some of the proofs in this text. We refer the interested reader to [35] for the omitted proofs.

Consider the three infinite-state DTMCs shown in Figure 26.4.

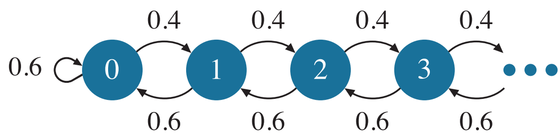

(a) Positive recurrent chain

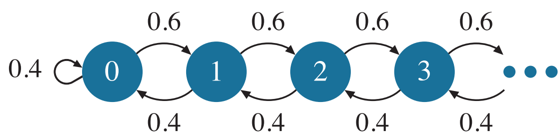

(b) Transient chain

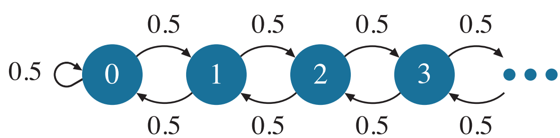

(c) Null recurrent chain

Figure 26.4 *Examples of three chains.*

Question: Which of these chains are aperiodic and irreducible?

Answer: All of them.

Question: For *finite-state* DTMCs that are aperiodic and irreducible, does a limiting distribution always exist?

Answer: Yes, by Theorem 25.6.

Question: Does a limiting distribution exist for all the chains in Figure 26.4?

Answer: We will see that a limiting distribution exists only for chain (a). For chain (a), we saw in Section 26.2 that there is a well-defined stationary probability of being in each state, and these stationary probabilities sum to 1. For the other two chains, we will show that the limiting probability of being in each state is 0, and the limiting probabilities do not sum to 1; hence there does not exist a limiting distribution. Chain (a) has a property called "positive recurrent." Chain (b) is what we call "transient," and chain (c) is "null recurrent." We explain all these terms in this chapter and how they relate to the existence of limiting distributions.

Question: Intuitively, what is the problem with chains (b) and (c) in Figure 26.4?

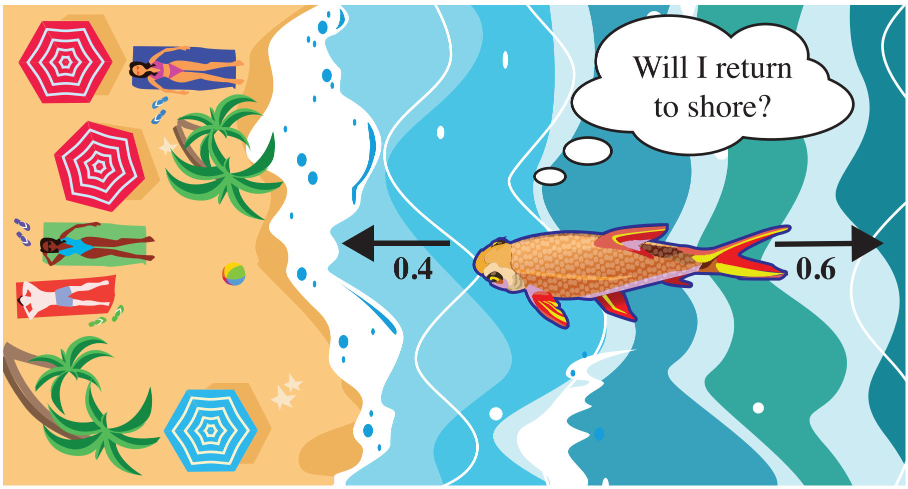

Figure 26.5 *Will the fish return to shore?*

Answer: To get some intuition, it helps to think about π_0, the limiting probability of being in state 0.

Chain (b) can be viewed as an ocean, where the shore is at state 0. Imagine you're a little fish swimming in the ocean. There is a drift away from shore. Think of this as a strong tide, pulling you deeper and deeper into the ocean. Given this drift, it is not obvious that you will keep returning to shore. In fact, we will show that after some point you never return to the shore. Thus, $\pi_0 = 0$. But this same argument holds for any state k that we call the "shore," so $\pi_k = 0$ for all k.

Chain (c) is the most confusing. It's not obvious whether the fish keeps returning to shore. We will show that the fish does in fact always return to shore. However, we will see that the time it takes for the fish to return is infinite. This ends up again resulting in $\pi_k = 0$ for all states k.

To formalize all of this, it helps to first understand the difference between a "recurrent" chain and a "transient" one.

26.5 Recurrent versus Transient: Will the Fish Return to Shore?

Definition 26.2 *We define f_j = probability that a chain starting in state j ever returns to state j.*

Definition 26.3 *A state j is either recurrent or transient:*

- *If $f_j = 1$, then j is a* **recurrent** *state.*
- *If $f_j < 1$, then j is a* **transient** *state.*

Question: What is the distribution of the number of visits to a transient state j?

Answer: Every time we visit state j we have probability $1 - f_j$ of never visiting it again. Hence the number of visits is a Geometric random variable (r.v.) with mean $1/(1 - f_j)$.

Theorem 26.4 *With probability* 1*, the number of visits to a* **recurrent** *state is infinite. With probability* 1*, the number of visits to a* **transient** *state is finite.*

Proof: If a state j is recurrent, then starting in state j, with probability 1 (w.p.1) we will visit j again. Thus, repeating this argument, we see that w.p.1 state j will be visited an infinite number of times. In contrast, if state j is transient, then every time we visit state j, there is some probability $(1 - f_j)$ that we will never again visit j. With probability 1, that $1 - f_j$ probability event will eventually happen. That is, w.p.1, after some point we will never again revisit state j. ■

Theorem 26.5 *Let $P^n_{ij} = (\mathbf{P}^n)_{ij}$ denote the probability that the chain will be in state j after n steps, given that the chain is in state i now.*

- *If state i is recurrent, then $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} = \infty$.*
- *If state i is transient, then $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} < \infty$.*

Proof: Observe that $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii} = \mathbf{E}\left[\text{Number visits to state } i\right]$.

To see this, note that if N is the number of visits to state i, then we can write

$$N = I_1 + I_2 + I_3 + \cdots ,$$

where I_n is an indicator r.v. which equals 1 if we're in state i at the nth time step.

Thus:

$$\begin{aligned} \mathbf{E}[N] &= \mathbf{E}[I_1] + \mathbf{E}[I_2] + \mathbf{E}[I_3] + \cdots \qquad \text{(Linearity of Expectation)} \\ &= \left(\mathbf{P}^1\right)_{ii} + \left(\mathbf{P}^2\right)_{ii} + \left(\mathbf{P}^3\right)_{ii} + \cdots \\ &= \sum_{n=0}^{\infty} (\mathbf{P}^n)_{ii}. \end{aligned}$$

Finally, by Theorem 26.4, for a recurrent state $\mathbf{E}[N] = \infty$, while for a transient one, $\mathbf{E}[N] < \infty$. ■

Theorem 26.6 (Recurrence class property) *If state i is recurrent and i communicates with j, (written $i \longleftrightarrow j$), then j is recurrent.*

We start with the intuition for Theorem 26.6. Consider Figure 26.6. We know that we come back to i infinitely many times. By the definition of "communicates," every time we are in i, we have some probability of taking the road to j, and once we are in j, we have some probability of taking the road to i. So, for every visit to i, there's some non-zero probability that we'll also visit j. Therefore the number of visits to j is proportional to the number of visits to i. Because the number of visits to i is infinite, so is the number of visits to j.

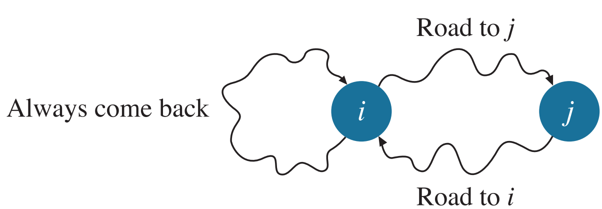

Figure 26.6 *Proof of Theorem 26.6.*

Now for the formal proof.

Proof: We know that i communicates with j. Thus, there exists an m such that $(\mathbf{P}^m)_{ji} > 0$ and there exists n such that $(\mathbf{P}^n)_{ij} > 0$. We also know that $\sum_{s=0}^{\infty} (\mathbf{P}^s)_{ii} = \infty$, because state i is recurrent.

What we want to show is that $\sum_{t=0}^{\infty} (\mathbf{P}^t)_{jj} = \infty$.

Now

$$\sum_{t=0}^{\infty} (\mathbf{P}^t)_{jj} \geq \sum_{s=0}^{\infty} (\mathbf{P}^{m+s+n})_{jj}, \tag{26.2}$$

since the left-hand side of (26.2) considers all j-to-j paths, while the right-hand side considers only those of length at least $m + n$.

We can now further constrain our j-to-j paths by insisting that we must use the first m steps of our path to go from j to i and the last n steps to go from i to j. Specifically:

$$\begin{aligned}
\sum_{t=0}^{\infty} \left(\mathbf{P}^t\right)_{jj} &\geq \sum_{s=0}^{\infty} \left(\mathbf{P}^{m+s+n}\right)_{jj} \\
&\geq \sum_{s=0}^{\infty} \left(\mathbf{P}^m\right)_{ji} \left(\mathbf{P}^s\right)_{ii} \left(\mathbf{P}^n\right)_{ij} \\
&= \left(\mathbf{P}^m\right)_{ji} \left(\mathbf{P}^n\right)_{ij} \sum_{s=0}^{\infty} \left(\mathbf{P}^s\right)_{ii} \qquad \text{(pulling out positive constants)} \\
&= \infty \qquad \text{(because state } i \text{ is recurrent).}
\end{aligned}$$

We have thus proven that state j is recurrent. ■

Theorem 26.7 (Transience class property) *If state i is transient and i communicates with j, ($i \longleftrightarrow j$), then j is transient.*

Proof: This follows directly from the previous Theorem 26.6. Suppose by contradiction that state j is recurrent. Then because j and i communicate, i is recurrent as well, which is a contradiction to the assumption. ■

We have thus seen that *in an irreducible Markov chain, either all states are transient, or all are recurrent!*

Theorem 26.8 *For a transient Markov chain,*

$$\lim_{n\to\infty} \left(\mathbf{P}^n\right)_{ij} = 0, \quad \forall j.$$

Hence, the limiting distribution does not exist.

Proof: As we have seen, in a transient Markov chain there is some point after which we never visit state j again. So the probability of being in state j after n steps is zero as $n \to \infty$, that is,

$$\lim_{n\to\infty} \left(\mathbf{P}^n\right)_{ij} = 0,$$

and this holds for every state j.

Now

$$\sum_{j=0}^{\infty} \pi_j = 0$$

because the sum of a countable number of 0's is still 0. Thus the limiting distribution does not exist. ■

Theorem 26.9 *For Markov chains where the limiting probabilities are all zero, no stationary distribution exists.*

Proof: The fact that no stationary distribution exists follows from an argument similar to that in the proof of Theorem 26.1. For details, see [35]. ■

26.6 Infinite Random Walk Example

It's not so obvious how to argue whether a chain is transient or recurrent. The following example illustrates how this is done.

Consider the random walk shown in Figure 26.7, where at each step a gambler either gains a dollar (with probability p) or loses a dollar (with probability $q = 1 - p$). We'd like to determine whether the chain is transient or recurrent.

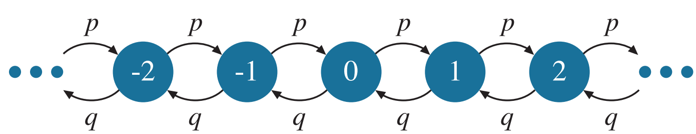

Figure 26.7 *Gambler's walk.*

Because all states communicate, it follows from Theorems 26.6 and 26.7 that either *all* states are transient or all are recurrent. Hence to determine whether the chain is recurrent or transient, it suffices to look at state 0.

To determine whether state 0 is transient or recurrent, we invoke Theorem 26.5. Let

$$V = \sum_{n=1}^{\infty} (\mathbf{P}^n)_{00}$$

denote the expected number of visits to state 0. If V is finite, then state 0 is transient. Otherwise it is recurrent.

Since one cannot get from 0 to 0 in an odd number of steps, it follows that

$$V = \sum_{n=1}^{\infty} (\mathbf{P}^n)_{00} = \sum_{n=1}^{\infty} \left(\mathbf{P}^{2n}\right)_{00} = \sum_{n=1}^{\infty} \binom{2n}{n} p^n q^n. \tag{26.3}$$

We now simplify this equation using Lavrov's lemma.

Lemma 26.10 *(Due to Misha Lavrov) For $n \geq 1$,*

$$\frac{4^n}{2n+1} < \binom{2n}{n} < 4^n. \tag{26.4}$$

Proof: By simple binomial expansion,

$$\sum_{k=0}^{2n} \binom{2n}{k} = (1+1)^{2n} = 2^{2n} = 4^n.$$

Since $\binom{2n}{n}$ is the largest term in the sum, it follows that it is bigger than the average term, $4^n/(2n+1)$. However, it is also smaller than the total sum, 4^n. ■

Substituting (26.4) into (26.3), we get that

$$\sum_{n=1}^{\infty} \frac{4^n}{2n+1} p^n q^n < V < \sum_{n=1}^{\infty} 4^n p^n q^n. \tag{26.5}$$

If we substitute $p = q = \frac{1}{2}$ into the left-hand side of (26.5), we get that

$$V > \sum_{n=1}^{\infty} \frac{4^n}{2n+1} \cdot \frac{1}{4^n} = \sum_{n=1}^{\infty} \frac{1}{2n+1} = \infty. \tag{26.6}$$

If instead we assume $p \neq q$ and consider the right-hand side of (26.5), we get that

$$V < \sum_{n=1}^{\infty} (4pq)^n < \infty \quad (\text{since } 4pq < 1). \tag{26.7}$$

Thus by (26.6) and (26.7) we see that $V = \sum_{n=1}^{\infty} (\mathbf{P}^n)_{00}$ is infinite if and only if $p = \frac{1}{2}$. So the chain is recurrent if and only if $p = \frac{1}{2}$.

We have thus proven Theorem 26.11.

Theorem 26.11 *The Gambler's walk shown in Figure 26.7 is recurrent only when $p = \frac{1}{2}$ and is transient otherwise.*

26.7 Back to the Three Chains and the Ergodicity Question

Let's return to the three infinite-state chains in Figure 26.4, repeated in Figure 26.8:

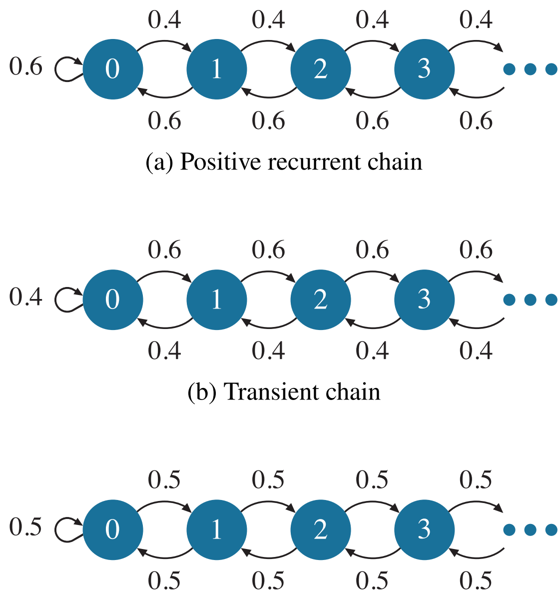

(a) Positive recurrent chain

(b) Transient chain

(c) Null-recurrent chain

Figure 26.8 *Three chains of Figure 26.4.*

26.7.1 Figure 26.8(a) is Recurrent

From what we've learned, chain (a) is recurrent, simply by virtue of the fact that we know that it has a stationary distribution (recall from Theorem 26.9 that, for a transient chain, no stationary distribution exists).

26.7.2 Figure 26.8(b) is Transient

Chain (b) is transient. Intuitively, imagine that j is very high. If you're in state j, the world looks very much like a 2D-infinite Gambler's walk where the drift goes to the right. However, this is not a formal proof of transience. Theorem 26.12 provides a proof by precisely relating chain (b) to the two-way Gambler's walk.

Theorem 26.12 *Chain (b) in Figure 26.8 is transient.*

Proof: [This cute proof is due to Misha Ivkov.] The proof relies on looking at two other chains which we have already analyzed, shown in Figure 26.9.

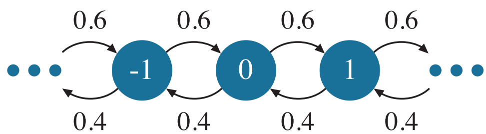

(b′) Two-way transient chain

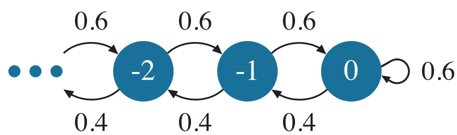

(b″) Recurrent chain

Figure 26.9 *Some helper chains for proving Theorem 26.12.*

We define a few quantities:

$$\begin{aligned}
f_0 &= \mathbf{P}\{\text{return to 0 in chain (b)}\}\\
f_0' &= \mathbf{P}\{\text{return to 0 in chain (b}')\}\\
f_0'' &= \mathbf{P}\{\text{return to 0 in chain (b}'')\}\\
f_{ij} &= \mathbf{P}\{\text{eventually visit } j\text{, given currently in } i \text{ in chain (b)}\}\\
f_{ij}' &= \mathbf{P}\{\text{eventually visit } j\text{, given currently in } i \text{ in chain (b}')\}\\
f_{ij}'' &= \mathbf{P}\{\text{eventually visit } j\text{, given currently in } i \text{ in chain (b}'')\}.
\end{aligned}$$

To show that chain (b) is transient, it suffices to show that $f_0 < 1$. By conditioning we have:

$$f_0 = (0.4) \cdot 1 + (0.6) \cdot f_{10}. \tag{26.8}$$

Now observe that

$$f_{10} = f_{10}'. \tag{26.9}$$

This is due to the fact that chain (b′) looks identical to chain (b) except for the states to the left of state 0; however, those states left of 0 don't matter in computing f_{10} or f_{10}'. Substituting (26.9) into (26.8) we have:

$$f_0 = (0.4) \cdot 1 + (0.6) \cdot f_{10}'. \tag{26.10}$$

Now observe that

$$f_{10}' < 1. \tag{26.11}$$

Equation (26.11) follows from this series of equations:

$$\begin{aligned} 1 &> f_0' \quad \text{because chain (b}'\text{) is transient} \\ &= (0.4) \cdot f_{-1,0}' + (0.6) \cdot f_{10}' \quad \text{by conditioning} \\ &= (0.4) \cdot f_{-1,0}'' + (0.6) \cdot f_{10}' \quad \text{what's right of state 0 doesn't matter} \\ &= (0.4) \cdot 1 + (0.6) \cdot f_{10}' \quad \text{because (b}''\text{) is recurrent} \end{aligned}$$

Thus $0.6 > (0.6) \cdot f_{10}'$, and hence $f_{10}' < 1$.

Combining (26.11) and (26.10), we have:

$$f_0 = (0.4) \cdot 1 + (0.6) \cdot f_{10}' < (0.4) \cdot 1 + (0.6)1 = 1. \quad \blacksquare$$

26.7.3 Figure 26.8(c) is Recurrent

Chain (c) is recurrent. This follows from the fact that f_0, the probability of returning to state 0, is at least as high in chain (c) as in the Gambler's walk of Figure 26.7 with $p = 0.5$, and we've shown that $f_0 = 1$ for the Gambler's walk with $p = 0.5$.

Question: Given that chain (c) is recurrent, does this mean that the limiting distribution exists for this third chain? If so, what is it?

Answer: Although this chain (c) is recurrent, and irreducible and aperiodic, it turns out that these are *not* enough to guarantee the existence of the limiting distribution. To see why, we turn to Theorem 26.13, known as the Ergodic Theorem of Markov Chains.

26.8 Why Recurrence Is Not Enough

Theorem 26.13 (Ergodic Theorem of Markov Chains) *Given a recurrent, aperiodic, irreducible DTMC,* $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij}$ *exists and*

$$\pi_j = \frac{1}{m_{jj}}, \ \forall j.$$

The Ergodic Theorem of Markov Chains is saying the same thing that we saw in Theorem 25.12, about π_j being the reciprocal of m_{jj}. However, those theorems were restricted to *finite*-state chains. The fact that we now allow for infinite-state chains makes the proof *much* more technical than for the case of a finite number of states, and we refer the reader to [35, section 9.10].

Question: The Ergodic Theorem (Theorem 26.13) *seems* to suggest that recurrent + aperiodic + irreducible suffices for the limiting distribution to exist. What's wrong with this?

Answer: There's an important distinction. While π_j exists, it is not necessarily positive. We're told that

$$\pi_j = \frac{1}{m_{jj}},$$

but m_{jj} can be infinite!

In the case of a finite irreducible chain, we were guaranteed that m_{jj} is finite, but that's not necessarily true for an infinite state chain. In particular for the chain in Figure 26.8(c), Theorem 26.14 shows that $m_{jj} = \infty$ for all states j. Hence the limiting probability of being in state j exists, but is zero. Consequently, the limiting distribution does not exist (since a countable number of 0's can't sum to 1). Furthermore, by Theorem 26.9 no stationary distribution exists for this chain either.

Theorem 26.14 *For chain (c) in Figure 26.8, $m_{jj} = \infty$, for all states j.*

Proof: We show that $m_{00} = \infty$. Suppose by contradiction that m_{00} is finite.

Observe that

$$m_{00} = 1 + \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot m_{10}.$$

Thus, given that m_{00} is finite, it must also be the case that m_{10} is finite. Now observe that

$$\begin{aligned} m_{10} &= 1 + \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot m_{20} \\ &= 1 + \frac{1}{2} \cdot (m_{21} + m_{10}) \\ &= 1 + \frac{1}{2} \cdot 2m_{10} \\ &= 1 + m_{10}. \end{aligned}$$

But the only way that

$$m_{10} = 1 + m_{10}$$

is if $m_{10} = \infty$, which is a contradiction. Hence,

$$m_{00} = \infty.$$

The argument is very similar to show that $m_{jj} = \infty$, where $j \geq 1$, and we leave it as Exercise 26.12. ■

We have seen that while chains (a) and (c) in Figure 26.4 are both recurrent, they differ in the mean time to return to a state.

Definition 26.15 *Recurrent Markov chains fall into two types:* **positive recurrent** *and* **null recurrent**. *In a positive-recurrent MC, the mean time between recurrences (returning to the same state) is finite. In a null-recurrent MC, the mean time between recurrences is infinite.*

Both positive recurrence and null recurrence are class properties.

Theorem 26.16 (More class properties) *If state i is positive recurrent and $i \longleftrightarrow j$, then j is positive recurrent. If state i is null recurrent and $i \longleftrightarrow j$, then j is null recurrent.*

Proof: See Exercise 26.23. ■

26.9 Ergodicity for Infinite-State Chains

Definition 26.17 *An* **ergodic** *DTMC is one that has all three desirable properties: aperiodicity, irreducibility, and positive recurrence.*

Theorem 26.18 *For an ergodic DTMC, the limiting distribution exists.*

Proof: By Theorem 26.13, the limiting *probabilities* (the π_j's) exist. By positive recurrence, they are all positive. All that remains is to show that $\sum_j \pi_j = 1$. To see this, recall p_j, the time-average fraction of time that the chain spends in state j. We proved in Section 25.6 that, for finite-state irreducible, aperiodic chains,

$$p_j = \frac{1}{m_{jj}} \quad \text{w.p.1.}$$

This was proven via invoking the Strong Law of Large Numbers (SLLN).

Question: What was the one thing needed for SLLN to hold?

Answer: We needed the mean time between renewals, m_{jj}, to be finite.

For the case of infinite-state DTMCs, the same argument as in Section 25.6 goes through, provided that m_{jj} is finite, which it is for a positive recurrent chain.

Now, observe that

$$\sum_{j=0}^{\infty} p_j = 1,$$

since a random walk must be in *some* state at all time steps, so the fraction of time it spends in each state must total to 1. Hence, since

$$p_j = \frac{1}{m_{jj}} = \pi_j,$$

it also follows that

$$\sum_{j=0}^{\infty} \pi_j = 1. \qquad \blacksquare$$

Remark: For a finite-state DTMC, positive recurrence is a consequence of irreducibility. This fact was proven in Exercise 25.19. Hence, for finite-state chains, aperiodicity and irreducibility suffice for ergodicity.

To summarize, infinite-state DTMCs are much more complicated than finite-state DTMCs because positive recurrence is required for the limiting distribution (and stationary distribution) to exist and we don't always have positive recurrence. Fortunately, as explained in the Theorem 26.19 and the associated Remark, we never need to check for positive recurrence.

Theorem 26.19 (Summary theorem) *An irreducible, aperiodic DTMC belongs to* **one** *of the following two classes:*

Either:

(i) All the states are transient, or all are null recurrent. In this case $\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij} = 0, \forall j$, *and there does not exist a limiting distribution or a stationary distribution.*

Or:

(ii) All states are positive recurrent. Then the limiting distribution $\vec{\pi} = (\pi_0, \pi_1, \pi_2, \ldots)$ *exists, and there is a positive probability of being in each state. Here,*

$$\pi_j = \lim_{n\to\infty} (\mathbf{P}^n)_{ij} > 0, \quad \forall i$$

is the limiting probability of being in state j*. In this case* $\vec{\pi}$ *is a stationary distribution, and no other stationary distribution exists. Also,* $\pi_j = \frac{1}{m_{jj}}$*, where* m_{jj} *is the mean number of steps between visits to state* j*.*

Proof: We know by Theorems 26.16 and 26.7 that transience, null recurrence, and positive recurrence are class properties, meaning that in an irreducible Markov chain all the states are of the same one type.

If all states are transient, then by Theorem 26.8, the limiting probabilities are all zero and no limiting distribution exists. Further, by Theorem 26.9, no stationary distribution exists.

If all states are null recurrent, then by Theorem 26.13, all the limiting probabilities are zero, so they can't add up to 1, hence no limiting distribution exists. Also, again by Theorem 26.9, no stationary distribution exists.

If all states are positive recurrent, then by Theorem 26.18, the limiting distribution exists. Finally, by Theorem 26.1, when the limiting distribution exists, it is equal to the unique stationary distribution. ■

Important Remark: What is nice about Theorem 26.19 is that it tells us that we never have to actually determine whether our DTMC is positive recurrent. It suffices to simply check for irreducibility and aperiodicity and then solve the stationary equations. If these stationary equations yield a distribution, then that distribution is also the limiting probability distribution.

26.10 Exercises

26.1 **Irreducibility, aperiodicity, and positive recurrence**
For each of the following transition matrices: (i) Is the DTMC irreducible? (ii) Is it aperiodic? (iii) Is it positive recurrent? [Note: If the period is not defined, then the chain is *not* aperiodic.]

$$\text{(a)} \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{(b)} \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{(c)} \begin{pmatrix} \frac{1}{3} & 0 & \frac{2}{3} \\ \frac{1}{4} & \frac{3}{4} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{(d)} \begin{pmatrix} 0 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

26.2 **Time to empty**
Consider a router where, at each time step, the number of packets increases by 1 with probability 0.4 and decreases by 1 with probability 0.6. How long does it take for the router to empty? The Markov chain depicting the number of packets is shown in Figure 26.10. Let $T_{1,0}$ denote the time to get from state 1 to state 0. (a) Compute $\mathbf{E}\left[T_{1,0}\right]$. (b) Compute $\mathbf{Var}(T_{1,0})$. [Hint: The variance computation is a little tricky. Be careful not to lump together distinct random variables.]

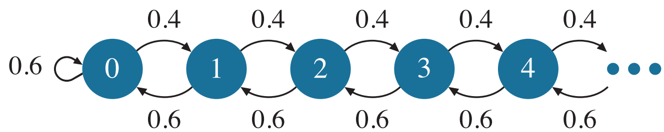

Figure 26.10 *Number of packets at router.*

.

26.3 **Time to empty – extra strength**
Consider the same setup as in Exercise 26.2. Let $T_{n,0}$ denote the time to get from state n to state 0. (a) Compute $\mathbf{E}\left[T_{n,0}\right]$. (b) Compute $\mathbf{Var}(T_{n,0})$.

26.4 **Gambling game**
Dafna starts out with zero dollars. Every day she gains a dollar with probability p, stays put with probability s, or loses all her money (goes broke) with probability b, where $p + s + b = 1$. Dafna plays the game forever. Use a DTMC to determine the stationary probability that Dafna has i dollars. What happens to your stationary probability when $s = 0$? What is Dafna's long-run expected money (for any general s)?

26.5 **Reviewing the definitions**
For the DTMC shown in Figure 26.11, circle all the statements that are *true*. Provide a one-line explanation for every item that you circled.

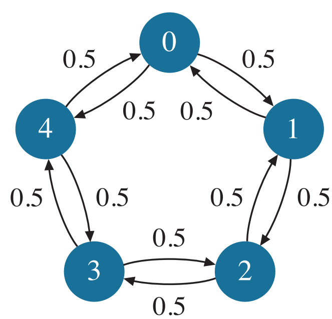

Figure 26.11 *Chain for Exercise 26.5.*

(a) The chain is null recurrent.
(b) The chain is positive recurrent.
(c) The chain is time-reversible.
(d) $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{00}$ is finite.
(e) The chain is irreducible.
(f) The chain is aperiodic.
(g) There are an infinite number of stationary distributions .

(h) $m_{00} = 0.5$.
(i) $\lim_{t\to\infty} \frac{N_0(t)}{t} = 0$.
(j) The chain is ergodic.
(k) $m_{00} = \infty$.
(l) The limiting distribution exists.
(m) $f_0 < 1$.
(n) $\left(\mathbf{P}^5\right)_{00} = 2^{-5}$.

Glossary:
P is the transition probability matrix
f_j = probability that we ever return to state j given that we start in state j.
$N_j(t)$ = number of visits to state j by time t.
m_{jj} = mean number of time steps to return to j given we're in state j.

26.6 **Equivalent definitions**
Given an ergodic DTMC with transition matrix **P**, make as few equivalence classes as you can out of the expressions below. For example, your answer might be:

$$a = b = d = i = j; \qquad c = g = h; \qquad e; \qquad f = j.$$

(a) $\lim_{n\to\infty} (\mathbf{P}^n)_{jj}$
(b) $\frac{1}{p_j}$
(c) $\frac{1}{m_{jj}}$
(d) π_j
(e) $\sum_{n=0}^{\infty} (\mathbf{P}^n)_{jj}$
(f) $1 + \sum_{k\neq j} P_{jk} \cdot m_{kj}$
(g) $\sum_k \pi_k P_{kj}$
(h) $\lim_{t\to\infty} \frac{N_j(t)}{t}$
(i) $\lim_{n\to\infty} (\mathbf{P}^n)_{kj}$
(j) f_j

Glossary:
P_{ij} = (i, j)th entry of transition matrix **P**.
m_{jj} = mean number of time steps to return to j given we're in state j.
π_j = limiting probability of being in state j.
p_j = time-average fraction of time that chain spends in state j.
f_j = probability that a chain starting in state j ever returns to state j.
$N_j(t)$ = number of visits to state j by time t.

26.7 **A positive recurrent chain**
Recall the chain in Figure 26.12 that we've seen many times before.
(a) Use Theorem 26.19 to explain how we know that the chain is positive recurrent.

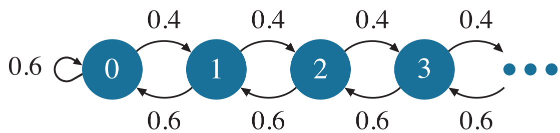

Figure 26.12 *Chain for Exercise 26.7.*

(b) Derive m_{00} via conditioning on the next step. Then use a theorem to explain why your answer makes sense.

26.8 **Stationary but not limiting**
We've seen several examples of finite-state DTMCs for which the stationary distribution exists, but the limiting distribution does not. Provide an example of an *infinite-state*, irreducible DTMC for which there is a unique stationary distribution, but no limiting distribution exists. Solve for the stationary distribution.

26.9 **Expected time until k failures**
This is a repeat of Exercise 4.18, where we want to derive the expected number of minutes until there are k consecutive failures in a row, assuming that a failure occurs independently every minute with probability p. However, this time, solve the problem by finding the limiting probability of some Markov chain. Include a picture of your Markov chain. [Hint: You will have to think a bit to see how to convert from the limiting probabilities of the Markov chain to what you really want.]

26.10 **Threshold queue**
Figure 26.13 depicts a "threshold queue" with integer parameter t.

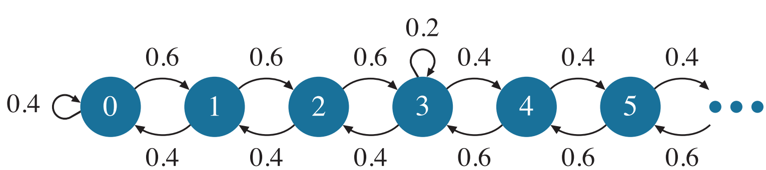

Figure 26.13 *For Exercise 26.10. Markov chain for threshold queue with* $t = 3$.

When the number of jobs is $< t$, then the number of jobs decreases by 1 with probability 0.4 and increases by 1 with probability 0.6 at each time step. However, when the number of jobs increases to $> t$, then the reverse is true and the number of jobs increases by 1 with probability 0.4 and decreases by 1 with probability 0.6 at each time step.

(a) Derive the stationary probability distribution as a function of t, for arbitrary threshold t.

(b) Given that you have a stationary distribution, explain why it follows that this distribution is the limiting distribution.
(c) Compute the mean number of jobs, $\mathbf{E}[N]$, as a function of t.
(d) What happens to $\mathbf{E}[N]$ when $t = 0$? Does this answer make sense?

26.11 **I am one with the chain**
[Proposed by Misha Ivkov] For the two chains in Figure 26.14, determine if each is positive recurrent, transient, or null recurrent. (Do not do anything complicated – just look and make a simple argument.)

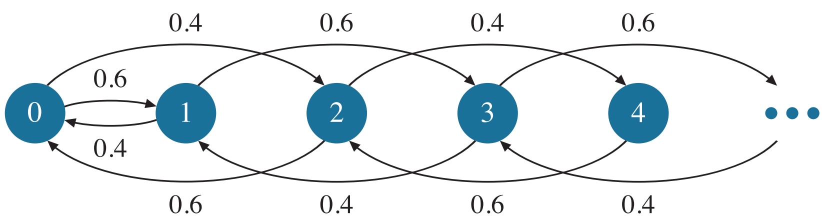

(a)

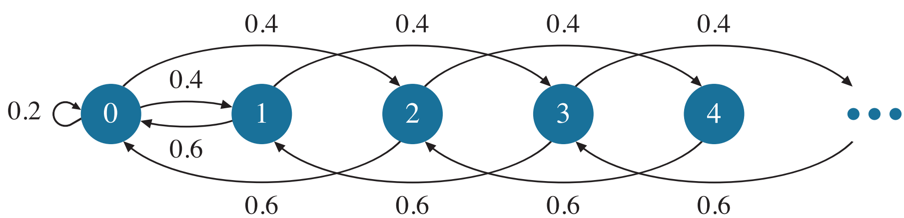

(b)

Figure 26.14 *Markov chains for Exercise 26.11.*

26.12 **Finish proof of Theorem 26.14**
Complete the proof of Theorem 26.14 in the chapter.

26.13 **Deriving the mean time between visits**
Consider the two DTMCs in Figure 26.15. For each chain, derive m_{00}, the mean number of time steps between visits to state 0. If you claim that $m_{00} = \infty$, you need to prove it. If you claim that $m_{00} < \infty$, you need to specify what m_{00} is.

26.14 **Walking in a winter wonderland**
[Proposed by Misha Ivkov] Figure 26.16 shows an infinite binary tree

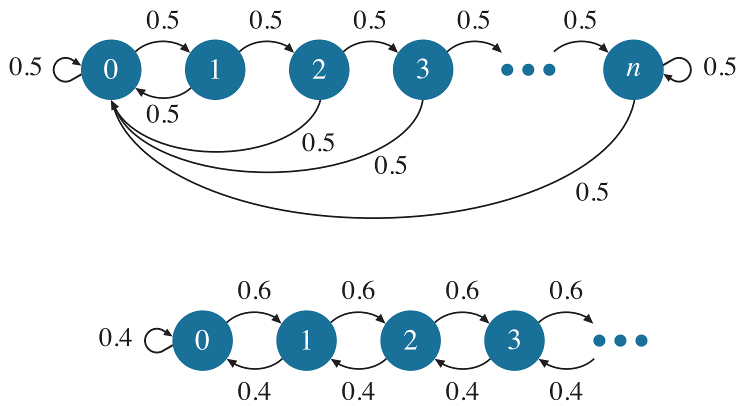

Figure 26.15 *Two chains for Exercise 26.13.*

representing a DTMC, where $p + q + r = 1$. Label the layers where node 1 is layer 0, nodes 2 and 3 are layer 1, and nodes 2^k through $2^{k+1} - 1$ comprise layer k.

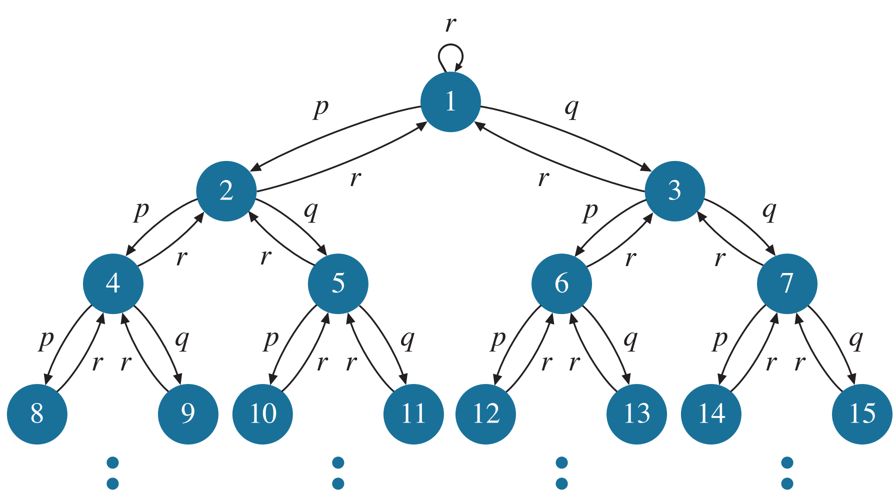

Figure 26.16 *Markov chain for Exercise 26.14.*

(a) Under what conditions does the limiting distribution exist (explain)? What is that limiting distribution? [Hint: It suffices to derive the following: (i) Express π_{2i} in terms of π_i; (ii) Express π_{2i+1} in terms of π_i; (iii) Derive π_1.]

(b) What is the long-run expected layer?

26.15 **Pricing model**

You are the market maker for GOGO. You have no clue whether GOGO stock will rise or fall, but you are obligated to buy or sell single shares from customers at all times. However, you do get to set the share price. To control the size of your position (number of shares of GOGO you own), when you are long (that is, own) GOGO, you set the price so that with probability $p < \frac{1}{2}$ your next trade is a buy, and with probability $q = 1 - p$ your next trade is a sell. In contrast, if you are short (that is, owe) GOGO, you set the price so that with probability p your next trade is a sell, and with probability q your next trade is a buy.

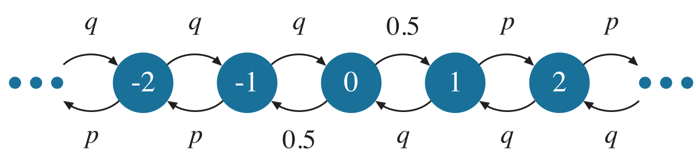

Figure 26.17 *Bidirectional chain for pricing.*

Your position is represented by the bidirectional chain in Figure 26.17. A negative state indicates how many shares you owe, and a positive state indicates how many shares you own.

(a) Given this pricing, what does your position tend to revert to?
(b) Derive the time-average fraction of time spent in each state.
(c) Why weren't you asked to find the limiting probabilities?
(d) What is the expected (absolute value) size of your position?

26.16 **Brownian motion**

Brownian motion models the walk of a drunkard, as depicted by Figure 26.18. Assume that the drunkard starts in state 0 and makes one move per day.

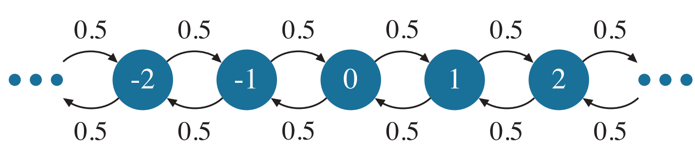

Figure 26.18 *Brownian motion.*

(a) Let $T_{0,n}$ denote the number of days it takes the drunkard to get from state 0 to state n. What is $\mathbf{E}\left[T_{0,n}\right]$? Prove it formally.
(b) Let $D_{0,n}$ denote the number of days until the drunkard first achieves distance n from its origin (that is, the drunkard first hits either state n or $-n$). Prove that $\mathbf{E}\left[D_{0,n}\right] = n^2$. Provide any proof that you like. The steps below are (optional) helping steps:

(i) Define $D_{i,j}$ to be the number of days until the drunkard first gets to either state j or state $-j$, given that the drunkard starts in state i. Argue that $D_{0,n} = D_{0,1} + D_{1,2} + D_{2,3} + \cdots + D_{n-1,n}$.
(ii) What is $\mathbf{E}\left[D_{0,1}\right]$? Derive $\mathbf{E}\left[D_{1,2}\right]$. Derive $\mathbf{E}\left[D_{2,3}\right]$. Do you see a pattern?
(iii) Guess a formula for $\mathbf{E}\left[D_{i,i+1}\right]$ and verify that your guess is correct.
(iv) Returning to step (i), derive $\mathbf{E}\left[D_{0,n}\right]$.

26.17 **Wandering around the Pittsburgh airport**
[Proposed by Adrian Abedon] At the Pittsburgh international airport, each of the terminals A, B, C, and D now have an infinite number of gates. A weary traveler in the airport wanders the gates at random, starting from the central hub (0). The traveler's movement is modeled by the Markov chain in Figure 26.19.

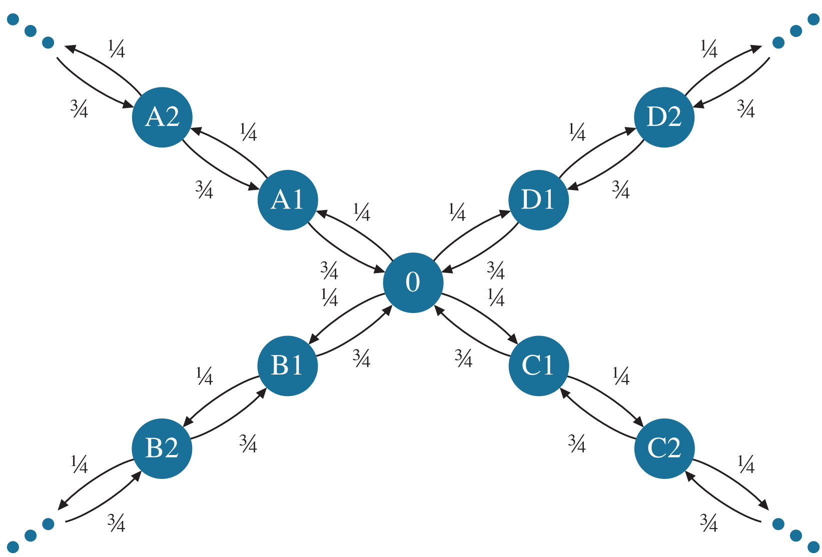

Figure 26.19 *DTMC for Exercise 26.17.*

(a) Find the stationary distribution $\pi_{Ai}, \pi_{Bi}, \pi_{Ci}, \pi_{Di}$, and π_0.
(b) Find $m_{0,A2}$, the expected time for the traveler to get to their gate $A2$.

26.18 **Gambler ruin problem**
Imagine a gambler who is equally likely to win a dollar or to lose a dollar every day (see Figure 26.20). The gambler starts out with i dollars. What is $P_{i,n}$, the probability that he makes it to n dollars before going bankrupt? [Hint: It helps to guess an expression for $P_{i,n}$ in terms of $P_{i+1,n}$.]

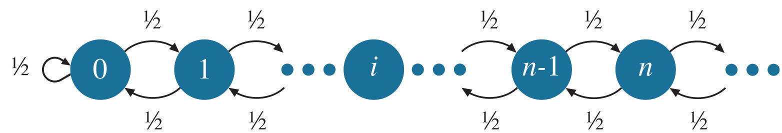

Figure 26.20 *State i indicates that there are i more tails than heads.*

26.19 **Mouse in infinite maze**
[Proposed by Misha Ivkov] A mouse is trapped in a maze with an infinite number of layers. At each time step, with probability $\frac{2}{3}$, the mouse decreases its layer by 1, and with probability $\frac{1}{3}$ it increases its layer by 1, as shown in Figure 26.21. The mouse can only escape from layer 0. Suppose that we drop the mouse into this maze at a random layer ≥ 1, where the mouse is dropped at layer i with probability $\frac{1}{2^i}$. Let T denote the number of steps until the mouse escapes (gets to layer 0).

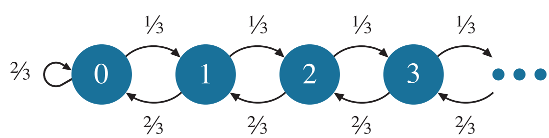

Figure 26.21 *Markov chain for Exercise 26.19, where the state represents the mouse's current layer.*

(a) Derive the z-transform, $\widehat{T}(z)$. It helps to define T in terms of T_i, where T_i is the time to escape when starting in layer i. Follow these steps:
  (i) Start by deriving the z-transform of T_1. [Hint: You will need to use the fact that $\widehat{X}(z) = 1$ when $z = 1$ for any discrete r.v. X.]
  (ii) Now derive the z-transform of T_i in terms of T_1.
  (iii) Finally derive the z-transform of T by conditioning on the starting state i.

(b) Differentiate your answer to (a) to get $\mathbf{E}[T]$.

26.20 **2D gambler's walk**
[Proposed by Weina Wang] A drunkard walks on the two-dimensional plane depicted in Figure 26.22. Formally prove or disprove that this is a recurrent chain. You will need to make use of the following equations:[1]

$$\sum_{k=0}^{n} \binom{n}{k}^2 = \binom{2n}{n} \geq \frac{4^{n-1}}{\sqrt{n}}.$$

[1] The equality is a special case of Vandermonde's identity. The inequality can be derived from Stirling's approximation. See Section 1.5 for a discussion of both.

[Hint: This will look a lot like Section 26.6. In expressing the V quantity, it helps to use a *single* summation over k, rather than a double summation.]

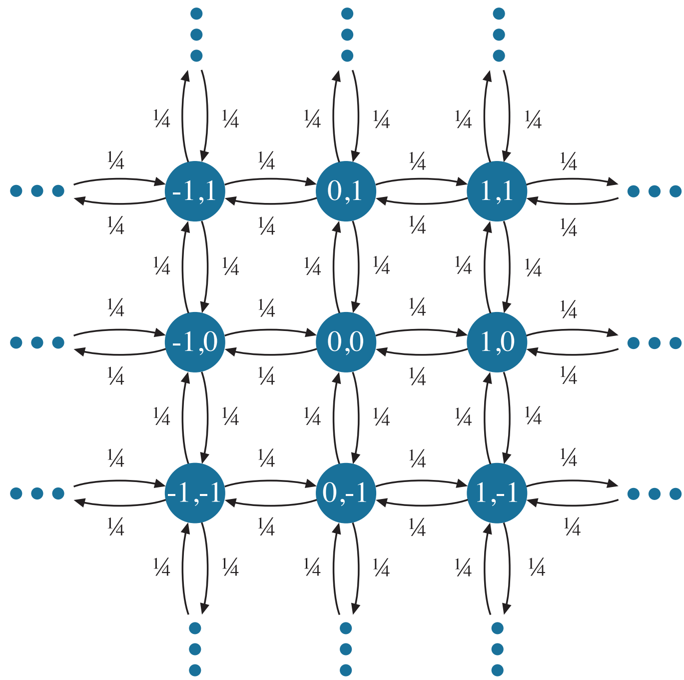

Figure 26.22 *Markov chain for Exercise 26.20.*

26.21 **Hellbound**

[Proposed by Alec Sun] Every lifetime Iggy is reincarnated into either heaven or hell. Since Iggy is a bad boy, reincarnations occur as follows:

– If Iggy is in heaven, then he will always be reincarnated into hell.
– If Iggy is in hell and has been in hell for $j \geq 1$ consecutive lifetimes since last being in heaven, then with probability $0 < p_j < 1$ he is reincarnated into heaven and with probability $1 - p_j$ he is reincarnated into hell.

Figure 26.23 depicts the infinite-state DTMC showing Iggy's state:

(a) Is the DTMC in Figure 26.23 irreducible, assuming that every value of p_j satisfies $0 < p_j < 1$?
(b) Let $c \in (0, 1)$ be a constant and suppose $p_j = c$ for all $j \geq 1$. Is our DTMC transient, positive recurrent, or null recurrent? Prove your answer.
(c) Suppose $p_j = \frac{1}{j+1}$ for all $j \geq 1$. Is this DTMC transient, positive recurrent, or null recurrent? Prove your answer. [Hint: It may be easier to consider $1 - f_0$.]

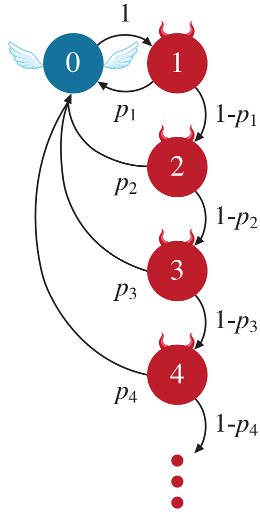

Figure 26.23 *DTMC for Exercise 26.21.*

(d) Suppose $p_j = 2^{-j}$ for all $j \geq 1$. Is this DTMC transient, positive recurrent, or null recurrent? Prove your answer. [Hint: Compute f_0.]

26.22 **Irreducible finite-state chains are positive recurrent ... again**
This is a repeat of Exercise 25.19. Once again, you're being asked to prove that in a finite-state, irreducible DTMC, all states are positive recurrent. This time, follow these steps:
(a) First show that all states are recurrent.
(b) Now show that there exists at least one positive recurrent state.
(c) Now make use of Theorem 26.16 to finish the proof.

26.23 **Proving that positive recurrence and null recurrence are class properties**
Prove Theorem 26.16, which states that positive recurrence and null recurrence are class properties. [Warning: This is a difficult exercise.]

26.24 **Processor with failures**

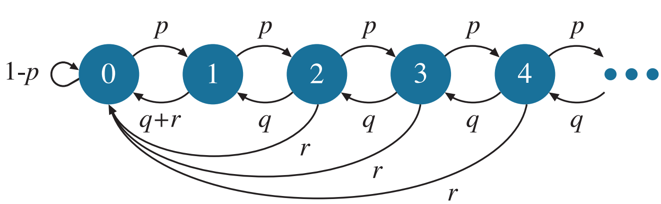

Figure 26.24 *DTMC for processor with failures.*

The DTMC in Figure 26.24 is used to model a processor with failures. The chain tracks the number of jobs in the system. At any time step, either the number of jobs increases by 1 (with probability p), or decreases by 1 (with probability q), or a processor failure occurs (with probability r), where $p + q + r = 1$. In the case of a processor failure, all jobs in the system are lost. Derive the limiting probability, π_i, of there being i jobs in the system.

You will want to use the z-transform approach that you learned in Chapter 6. Here are some steps to help remind you how this works:

(a) Write the balance equation for state 0. Now express π_1 in terms of π_0.

(b) Write the balance equations for state $i \geq 1$.

(c) Let $\hat{\Pi}(z) = \sum_{i=0}^{\infty} \pi_i z^i$. Derive an expression for $\hat{\Pi}(z)$ in terms of π_0. You should get

$$\hat{\Pi}(z) = \frac{\pi_0 - z\pi_0 - z\frac{r}{q}}{z^2\frac{p}{q} - z\frac{1}{q} + 1}.$$

(d) Rewrite $\hat{\Pi}(z)$ with its denominator factored into $\left(1 - \frac{z}{r_1}\right)\left(1 - \frac{z}{r_2}\right)$, where r_1 and r_2 are roots that you specify, where $r_1 < r_2$.

(e) Determine π_0. You will need three steps:

  (i) Explain why $\hat{\Pi}(z)$ is bounded for all $0 \leq z \leq 1$.

  (ii) Now show that $0 \leq r_1 < 1$.

  iii. We thus can conclude that $\hat{\Pi}(r_1) < \infty$. Thus, since r_1 is a root of the denominator of $\hat{\Pi}(z)$, it must also be a root of the numerator of $\hat{\Pi}(z)$. Use this to get π_0. [Note: Although you now have π_0, wait until the very end of the problem to substitute in this value.]

(f) Apply partial fraction decomposition to $\hat{\Pi}(z)$.

(g) $\hat{\Pi}(z)$ should now be very simple. Rewrite $\hat{\Pi}(z)$ as a geometric series.

(h) Match coefficients to get the π_i's.

(i) Verify that your solution for π_i satisfies the balance equations.