7. Referring to the Markov chain of problem 5, §16, use the methodology of Example 17.5 above to compute the expected number of steps required for the ant to return to vertex 1 if it starts at vertex 1. Square this with the invariant distribution you obtained in problem 5 above.

8. Prove Corollary 17.3.

9. Prove that if $\mathbf{P}_{ii} > 0$ for some state i then $\mathbf{P}^k_{ii} > 0$ for all $k \geq 1$.

10. Consider the Markov chain $(R_n : n \geq 0)$ of problem 7, §16. For what values of p $(0 \leq p \leq 1)$ is this Markov chain irreducible?

11. For those p for which R_n is irreducible, what is the invariant distribution?

12.* A Markov chain on the *infinite* state space $S = \{0, 1, 2, 3, \dots\}$ works as follows: $\mathbf{P}_{01} = 1$ and, for $n > 0$, $\mathbf{P}_{n,n+1} = \frac{1}{n+1}$ and $\mathbf{P}_{n,n-1} = \frac{n}{n+1}$. Find the invariant distribution for this Markov chain.

§18. Example: An Interacting Particle System

We did a lot of work in the last section — here we have some fun. The following beautiful example is called an *interacting particle system.* We start with a very simple case to illustrate the idea. Consider a grid with three cells and two indistinguishable particles. Each cell can be occupied by only one particle – which we'll call the 'single-occupancy rule'. A state (we'll also call it a *configuration*) consists of any way to allocate the two particles to the three cells. There are only three such configurations as shown in Figure 18.1, where they are labeled x, y, and z.

Figure 18.1. A three-state interacting particle system.

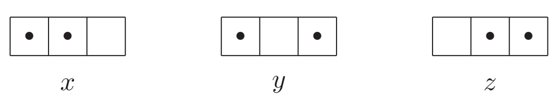

The transition rule is as follows. At each step of the Markov chain, pick a particle at random. (Here 'at random' means each possibility with equal likelihood and independently of any other choices.) Then pick a direction (east or west) at random. If the chosen particle can be moved one cell in the chosen direction, move it. Otherwise the configuration remains unchanged until (possibly) the next step. Suppose the current configuration is x and the rightmost particle is selected. If the direction is 'east', the transition is to configuration y; if the direction is 'west', the configuration remains unchanged due to the single-occupancy rule. If the leftmost particle is selected, then either direction leaves the configuration unchanged. A move to the west puts it off of the grid (not permitted); a move to the east violates single-occupancy. This shows that $\mathbf{P}_{xy} = \frac{1}{2} \cdot \frac{1}{2} = 1/4$ — the first $\frac{1}{2}$ is for picking the correct particle, the second $\frac{1}{2}$ is for picking the correct direction. Clearly $\mathbf{P}_{xz} = 0$, as z is not obtainable from

x by moving one particle one cell. The reader should verify that the PTM for this Markov chain is given by

$$\mathbf{P} \;=\; \begin{pmatrix} \mathbf{P}_{xx} & \mathbf{P}_{xy} & \mathbf{P}_{xz} \\ \mathbf{P}_{yx} & \mathbf{P}_{yy} & \mathbf{P}_{yz} \\ \mathbf{P}_{zx} & \mathbf{P}_{zy} & \mathbf{P}_{zz} \end{pmatrix} \;=\; \begin{pmatrix} 3/4 & 1/4 & 0 \\ 1/4 & 1/2 & 1/4 \\ 0 & 1/4 & 3/4 \end{pmatrix}.$$

Ergodicity is easily established as $\mathbf{P}^2$ has all positive entries.

Note that this PTM is *symmetric* in the sense that $\mathbf{P}_{ij} = \mathbf{P}_{ji}$ for any two states i and j. Since the rows of a PTM sum to 1 (always), if the PTM is symmetric the columns do as well. This makes the unique invariant distribution particularly easy to compute: $\boldsymbol{\mu} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. To verify that the second component of $\boldsymbol{\mu}\mathbf{P}$ is $\frac{1}{3}$, for example, we compute

$$(\boldsymbol{\mu}\mathbf{P})_2 \;=\; \frac{1}{3}\cdot\frac{1}{4} + \frac{1}{3}\cdot\frac{1}{2} + \frac{1}{3}\cdot\frac{1}{4} \;=\; \frac{1}{3}\cdot\left(\frac{1}{4}+\frac{1}{2}+\frac{1}{4}\right) \;=\; \frac{1}{3}. \checkmark$$

This holds in general. If the columns of the $s \times s$ PTM $\mathbf{P} = (\mathbf{P}_{ij})$ sum to one and $\boldsymbol{\mu} = (\,\frac{1}{s} \quad \frac{1}{s} \quad \cdots \quad \frac{1}{s}\,)$, then necessarily $\boldsymbol{\mu}\mathbf{P} = \boldsymbol{\mu}$ (see exercises).

We may generalize this interacting particle system to arbitrary dimension, grid, and number of particles. For example, consider a 4×4 2-dimensional grid with 4 particles. Again we impose a single-occupancy rule for the cells so there are $s = \binom{16}{4} = 1820$ possible particle configurations. We will consider any two configurations x and y to be *neighbors* if y can be obtained from x by moving one particle exactly one cell to the north, east, south, or west. Clearly this is a symmetric relation, as x can be obtained from y by reversing the move. If x and y are neighbors we write $x \sim y$. Two such neighboring configurations are shown in Figure 18.2 below.

Figure 18.2. Two neighboring configurations on the 4×4 *2-dimensional grid.*

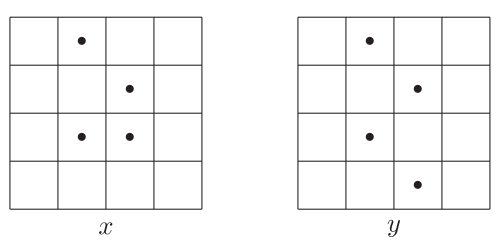

The Markov chain dynamics work as follows. At each step, pick one of the four particles randomly and one of the four directions (north, east, south, or west) randomly. If it is legitimate, move the chosen particle one cell in the chosen direction. If this move is not legitimate, leave the configuration unchanged. As before, a move is not legitimate if it would move the particle off of the grid or to an already occupied cell.

It is easy to see that this Markov chain satisfies the criteria for ergodicity of Theorem 17.17. Clearly we may proceed from any configuration to any other configuration by a sequence of single-particle moves, establishing irreducibility. (Although rigorously spelling out an algorithm to do this is harder than you might think.) Also, any configuration z with two adjacent particles has $\mathbf{P}_{zz} > 0$. For example, in Figure 18.2 we have $\mathbf{P}_{xx} > 0$ for this reason. (We also have $\mathbf{P}_{yy} > 0$. Why? — it has no adjacent particles.)

The PTM $\mathbf{P}$ for this process is a 1820×1820 matrix. Clearly we cannot write it down, let alone solve the equations $\boldsymbol{\mu}\mathbf{P} = \boldsymbol{\mu}$ to determine the invariant distribution! However we can deduce its invariant distribution. If $y \neq x$ and y is not a neighbor of x, we must have $\mathbf{P}_{xy} = 0$ by virtue of the system's dynamics. On the other hand, if $x \sim y$, to transition from x to y we must: (i) select the correct particle to move; and (ii) select the correct direction to move it. Hence $\mathbf{P}_{xy} = \frac{1}{4} \cdot \frac{1}{4} = \frac{1}{16}$ in this case. Summarizing we have:

$$\text{for } y \neq x, \quad \mathbf{P}_{xy} \;=\; \begin{cases} \frac{1}{16} & \text{if } x \sim y \\ 0 & \text{otherwise.} \end{cases} \tag{18.3}$$

Of course $\mathbf{P}_{xx}$ must get the balance of the probability:

$$\mathbf{P}_{xx} \;=\; 1 - \sum_{y:y\neq x} \mathbf{P}_{xy}, \tag{18.4}$$

since the rows of the PTM sum to 1. Equations (18.3) and (18.4) give formulas for the entries of the PTM $\mathbf{P}$ — equation (18.4) is for the diagonal entries of $\mathbf{P}$ and equation (18.3) is for the off-diagonal entries. The crucial point here is that $\mathbf{P}$ is a symmetric matrix, just as in the 1-dimensional 2-particle model. To show this we must show that if $x \neq y$ then $\mathbf{P}_{xy} = \mathbf{P}_{yx}$. This works because the neighbor relation $x \sim y$ is symmetric. If $x \sim y$ then $y \sim x$ and $\mathbf{P}_{xy} = \mathbf{P}_{yx} = \frac{1}{16}$. On the other hand, if $x \not\sim y$ then $y \not\sim x$ and $\mathbf{P}_{xy} = \mathbf{P}_{yx} = 0$. Since the PTM is symmetric, we know immediately that its invariant distribution is uniformly distributed on all 1820 configurations.

A Larger Version of the Same Model. Now for the fun part. Here we consider a 200×200 2-dimensional grid with 2500 particles. Originally they are confined to a 50×50 'container' as shown in the first frame of Figure 18.5. (The actual grid in Figure 18.5 cannot be shown as the entire picture would be black.) Then the top of the container is removed, but the sides are left in place. This is accomplished by imposing two vertical columns (1×50 each) of permanently occupied grid cells — these are discernible in the final configuration. The particles in these cells are not among the 2500 and are not subject to being moved. The Markov chain dynamics are exactly the same. At each step pick one of the 2500 movable particles and a direction at random. If the chosen particle can legitimately be moved one cell in the chosen direction, do so. Otherwise leave the configuration unchanged. (Here moving a particle to one of the 100 permanently occupied cells is illegitimate — so particles cannot pass through the walls of the container.)

Here the number of configurations is $\binom{40000-100}{2500} \approx 3.07 \times 10^{4056}$, which is a truly super-astronomical number. Furthermore the PTM $\mathbf{P}$ is a $(3.07 \times 10^{4056}) \times (3.07 \times 10^{4056})$ matrix! Yet we know from the reasoning above (which applies here as well) that $\mathbf{P}$ is a symmetric matrix. Thus the invariant distribution is uniformly distributed on all 3.07×10^{4056} states. Stated differently, in the long run the particles are uniformly distributed on the grid. A simulated realization of this Markov chain, shown at various stages in Figure 18.5, beautifully illustrates the convergence to the invariant distribution.

Figure 18.5. Simulated gas escaping from a container in a vacuum.

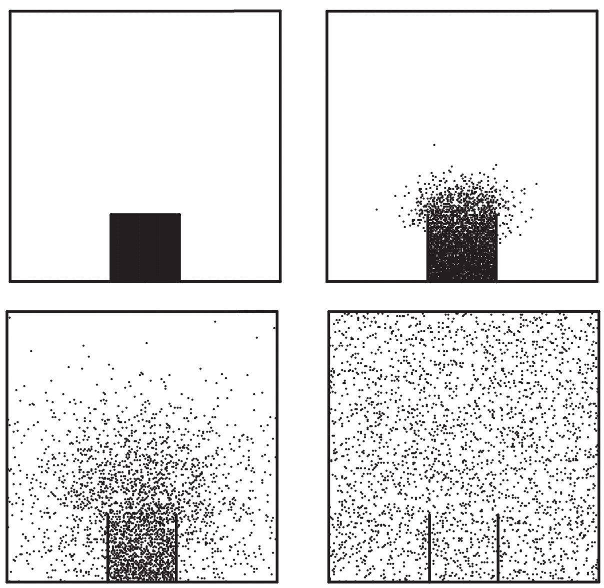

EXERCISES

1. Show that if the columns of the $s \times s$ PTM $\mathbf{P}$ sum to one and $\boldsymbol{\mu} = (\frac{1}{s} \quad \frac{1}{s} \quad \cdots \quad \frac{1}{s})$, then $\boldsymbol{\mu}\mathbf{P} = \boldsymbol{\mu}$. (A PTM where the columns also sum to one is called *doubly stochastic*.

2. Use the result in problem 1 above to compute the invariant distribution for R_n of problem 7, §16. Compare with your answer to problem 11, §17.

3. Jim does a 'random walk' on the face of a clock (a real clock, not a digital clock). Starting at 12:00, he repeatedly steps clockwise to the next hour with probability $\frac{2}{3}$ and counterclockwise to the previous hour with probability $\frac{1}{3}$.