Chapter 4. Irreducibility and Recurrence

Recurrence and Transience

Recall that the first return time in a set A is denoted τ_A^+ and the k-th passage time is denoted $\tau_A^{(k)}$, and that these are stopping times. The hitting probability is denoted $h_A(x)$. When $A = \{y\}$, we denote these functions without the curly brackets, e.g. $h_y(x) \equiv h_{\{y\}}(x)$.

We can also introduce the random variables

$$\mathcal{N}_t(x) := |\{0 < s \leq t \,:\, X_s = x\}| = \sup\{k \in \mathbb{N} \,:\, \tau_x^{(k)} \leq t\} \in \{0, \dots, t\}, \qquad \mathcal{N}(x) := \lim_{t \to \infty} \mathcal{N}_t(x) \in \mathbb{N} \cup \{\infty\}$$

as the number of times (up to time t and in total) that the process returns to the point x.

Lemma 0.1. *For $x \in S$, let $q_x := \mathbb{P}_x(\tau_x^+ < \infty)$.*

- *if $q_x = 1$, namely if there is a.s. return to x, then $\mathbb{P}_x(\mathcal{N}(x) = +\infty) = 1$.*
- *if $q_x < 1$, then, under $\mathbb{P}_x$, $\mathcal{N}(x)$ is a geometric random variable of parameter $1 - q_x$, namely $\mathbb{P}_x(\mathcal{N}(x) = k) = (1 - q_x)q_x^k$ for $k \geq 0$. In particular $\mathbb{P}_x(\mathcal{N}(x) < \infty) = 1$.*

In particular

$$\sum_{t=0}^{\infty} \mathbb{P}_x(X_t = x) = \mathbb{E}_x[\mathcal{N}(x)] = \begin{cases} +\infty & \text{if } q_x = 1 \\ 1/(1 - q_x) & \text{if } q_x < 1 \end{cases} \tag{1}$$

Proof. Fix $x \in S, k \geq 1$

$$q_x^{(k+1)} := \mathbb{P}_x(\tau_x^{(k+1)} < \infty) = \mathbb{P}_x(\tau_x^{(k+1)} < \infty \mid \tau_x^{(k)} < \infty)\mathbb{P}_x(\tau_x^{(k)} < \infty) = \mathbb{P}_x(\tau_x^+ < \infty)\mathbb{P}_x(\tau_x^{(k)} < \infty) = q_x \, q_x^{(k)}$$

where in the first equality we used $\{\tau_x^{(k+1)} < \infty\} \subset \{\tau_x^{(k)} < \infty\}$, while in the second equality we used the Strong Markov Property with $F = \mathbf{1}_{\tau_x^+ < \infty}$, $\mu = \delta_x$ and $\sigma = \tau_x^{(k)}$.

In particular, since $q_x^{(0)} = 1$, $q_x^{(k)} = q_x^k$, and thus

$$\mathbb{P}_x(\mathcal{N}(x) = +\infty) = \mathbb{P}_x(\cap_k \{\tau_x^{(k)} < \infty\}) = \lim_k \mathbb{P}_x(\tau_x^{(k)} < \infty) = \begin{cases} 1 & \text{if } q_x = 1 \\ 0 & \text{if } q_x < 1 \end{cases}$$

This concludes if $q_x = 1$. If $q_x < 1$

$$\mathbb{P}_x(\mathcal{N}(x) = k) = \mathbb{P}_x(\tau_x^{(k+1)} = \infty, \tau_x^{(k)} < \infty) = \mathbb{P}_x(\tau_x^{(k)} < \infty) - \mathbb{P}_x(\tau_x^{(k+1)} < \infty) = q_x^k - q_x^{k+1}$$

we necessarily have $\sum_{k \geq 1} q_x^{(k)} = 1$ and thus $c_x = 1 - q_x$.

We are left to check Equation 1. Indeed by monotone convergence

$$\sum_{t=0}^{\infty} \mathbb{P}_x(X_t = x) = \sum_{t=0}^{\infty} \mathbb{E}_x[\mathbf{1}_x(X_t)] = \mathbb{E}_x[\sum_{t=0}^{\infty} \mathbf{1}_x(X_t)] = \mathbb{E}_x[\mathcal{N}(x)]$$

$\square$

Definition 0.1 (Recurrent and transient states)**.** An element $x \in S$ is **recurrent** if $\mathbb{P}_x(\tau_x^+ < \infty) = 1$, namely (in view of Lemma 0.1) if the Markov chain recurs on x infinitely many times a.s..

$x \in S$ is **transient** if $\mathbb{P}_x(\tau_x^+ < \infty) < 1$, namely if the chain only passes on x finitely many times a.s..

Communicating classes

Given a time-homogeneous Markov chain with transition probabilities $(p_{x,y})$, for $x, y \in S$ we write $x \to y$ if any of the equivalent conditions is satisfied

- There exists $t \geq 0$ such that $p_{x,y}^{(t)} > 0$.
- There exists a path in $D(S)$ starting in x and passing through y.
- There exists $t \geq 0$, $x_0 = x, x_1, \dots, x_t = y$ such that $p_{x_{i-1},x_i} > 0$ for $i = 1, \dots, t$.
- $\mathbb{P}_x(X_t = y \text{ for some } t \geq 0) > 0$.

If $x \to y$ and $y \to x$, we write $x \leftrightarrow y$. It is easy to verify that $\leftrightarrow$ is an equivalence relation. The equivalence classes are given a name that matches the intuition well.

Definition 0.2 (Communicating classes)**.** The equivalence classes $S/\leftrightarrow$ are called the **communicating classes** of the Markov chain.

Consider the set $S_x := \{y \in S \; : \; x \to y\}$ of points reachable from x. If $x \to y$, then $S_x \supset S_y$, so that if $x \leftrightarrow y$ then $S_x = S_y$.

Definition 0.3 (Closed class)**.** A communicating class C is **closed** if any of the following equivalent conditions holds:

- $S_x = C$ for all $x \in C$.
- $S_x = C$ for some $x \in C$.
- $p_{x,y} = 0$ for all $x \in C$ and $y \notin C$.
- $\mathbb{P}_x(X_t \in C, \; \forall t \geq 0) = 1$ for all $x \in C$.

Remark. If there are finitely many communicating classes (in particular if S is finite), at least one of them is closed. Indeed let C_0 be any communicating class. If it is closed we are done, otherwise there exists $x_0 \in C_0, x_1 \in C_1 \neq C_0$ such that $x_0 \to x_1$ but $x_1 \nrightarrow x_0$. Iterating, we find a sequence of distinct classes $C_0, C_1, \dots, C_n$. This sequence is finite (since we assumed there are finitely many classes), and the last entry C_n is easily seen to be closed.

On the other hand, if there are infinitely many classes, a closed class may fail to exist. E.g. $S = \mathbb{N}$ and $p_{n,n+1} = 1$. In this case each singleton is a (non-closed) communicating class.

Exercise 0.1. In the following directed graph, the arrows corresponding to strictly positive transition probabilities are drawn. Find the corresponding communicating classes and identify the closed ones.

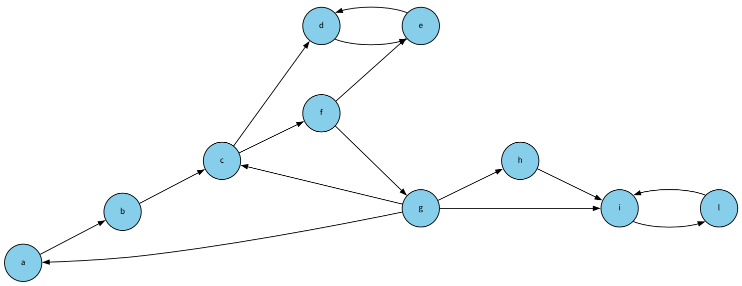

Figure 1

💡 Solution

The classes are $\{a, b, c, f, g\}$, $\{d, e\}$, $\{h\}$, $\{i, l\}$. $\{d, e\}$ and $\{i, l\}$ are closed.

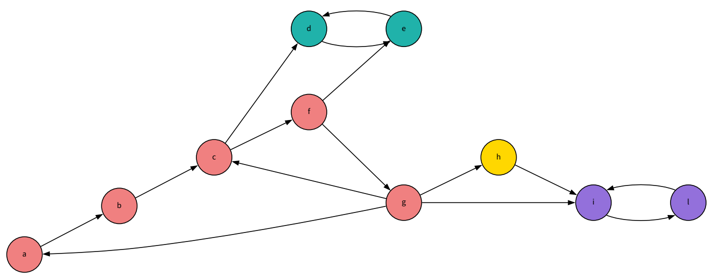

Figure 2

Recurrence of Communicating Classes

There is a strict relationship between recurrence and closure of a class.

Theorem 0.1 (Recurrence and Transience for communicating classes)**.** *Let C be a communicating class and recall $h_{C^c}(x) = \mathbb{P}_x(\tau_{C^c} < \infty)$.*

a. *If C is **not closed**, $h_{C^c}(x) > 0$ for all $x \in C$. In particular every point $x \in C$ is transient.*
b. *If $\inf_{x \in C} h_{C^c}(x) > 0$ then $\mathbb{P}_x(\tau_{C^c} < \infty) = 1$. In particular a Markov chain will eventually leave any **finite and not closed** class.*
c. *If C is **closed**, then $h_{C^c}(x) = 0$. And either every point $x \in C$ is recurrent, either every point $x \in C$ is transient. Thus recurrence and transience if a property of a closed class C, not each single point in C.*
d. *If C is **finite and closed**, then it is a recurrent class.*

Proof. Let's prove point by point.

a. Fix $x \in C$. If C is not closed, then there exists $y \in C^c$, $t > 0$ such that $p^{(t)}_{x,y} > 0$. However, since $y \nrightarrow C$

$$\mathbb{P}_x(\tau^c_C < \infty) \leq 1 - \mathbb{P}_x(X_t = y) = 1 - p^{(t)}_{x,y} < 1$$

b. Let $r(x) = 1 - h_{C^c}(x)$ be the probability that the chain never leaves the class C when starting from x, we are assuming $R := \sup_{x \in C} r(x) < 1$, and want to prove $r(x) \equiv 0$. Since h_{C^c} solves a linear problem, r solves

$$\begin{cases} (I - P)r = 0 & \text{on } C \\ r = 0 & \text{on } C^c \end{cases}$$

Now notice that, since $r(x) = 0$ outside C and $p_{x,y} = 0$ if $x \notin C$ and $y \in C$, $r = Pr$ holds everywhere on S. Iterating $r = P^t r$, namely for $t \geq 0$

$$r(x) = \sum_{y \in S} p^{(t)}_{x,y} r(y) = \sum_{y \in C} p^{(t)}_{x,y} r(y) \leq R \sum_{y \in C} p^{(t)}_{x,y}$$

where we used $r(y) = 0$ for $y \notin C$ in the second equality. However as $t \to \infty$, since probabilities are continuous on increasing sets

$$\sum_{y \in C} p^{(t)}_{x,y} = \mathbb{P}_x(\tau_{C^c} > t) \downarrow \mathbb{P}_x(\tau_{C^c} = \infty) = r(x)$$

Therefore we get, taking the limit $r(x) \leq Rr(x)$, and thus $r(x) = 0$ since $R < 1$.

c. If C is closed, then $r(x) = 0$ by the last point of Definition 0.3. Let $x, y \in C$, and assume that x is recurrent. It is enough to show that then also y is recurrent. Indeed, since $x \leftarrow y$, there exists $s, t > 0$ such that $p^{(s)}_{y,x}, p^{(t)}_{x,y} > 0$. However, the probability of not coming back to y, is smaller than the probability of reaching x in s steps, and not reaching y from x is t step each time we visit x:

$$\mathbb{P}_y(\tau^+_y \leq \tau^{(k)}_x + t + s) \geq 1 - p^{(s)}_{y,x}(1 - p^{(t)}_{x,y})^k \tag{2}$$

Since x is recurrent, $\tau^{(k)}_x < \infty$ a.s., thus $\{\tau^+_y < \infty\} = \cup_k \{\tau^+_y \leq \tau^{(k)}_x + t + s\}$. And therefore,

$$\mathbb{P}_y(\tau^+_y < \infty) = \lim_k \mathbb{P}_y(\tau^+_y \leq \tau^{(k)}_x + t + s) \geq 1$$

d. Since C is finite, at least one point recurs infinitely many times. And from the previous point, each does.

□

Exercise 0.2. Give an example of a Markov chain that features a non-closed class C such that $\mathbb{P}_x(\tau_{C^c} < \infty) > 0$ for all $x \in C$.

Exercise 0.3. Use the Strong Markov Property to give a detailed proof of Equation 2.

A Markov chain will eventually leave every finite non-closed class. Thus, at least on finite state spaces, it will eventually enter some closed class and remain there forever. In the previous chapter we have seen how to compute the probability that the chain enters a given set, in particular a given closed class. The next natural step is to understand the behavior within a closed class. To this aim, it is enough to study a Markov chain defined on a closed class, since closed classes constitute the independent building blocks of the chain.

Definition 0.4 (Irreducible Markov Chain)**.** A Markov chain is **irreducible** if its state space is a closed communicating class, namely if $x \leftrightarrow y$ for all $x, y \in S$.

Most of the times, we will focus on irreducible, homogeneous chains hereafter.

Recurrence on $\mathbb{Z}^d$

Thanks to Lemma 0.1, we have for an irreducible chain the following criterion:

An irreducible chain is recurrent iff $\sum_{t=0}^{\infty} \mathbb{P}_x(X_t = x) = \infty$ for some (iff for any) $x \in S$.

Before proceeding, recall the Stirling formula

$$n^n e^{-n}\sqrt{2\pi n}e^{1/(12n+1)} \leq n! \leq n^n e^{-n}\sqrt{2\pi n}e^{1/(12n)} \tag{3}$$

Example 0.1. Consider the Markov chain on $\mathbb{Z}$ where at each time-step X_t can increase by 1 with probability p and decrease with probability $1-p$. In other words, $p_{x,x+1} = 1 - p_{x,x-1} = p$. Then the chain

- is transient if $p \neq 1/2$.
- is recurrent if $p = 1/2$.

Indeed, the chain is irreducible, and we can start at $X_0 = 0$ for simplicity of notation. X_t can be 0 with positive probability iff t is even, say $t = 2n$. Then X_t will be 0 iff exactly n times we increased by 1 and n times we decreased by 1. Therefore by Equation 3

$$\mathbb{P}_0(X_{2n} = 0) = \binom{2n}{n} p^n (1-p)^n = \frac{4^n p^n (1-p)^n}{\sqrt{\pi n}}(1 + o_n(1))$$

Now $4^n p^n (1-p)^n < 1$ for $p \neq 1/2$, the series converges in this case. If $p = 1/2$ however, $\mathbb{P}_0(X_{2n} = 0) = \frac{1}{\sqrt{\pi n}}(1 + o_n(1))$ and thus the series diverges.

Remark. Consider a Markov Chain on $\mathbb{Z}^d$ where at each step we add some random vector in $\mathbb{Z}^d$:

$$X_t = x + \sum_{s=1}^{t} Z_s$$

where Z_s are i.i.d.. Assume that

$$\mathbb{E}[|Z_1|] < \infty, \mathbb{E}[Z_1] := m \neq 0$$

Then the random walk **X** is transient. Indeed, by the law of large numbers

$$\mathbb{P}(\lim_t X_t/t = m) = 1$$

Therefore with probability 1, X_t only visits each point in $\mathbb{Z}^d$ finitely many times.

If S is a countable undirected graph, we write $x \sim y$ to mean that x and y are neighbors. The degree $d(x)$ of a point x is the number of neighbors of x. A graph is called locally finite if $d(x)$ is finite for all $x \in S$.

Definition 0.5. Let S be a locally finite graph. The **simple random walk** on S, is a random walk with transition probabilities

$$p_{x,y} = \begin{cases} 1/d(x) & \text{if } y \sim x \\ 0 & \text{otherwise} \end{cases}$$

For instance the simple random walk on $\mathbb{Z}^d$ moves at each step with probability $1/2d$ on each of the $2d$ neighbors of X_t. For instance if $d = 2$, at each step we can add $(1,0)$, $(-1,0)$, $(0,1)$ or $(0,-1)$, each with probability $1/4$. The previous remark does not help us to decide if this random walk is recurrent or not, since the displacement expected value m vanishes. However we have already seen that for $d = 1$ the walk is recurrent.

Proposition 0.1. *The simple random walk on $\mathbb{Z}^d$ is*

- *recurrent if $d = 1, 2$.*
- *transient if $d \geq 3$.*

Or, as Pólya put it, a drunk person will eventually find their way home, but a drunk bird may be lost forever.

Proof. The proof is quite simple indeed. Again we can consider the case $X_0 = 0$ for simplicity of notation, since by translation invariance $\mathbb{P}_x(X_t = x)$ does not depend on x (or in any case transience and recurrence does not depend on x for irreducible walks). To return to 0 at time t, we certainly need $t = 2n$, moreover we should have that in each direction we have added $+1$ as many times as we removed -1. Thus

$$\mathbb{P}_0(X_{2n} = 0) = \sum_{k1+k2+\ldots+k_d=n} \binom{2n}{k_1, k_1, k_2, k_2, \ldots, k_d, k_d}(2d)^{-2n} = \sum_{k1+k2+\ldots+k_d=n} \frac{(2k_1 + 2k_2 + \ldots + 2k_d)!}{(k_1!)^2 \cdots (k_d!)^2}(2d)^{-2n}$$

We have already considered the case $d = 1$ in Example 0.1.

For $d = 2$, we can use a combinatorial identity $\sum_{k=0}^{n} \left(\binom{n}{k}\right)^2 = \binom{2n}{n}$ to get

$$\mathbb{P}_0(X_{2n} = 0) = \sum_{k=0}^{n} \frac{(2n)!}{(k!)^2((n-k)!)^2}(4)^{-2n} = 4^{-2n}\binom{2n}{n}\sum_{k=0}^{n}\left(\binom{n}{k}\right)^2 = \left(\binom{2n}{n}4^{-n}\right)^2$$

This is exactly the square of the one-dimensional computation, so that

$$\mathbb{P}_0(X_{2n} = 0) = \frac{1}{\pi n}(1 + o_n(1))$$

and therefore the series diverges, and the walk is recurrent.

For $d = 3$, this simple power rule fails (namely the probability is not an exact power of the one-dimensional formula). Yet a Stirlingation gives $\mathbb{P}_0(X_{2n} = 0) = cn^{-3/2}(1 + o_n(1))$ and thus the series converges and the walk is transient.

For $d \geq 4$, it already becomes cumbersome to use Stirling, although it is possible to prove $\mathbb{P}_0(X_{2n} = 0) = c_d n^{-d/2}(1 + o_d(1))$. But we do not need such a computation. Indeed, it is easy to see that the higher the dimension, the more the walk is transient, so to speak. Let X_t be the simple random walk in $\mathbb{Z}^4$. If we only look at the first three components of X_t, and only consider the times when it moves, we get a sequence Y_t which is exactly a 3-dimensional simple random walk. Since Y_t passes finitely many times at 0, so does X_t. The same argument works for all the higher dimensions. □

Exercise 0.4. Consider an irreducible, transient random walk with $p_{x,y} = p_{y,x}$. Prove that

$$d(x, y) = -\log \mathbb{P}_x(\tau_y < \infty)$$

is a distance on S.

Abstraction

We can mention some more advanced results.
A. Consider a random walk in $\mathbb{Z}^d$, $X_t = \sum_{s=1}^t Z_s$ with (Z_s) an i.i.d. sequence in $\mathbb{Z}^d$ such that $\mathbb{P}(Z_t = x) = \mathbb{P}(Z_t = -x) = \Theta(\|x\|^{-s})$ (meaning that this probability decays as $\|x\|^{-s}$ for $\|x\|$ large). Then the walk is recurrent iff $d = 1, 2$ and $s < d$ or $s > 2d$.
B. Consider a random walk on a finitely generated group S such that

- It is group invariant: $p_{x,y} = p_{gx,gy}$ for all $g \in S$.
- It has finite second moment: $\sum_x d(e,x)^2 p_{e,x} < \infty$ (it is easily seen that this condition does not depend on the finite set of generators used to define the Cayley distance $d(x,y)$). Then the random walk is recurrent iff the group S is finite or is virtually isomorphic to $\mathbb{Z}$ or $\mathbb{Z}^2$.