8

Transience and Recurrence

We have developed substantial structural results for ψ-irreducible Markov chains in Part I of this book. Part II is devoted to stability results of ever-increasing strength for such chains.

In Chapter 1, we discussed in a heuristic manner two possible approaches to the stability of Markov chains. The first of these discussed basic ideas of stability and instability, formulated in terms of recurrence and transience for ψ-irreducible Markov chains. The aim of this chapter is to formalize those ideas.

In many ways it is easier to tell when a Markov chain is unstable than when it is stable: it fails to return to its starting point, it eventually leaves any "bounded" set with probability one, it returns only a finite number of times to a given set of "reasonable size". Stable chains are then conceived of as those which do not vanish from their starting points in at least some of these ways. There are many ways in which stability may occur, ranging from weak "expected return to origin" properties, to convergence of all sample paths to a single point, as in global asymptotic stability for deterministic processes. In this chapter we concentrate on rather weak forms of stability, or conversely on strong forms of instability.

Our focus is on the behavior of the occupation time random variable $\eta_A := \sum_{n=1}^{\infty} \mathbb{1}\{\Phi_n \in A\}$ which counts the number of visits to a set A. In terms of η_A we study the stability of a chain through the transience and recurrence of its sets.

Uniform Transience and Recurrence

The set A is called *uniformly transient* if for there exists $M < \infty$ such that $\mathsf{E}_x[\eta_A] \le M$ for all $x \in A$.

The set A is called *recurrent* if $\mathsf{E}_x[\eta_A] = \infty$ for all $x \in A$.

The highlight of this approach is a solidarity, or dichotomy, theorem of surprising strength.

Theorem 8.0.1 *Suppose that Φ is ψ-irreducible. Then either*

(i) *every set in $\mathcal{B}^+(\mathsf{X})$ is recurrent, in which case we call Φ recurrent; or*

(ii) *there is a countable cover of X with uniformly transient sets, in which case we call Φ transient; and every petite set is uniformly transient.*

PROOF This result is proved through a splitting approach in Section 8.2.3. We also give a different proof, not using splitting, in Theorem 8.3.4, where the cover with uniformly transient sets is made more explicit, leading to Theorem 8.3.5 where all petite sets are shown to be uniformly transient if there is just one petite set in $\mathcal{B}^+(\mathsf{X})$ which is not recurrent. □

The other high point of this chapter is the first development of one of the themes of the book: the existence of so-called *drift criteria*, couched in terms of the expected change, or drift, defined by the one-step transition function P, for chains to be stable or unstable in the various ways this is defined.

Drift for Markov Chains

The (possibly extended valued) drift operator Δ is defined for any non-negative measurable function V by

$$\Delta V(x) := \int P(x,dy)V(y) - V(x), \qquad x \in \mathsf{X}. \tag{8.1}$$

A second goal of this chapter is the development of criteria based on the drift function for both transience and recurrence.

Theorem 8.0.2 *Suppose Φ is a ψ-irreducible chain.*

(i) *The chain Φ is transient if and only if there exists a bounded non-negative function V and a set $C \in \mathcal{B}^+(\mathsf{X})$ such that for all $x \in C^c$,*

$$\Delta V(x) \ge 0 \tag{8.2}$$

and

$$D = \{V(x) > \sup_{y \in C} V(y)\} \in \mathcal{B}^+(\mathsf{X}). \tag{8.3}$$

(ii) *The chain Φ is recurrent if there exists a petite set $C \subset \mathsf{X}$, and a function V which is unbounded off petite sets in the sense that $C_V(n) := \{y : V(y) \le n\}$ is petite for all n, such that*

$$\Delta V(x) \le 0, \qquad x \in C^c. \tag{8.4}$$

PROOF The drift criterion for transience is proved in Theorem 8.4.2, whilst the condition for recurrence is in Theorem 8.4.3. □

Such conditions were developed by Lyapunov as criteria for stability in deterministic systems, by Khas'minskii and others for stochastic differential equations [134, 149], and by Foster as criteria for stability for Markov chains on a countable space: Theorem 8.0.2 is originally due (for countable spaces) to Foster [82] in essentially the form given above.

There is in fact a converse to Theorem 8.0.2 (ii) also, but only for ψ-irreducible Feller chains (which include all countable space chains): we prove this in Section 9.4.2. It is not known whether a converse holds in general.

Recurrence is also often phrased in terms of the hitting time variables $\tau_A = \inf\{k \geq 1 : \Phi_k \in A\}$, with "recurrence" for a set A being defined by $L(x, A) = \mathsf{P}_x(\tau_A < \infty) = 1$ for all $x \in A$. The connections between this condition and recurrence as we have defined it above are simple in the countable state space case: the conditions are in fact equivalent when A is an atom. In general spaces we do not have such complete equivalence. Recurrence properties in terms of τ_A (which we call Harris recurrence properties) are much deeper and we devote much of the next chapter to them. In this chapter we do however give some of the simpler connections: for example, if $L(x, A) = 1$ for all $x \in A$ then $\eta_A = \infty$ a.s. when $\Phi_0 \in A$, and hence A is recurrent (see Proposition 8.3.1).

8.1 Classifying chains on countable spaces

8.1.1 The countable recurrence/transience dichotomy

We turn as before to the countable space to guide and motivate our general results, and to aid in their interpretation.

When $\mathsf{X} = \mathbb{Z}_+$, we initially consider the stability of an *individual state* $\boldsymbol{\alpha}$. This will lead to a global classification for irreducible chains.

The first, and weakest, stability property involves the *expected number of visits* to $\boldsymbol{\alpha}$. The random variable $\eta_\alpha = \sum_{n=1}^\infty \mathbb{1}\{\Phi_n = \boldsymbol{\alpha}\}$ has been defined in Section 3.4.3 as the number of visits by $\boldsymbol{\Phi}$ to $\boldsymbol{\alpha}$: clearly η_α is a measurable function from Ω to $\mathbb{Z}_+ \cup \{\infty\}$.

Classification of States

The state $\boldsymbol{\alpha}$ is called *transient* if $\mathsf{E}_\alpha(\eta_\alpha) < \infty$, and *recurrent* if $\mathsf{E}_\alpha(\eta_\alpha) = \infty$.

From the definition $U(x, y) = \sum_{n=1}^\infty P^n(x, y)$ we have immediately that for any states $x, y \in \mathsf{X}$

$$\mathsf{E}_x[\eta_y] = U(x, y). \tag{8.5}$$

The following result gives a structural dichotomy which enables us to consider, not just the stability of states, but of chains as a whole.

Proposition 8.1.1 *When* X *is countable and* $\boldsymbol{\Phi}$ *is irreducible, either* $U(x, y) = \infty$ *for all* $x, y \in \mathsf{X}$ *or* $U(x, y) < \infty$ *for all* $x, y \in \mathsf{X}$.

PROOF This relies on the definition of irreducibility through the relation $\leftrightarrow$.

If $\sum_n P^n(x, y) = \infty$ for some x, y, then since $u \to x$ and $y \to v$ for any u, v, we have r, s such that $P^r(u, x) > 0$, $P^s(y, v) > 0$ and so

$$\sum_n P^{r+s+n}(u, v) > P^r(u, x)\Big[\sum_n P^n(x, y)\Big] P^s(y, v) = \infty. \tag{8.6}$$

Hence the series $U(x, y)$ and $U(u, v)$ all converge or diverge simultaneously, and the result is proved. □

Now we can extend these stability concepts for states to the whole chain.

Transient and recurrent chains

If every state is transient the chain itself is called *transient*.

If every state is recurrent, the chain is called *recurrent*.

The solidarity results of Proposition 8.1.3 and Proposition 8.1.1 enable us to classify irreducible chains by the property possessed by one and then all states.

Theorem 8.1.2 *When* $\boldsymbol{\Phi}$ *is irreducible, then either* $\boldsymbol{\Phi}$ *is transient or* $\boldsymbol{\Phi}$ *is recurrent.* □

We can say, in the countable case, exactly what recurrence or transience means in terms of the return time probabilities $L(x, x)$. In order to connect these concepts, for a fixed n consider the event $\{\Phi_n = \boldsymbol{\alpha}\}$, and decompose this event over the mutually exclusive events $\{\Phi_n = \boldsymbol{\alpha}, \tau_\alpha = j\}$ for $j = 1, \ldots, n$. Since $\boldsymbol{\Phi}$ is a Markov chain, this provides the first-entrance decomposition of P^n given for $n \geq 1$ by

$$P^n(x, \boldsymbol{\alpha}) = \mathsf{P}_x\{\tau_\alpha = n\} + \sum_{j=1}^{n-1} \mathsf{P}_x\{\tau_\alpha = j\} P^{n-j}(\boldsymbol{\alpha}, \boldsymbol{\alpha}). \tag{8.7}$$

If we introduce the generating functions for the series P^n and ${}_\alpha P^n$ as

$$U^{(z)}(x, \boldsymbol{\alpha}) := \sum_{n=1}^{\infty} P^n(x, \boldsymbol{\alpha}) z^n, \qquad |z| < 1 \tag{8.8}$$

$$L^{(z)}(x, \boldsymbol{\alpha}) := \sum_{n=1}^{\infty} \mathsf{P}_x(\tau_\alpha = n) z^n, \qquad |z| < 1 \tag{8.9}$$

then multiplying (8.7) by z^n and summing from $n = 1$ to ∞ gives for $|z| < 1$

$$U^{(z)}(x, \boldsymbol{\alpha}) = L^{(z)}(x, \boldsymbol{\alpha}) + L^{(z)}(x, \boldsymbol{\alpha}) U^{(z)}(\boldsymbol{\alpha}, \boldsymbol{\alpha}). \tag{8.10}$$

From this identity we have

Proposition 8.1.3 *For any $x \in \mathsf{X}$, $U(x, x) = \infty$ if and only if $L(x, x) = 1$.*

PROOF Consider the first entrance decomposition in (8.10) with $x = \boldsymbol{\alpha}$: this gives

$$U^{(z)}(\boldsymbol{\alpha}, \boldsymbol{\alpha}) = L^{(z)}(\boldsymbol{\alpha}, \boldsymbol{\alpha}) \Big/ \Big[1 - L^{(z)}(\boldsymbol{\alpha}, \boldsymbol{\alpha})\Big]\,. \tag{8.11}$$

Letting $z \uparrow 1$ in (8.11) shows that

$$L(\boldsymbol{\alpha}, \boldsymbol{\alpha}) = 1 \Longleftrightarrow U(\boldsymbol{\alpha}, \boldsymbol{\alpha}) = \infty.$$

□

This gives the following interpretation of the transience/recurrence dichotomy of Proposition 8.1.1.

Proposition 8.1.4 *When $\boldsymbol{\Phi}$ is irreducible, either $L(x, y) = 1$ for all $x, y \in \mathsf{X}$ or $L(x, x) < 1$ for all $x \in \mathsf{X}$.*

PROOF From Proposition 8.1.3 and Proposition 8.1.1, we have $L(x, x) < 1$ for all x or $L(x, x) = 1$ for all x. Suppose in the latter case, we have $L(x, y) < 1$ for some pair x, y: by irreducibility, $U(y, x) > 0$ and thus for some n we have $\mathsf{P}_y(\Phi_n = x, \tau_y > n) > 0$, from which we have $L(y, y) < 1$, which is a contradiction. □

In Chapter 9 we will define *Harris recurrence* as the property that $L(x, A) \equiv 1$ for all $x \in A$ and $A \in \mathcal{B}^+(\mathsf{X})$: for countable chains, we have thus shown that recurrent chains are also Harris recurrent, a theme we return to in the next chapter when we explore stability in terms of $L(x, A)$ in more detail.

8.1.2 Specific models: evaluating transience and recurrence

Calculating the quantities $U(x, y)$ or $L(x, x)$ directly for specific models is non-trivial except in the simplest of cases. However, we give as examples two simple models for which this is possible, and then a deeper proof of a result for general random walk.

Renewal processes and forward recurrence time chains Let the transition matrix of the forward recurrence time chain be given as in Section 3.3. Then it is straightforward to see that for all states $n > 1$,

$${}_1P^{n-1}(n, 1) = 1.$$

This gives

$$L(1, 1) = \sum_{n \geq 1} p(n)\, {}_1P^{n-1}(n, 1) = 1$$

also. Hence the forward recurrence time chain is always recurrent if p is a proper distribution.

The calculation in the proof of Proposition 8.1.3 is actually a special case of the use of the *renewal equation*. Let Z_n be a renewal process with increment distribution p as defined in Section 2.4. By breaking up the event $\{Z_k = n\}$ over the last time before n that a renewal occurred we have

$$u(n) := \sum_{k=0}^{\infty} \mathsf{P}(Z_k = n) = 1 + u * p(n)$$

and multiplying by z^n and summing over n gives the form

$$U(z) = [1 - P(z)]^{-1} \tag{8.12}$$

where $U(z) := \sum_{n=0}^{\infty} u(n) z^n$ and $P(z) := \sum_{n=0}^{\infty} p(n) z^n$.

Hence a renewal process is also called recurrent if p is a proper distribution, and in this case $U(1) = \infty$.

Notice that the renewal equation (8.12) is identical to (8.11) in the case of the specific renewal chain given by the return time $\tau_\alpha(n)$ to the state $\boldsymbol{\alpha}$.

Simple random walk on $\mathbb{Z}_+$ Let P be the transition matrix of random walk on a half line in the simplest irreducible case, namely $P(0,0) = p$ and

$$\begin{array}{rcll} P(x, x-1) & = & p, & x > 0 \\ P(x, x+1) & = & q, & x \geq 0. \end{array} \tag{8.13}$$

where $p + q = 1$. This is known as the simple, or Bernoulli, random walk.

We have that

$$\begin{array}{rcl} L(0,0) & = & p + qL(1,0), \\ L(1,0) & = & p + qL(2,0). \end{array} \tag{8.14}$$

Now we use two tricks specific to chains such as this. Firstly, since the chain is skip-free to the left, it must reach $\{0\}$ from $\{2\}$ only by going through $\{1\}$, so that we have

$$L(2,0) = L(2,1)L(1,0).$$

Secondly, the translation invariance of the chain, which implies $L(j, j-1) = L(1,0), j \geq 1$, gives us

$$L(2,0) = [L(1,0)]^2.$$

Thus from (8.14), we find that

$$L(1,0) = p + q[L(1,0)]^2 \tag{8.15}$$

so that $L(1,0) = 1$ or $L(1,0) = p/q$.

This shows that $L(1,0) = 1$ if $p \geq q$, and from (8.14) we derive the well-known result that $L(0,0) = 1$ if $p \geq q$.

Random walk on $\mathbb{Z}$ In order to classify general random walk on the integers we will use the laws of large numbers. Proving these is outside the scope of this book: see, for example, Billingsley [25] or Chung [50] for these results.

Suppose that Φ_n is a random walk such that the increment distribution Γ has a mean which is zero. The form of the Weak Law of Large Numbers that we will use can be stated in our notation as

$$P^n(0, A(\varepsilon n)) \to 1 \tag{8.16}$$

for any ε, where the set $A(k) = \{y : |y| \leq k\}$. From this we prove

Theorem 8.1.5 *If Φ is an irreducible random walk on $\mathbb{Z}$ whose increment distribution Γ has mean zero, then Φ is recurrent.*

PROOF First note that from (8.7) we have for any x

$$\begin{aligned}\sum_{m=1}^N P^m(x,0) &= \sum_{k=1}^N \sum_{j=0}^k \mathsf{P}_x(\tau_0 = k-j)P^j(0,0)\\ &= \sum_{j=0}^N P^j(0,0)\sum_{i=0}^{N-j} \mathsf{P}_x(\tau_0 = i)\\ &\le \sum_{j=0}^N P^j(0,0).\end{aligned} \tag{8.17}$$

Now using this with the symmetry that $\sum_{m=1}^N P^m(x,0) = \sum_{m=1}^N P^m(0,-x)$ gives

$$\begin{aligned}\sum_{m=0}^N P^m(0,0) &\ge [2M+1]^{-1}\sum_{|x|\le M}\sum_{j=0}^N P^j(0,x)\\ &\ge [2M+1]^{-1}\sum_{j=0}^N P^j(0,A(jM/N))\\ &= [2aN+1]^{-1}\sum_{j=0}^N P^j(0,A(aj))\end{aligned} \tag{8.18}$$

where we choose $M = Na$ where a is to be chosen later.

But now from the Weak Law of Large Numbers (8.16) we have

$$P^k(0,A(ak)) \to 1$$

as $k \to \infty$; and so from (8.18) we have

$$\begin{aligned}\liminf_{N\to\infty}\sum_{m=0}^N P^m(x,0) &\ge \liminf_{N\to\infty}[2aN+1]^{-1}\sum_{j=0}^N P^j(0,A(aj))\\ &= [2a]^{-1}.\end{aligned} \tag{8.19}$$

Since a can be chosen arbitrarily small, we have $U(0,0) = \infty$ and the chain is recurrent. □

This proof clearly uses special properties of random walk. If Γ has simpler structure then we shall see that simpler procedures give recurrence in Section 8.4.3.

8.2 Classifying ψ-irreducible chains

The countable case provides guidelines for us to develop solidarity properties of chains which admit a single atom rather than a multiplicity of atoms. These ideas can then be applied to the split chain and carried over through the m-skeleton to the original chain, and this is the agenda in this section.

In order to accomplish this, we need to describe precisely what we mean by recurrence or transience of sets in a general space.

8.2.1 Transience and recurrence for individual sets

For general $A, B \in \mathcal{B}(\mathsf{X})$ recall from Section 3.4.3 the taboo probabilities given by

$${}_AP^n(x,B) = \mathsf{P}_x\{\Phi_n \in B, \tau_A \ge n\},$$

and by convention we set ${}_AP^0(x,A) = 0$. Extending the first entrance decomposition (8.7) from the countable space case, for a fixed n consider the event $\{\Phi_n \in B\}$ for arbitrary $B \in \mathcal{B}(\mathsf{X})$, and decompose this event over the mutually exclusive events $\{\Phi_n \in B, \tau_A = j\}$ for $j = 1, \dots, n$, where A is any other set in $\mathcal{B}(\mathsf{X})$. The general *first-entrance decomposition* can be written

$$P^n(x,B) = {}_AP^n(x,B) + \sum_{j=1}^{n-1} \int_A {}_AP^j(x,dw)P^{n-j}(w,B) \tag{8.20}$$

whilst the analogous *last-exit decomposition* is given by

$$P^n(x,B) = {}_AP^n(x,B) + \sum_{j=1}^{n-1} \int_A P^j(x,dw)\,{}_AP^{n-j}(w,B). \tag{8.21}$$

The first-entrance decomposition is clearly a decomposition of the event $\{\Phi_n \in A\}$ which could be developed using the Strong Markov Property and the stopping time $\zeta = \tau_A \wedge n$. The last exit decomposition, however, is not an example of the use of the Strong Markov Property: for, although the first entrance time τ_A is a stopping time for $\boldsymbol{\Phi}$, the last exit time is not a stopping time. These decompositions do however illustrate the same principle that underlies the Strong Markov Property, namely the decomposition of an event over the sub-events on which the random time takes on the (countable) set of values available to it.

We will develop classifications of sets using the generating functions for the series $\{P^n\}$ and $\{{}_AP^n\}$:

$$U^{(z)}(x,B) := \sum_{n=1}^{\infty} P^n(x,B)z^n, \qquad |z| < 1 \tag{8.22}$$

$$U_A^{(z)}(x,B) := \sum_{n=1}^{\infty} {}_AP^n(x,B)z^n, \qquad |z| < 1. \tag{8.23}$$

The kernel U then has the property

$$U(x,A) = \sum_{n=1}^{\infty} P^n(x,A) = \lim_{z \uparrow 1} U^{(z)}(x,A) \tag{8.24}$$

and as in the countable case, for any $x \in \mathsf{X}, A \in \mathcal{B}(\mathsf{X})$

$$\mathsf{E}_x(\eta_A) = U(x,A). \tag{8.25}$$

Thus uniform transience or recurrence is quantifiable in terms of the finiteness or otherwise of $U(x,A)$.

The return time probabilities $L(x,A) = \mathsf{P}_x\{\tau_A < \infty\}$ satisfy

$$L(x,A) = \sum_{n=1}^{\infty} {}_AP^n(x,A) = \lim_{z \uparrow 1} U_A^{(z)}(x,A). \tag{8.26}$$

We will prove the solidarity results we require by exploiting the convolution forms in (8.20) and (8.21). Multiplying by z^n in (8.20) and (8.21) and summing, the first entrance and last exit decompositions give, respectively, for $|z| < 1$

$$U^{(z)}(x,B) = U_A^{(z)}(x,B) + \int_A U_A^{(z)}(x,dw)U^{(z)}(w,B), \tag{8.27}$$

$$U^{(z)}(x,B) = U_A^{(z)}(x,B) + \int_A U^{(z)}(x,dw)U_A^{(z)}(w,B). \tag{8.28}$$

In classifying the chain $\boldsymbol{\Phi}$ we will use these relationships extensively.

8.2.2 The recurrence/transience dichotomy: chains with an atom

We can now move to classifying a chain Φ which admits an atom in a dichotomous way as either recurrent or transient. Through the splitting techniques of Chapter 5 this will then enable us to classify general chains.

Theorem 8.2.1 *Suppose that Φ is ψ-irreducible and admits an atom $\alpha \in \mathcal{B}^+(\mathsf{X})$. Then*

(i) *if α is recurrent, then every set in $\mathcal{B}^+(\mathsf{X})$ is recurrent.*

(ii) *if α is transient, then there is a countable covering of X by uniformly transient sets.*

PROOF **(i)** If $A \in \mathcal{B}^+(\mathsf{X})$ then for any x we have r, s such that $P^r(x, \alpha) > 0$, $P^s(\alpha, A) > 0$ and so

$$\sum_n P^{r+s+n}(x, A) \geq P^r(x, \alpha)\Big[\sum_n P^n(\alpha, \alpha)\Big] P^s(\alpha, A) = \infty. \tag{8.29}$$

Hence the series $U(x, A)$ diverges for every x, A when $U(\alpha, \alpha)$ diverges.

(ii) To prove the converse, we first note that for an atom, transience is equivalent to $L(\alpha, \alpha) < 1$, exactly as in Proposition 8.1.3.

Now consider the last exit decomposition (8.28) with $A, B = \alpha$. We have for any $x \in \mathsf{X}$

$$U^{(z)}(x, \alpha) = U_\alpha^{(z)}(x, \alpha) + U^{(z)}(x, \alpha) U_\alpha^{(z)}(\alpha, \alpha)$$

and so by rearranging terms we have for all $z < 1$

$$U^{(z)}(x, \alpha) = U_\alpha^{(z)}(x, \alpha)[1 - U_\alpha^{(z)}(\alpha, \alpha)]^{-1} \leq [1 - L(\alpha, \alpha)]^{-1} < \infty.$$

Hence $U(x, \alpha)$ is bounded for all x.

Now consider the countable covering of X given by the sets

$$\overline{\alpha}(j) = \{y : \sum_{n=1}^{j} P^n(y, \alpha) > j^{-1}\}.$$

Using the Chapman-Kolmogorov equations,

$$U(x, \alpha) \geq j^{-1} U(x, \overline{\alpha}(j)) \inf_{y \in \overline{\alpha}(j)} \sum_{n=1}^{j} P^n(y, \alpha) \geq j^{-2} U(x, \overline{\alpha}(j))$$

and thus $\{\overline{\alpha}(j)\}$ is the required cover by uniformly transient sets. □

We shall frequently find sets which are not uniformly transient themselves, but which can be covered by a countable number of uniformly transient sets. This leads to the definition

Transient sets

If $A \in \mathcal{B}(\mathsf{X})$ can be covered with a countable number of uniformly transient sets, then we call A *transient.*

8.2.3 The general recurrence/transience dichotomy

Now let us consider chains which do not have atoms, but which are strongly aperiodic.

We shall find that the split chain construction leads to a "solidarity result" for the sets in $\mathcal{B}^+(\mathsf{X})$ in the ψ-irreducible case, thus allowing classification of $\boldsymbol{\Phi}$ as a whole. Thus the following definitions will not be vacuous.

Stability Classification of ψ-irreducible Chains

(i) The chain $\boldsymbol{\Phi}$ is called *recurrent* if it is ψ-irreducible and $U(x, A) \equiv \infty$ for every $x \in \mathsf{X}$ and every $A \in \mathcal{B}^+(\mathsf{X})$.

(ii) The chain $\boldsymbol{\Phi}$ is called *transient* if it is ψ-irreducible and X is transient.

We first check that the split chain and the original chain have mutually consistent recurrent/transient classifications.

Proposition 8.2.2 *Suppose that $\boldsymbol{\Phi}$ is ψ-irreducible and strongly aperiodic. Then either both $\boldsymbol{\Phi}$ and $\check{\boldsymbol{\Phi}}$ are recurrent, or both $\boldsymbol{\Phi}$ and $\check{\boldsymbol{\Phi}}$ are transient.*

PROOF Strong aperiodicity ensures as in Proposition 5.4.5 that the Minorization Condition holds, and thus we can use the Nummelin Splitting of the chain $\boldsymbol{\Phi}$ to produce a chain $\check{\boldsymbol{\Phi}}$ on $\check{\mathsf{X}}$ which contains an accessible atom $\check{\boldsymbol{\alpha}}$.

We see from (5.9) that for every $x \in \mathsf{X}$, and for every $B \in \mathcal{B}^+(\mathsf{X})$,

$$\sum_{n=1}^{\infty} \int \delta_x^*(dy_i) \check{P}^n(y_i, B) = \sum_{n=1}^{\infty} P^n(x, B). \tag{8.30}$$

If $B \in \mathcal{B}^+(\mathsf{X})$ then since $\psi^*(B_0) > 0$ it follows from (8.30) that if $\check{\boldsymbol{\Phi}}$ is recurrent, so is $\boldsymbol{\Phi}$. Conversely, if $\check{\boldsymbol{\Phi}}$ is transient, by taking a cover of $\check{\mathsf{X}}$ with uniformly transient sets it is equally clear from (8.30) that $\boldsymbol{\Phi}$ is transient.

We know from Theorem 8.2.1 that $\check{\boldsymbol{\Phi}}$ is either transient or recurrent, and so the dichotomy extends in this way to $\boldsymbol{\Phi}$. □

To extend this result to general chains without atoms we first require a link between the recurrence of the chain and its resolvent.

Lemma 8.2.3 *For any $0 < \varepsilon < 1$ the following identity holds:*

$$\sum_{n=1}^{\infty} K_{a_\varepsilon}^n = \frac{1-\varepsilon}{\varepsilon} \sum_{n=0}^{\infty} P^n$$

PROOF From the generalized Chapman-Kolmogorov equations (5.46) we have

$$\sum_{n=1}^{\infty} K_{a_\varepsilon}^n = \sum_{n=1}^{\infty} K_{a_\varepsilon^{*n}} = \sum_{n=0}^{\infty} b(n) P^n$$

where we define $b(k)$ to be the kth term in the sequence $\sum_{n=1}^{\infty} a_\varepsilon^{*n}$. To complete the proof, we will show that $b(k) = (1-\varepsilon)/\varepsilon$ for all $k \geq 0$.

Let $B(z) = \sum b(k) z^k$, $A_\varepsilon(z) = \sum a_\varepsilon(k) z^k$ denote the power series representation of the sequences b and a_ε. From the identities

$$A_\varepsilon(z) = \Big(\frac{1-\varepsilon}{1-\varepsilon z}\Big) \qquad B(z) = \sum_{n=1}^{\infty} \Big(A_\varepsilon(z)\Big)^n$$

we see that $B(z) = ((1-\varepsilon)/\varepsilon)(1-z)^{-1}$. By uniqueness of the power series expansion it follows that $b(n) = (1-\varepsilon)/\varepsilon$ for all n, which completes the proof. □

As an immediate consequence of Lemma 8.2.3 we have

Proposition 8.2.4 *Suppose that $\boldsymbol{\Phi}$ is ψ-irreducible.*

(i) *The chain $\boldsymbol{\Phi}$ is transient if and only if each K_{a_ε}-chain is transient.*

(ii) *The chain $\boldsymbol{\Phi}$ is recurrent if and only if each K_{a_ε}-chain is recurrent.*

□

We may now prove

Theorem 8.2.5 *If $\boldsymbol{\Phi}$ is ψ-irreducible, then $\boldsymbol{\Phi}$ is either recurrent or transient.*

PROOF From Proposition 5.4.5 we are assured that the K_{a_ε}-chain is strongly aperiodic. Using Proposition 8.2.2 we know then that each K_{a_ε}-chain can be classified dichotomously as recurrent or transient.

Since Proposition 8.2.4 shows that the K_{a_ε}-chain passes on either of these properties to $\boldsymbol{\Phi}$ itself, the result is proved. □

We also have the following analogue of Proposition 8.2.4:

Theorem 8.2.6 *Suppose that $\boldsymbol{\Phi}$ is ψ-irreducible and aperiodic.*

(i) *The chain $\boldsymbol{\Phi}$ is transient if and only if one, and then every, m-skeleton $\boldsymbol{\Phi}^m$ is transient.*

(ii) *The chain $\boldsymbol{\Phi}$ is recurrent if and only if one, and then every, m-skeleton $\boldsymbol{\Phi}^m$ is recurrent.*

PROOF **(i)** If A is a uniformly transient set for the m-skeleton $\boldsymbol{\Phi}^m$, with $\sum_j P^{jm}(x,A) \le M$, then we have from the Chapman-Kolmogorov equations

$$\sum_{j=1}^{\infty} P^j(x,A) = \sum_{r=1}^{m} \int P^r(x,dy) \sum_j P^{jm}(y,A) \le mM. \tag{8.31}$$

Thus A is uniformly transient for $\boldsymbol{\Phi}$. Hence $\boldsymbol{\Phi}$ is transient whenever a skeleton is transient. Conversely, if $\boldsymbol{\Phi}$ is transient then every $\boldsymbol{\Phi}^k$ is transient, since

$$\sum_{j=1}^{\infty} P^j(x,A) \ge \sum_{j=1}^{\infty} P^{jk}(x,A).$$

(ii) If the m-skeleton is recurrent then from the equality in (8.31) we again have that

$$\sum P^j(x,A) = \infty, \qquad x \in \mathsf{X},\ A \in \mathcal{B}^+(\mathsf{X}) \tag{8.32}$$

so that the chain $\boldsymbol{\Phi}$ is recurrent.

Conversely, suppose that $\boldsymbol{\Phi}$ is recurrent. For any m it follows from aperiodicity and Proposition 5.4.5 that $\boldsymbol{\Phi}^m$ is ψ-irreducible, and hence by Theorem 8.2.5, this skeleton is either recurrent or transient. If it were transient we would have $\boldsymbol{\Phi}$ transient, from (i). □

It would clearly be desirable that we strengthen the definition of recurrence to a form of Harris recurrence in terms of $L(x,A)$, similar to that in Proposition 8.1.4. The key problem in moving to the general situation is that we do not have, for a general set, the equivalence in Proposition 8.1.3. There does not seem to be a simple way to exploit the fact that the atom in the split chain is not only recurrent but also satisfies $L(\check{\alpha},\check{\alpha}) = 1$, and the dichotomy in Theorem 8.2.5 is as far as we can go without considerably stronger techniques which we develop in the next chapter.

Until such time as we provide these techniques we will consider various partial relationships between transience and recurrence conditions, which will serve well in practical classification of chains.

8.3 Recurrence and transience relationships

8.3.1 Transience of sets

We next give conditions on hitting times which ensure that a set is uniformly transient, and which commence to link the behavior of τ_A with that of η_A.

Proposition 8.3.1 *Suppose that $\boldsymbol{\Phi}$ is a Markov chain, but not necessarily irreducible.*

(i) *If any set $A \in \mathcal{B}(\mathsf{X})$ is uniformly transient with $U(x,A) \le M$ for $x \in A$, then $U(x,A) \le 1 + M$ for every $x \in \mathsf{X}$.*

(ii) *If any set $A \in \mathcal{B}(\mathsf{X})$ satisfies $L(x,A) = 1$ for all $x \in A$, then A is recurrent. If $\boldsymbol{\Phi}$ is ψ-irreducible, then $A \in \mathcal{B}^+(\mathsf{X})$ and we have $U(x,A) \equiv \infty$ for $x \in \mathsf{X}$.*

(iii) *If any set $A \in \mathcal{B}(\mathsf{X})$ satisfies $L(x,A) \le \varepsilon < 1$ for $x \in A$, then we have $U(x,A) \le 1/[1-\varepsilon]$ for $x \in \mathsf{X}$, so that in particular A is uniformly transient.*

(iv) *Let $\tau_A(k)$ denote the k^{th} return time to A, and suppose that for some m*

$$\mathsf{P}_x(\tau_A(m) < \infty) \leq \varepsilon < 1, \qquad x \in A; \tag{8.33}$$

then $U(x,A) \leq 1 + m/[1-\varepsilon]$ for every $x \in \mathsf{X}$.

Proof **(i)** We use the first-entrance decomposition: letting $z \uparrow 1$ in (8.27) with $A = B$ shows that for all x,

$$U(x,A) \leq 1 + \sup_{y \in A} U(y,A), \tag{8.34}$$

which gives the required bound.

(ii) Suppose that $L(x,A) \equiv 1$ for $x \in A$. The last exit decomposition (8.28) gives

$$U^{(z)}(x,A) = U_A^{(z)}(x,A) + \int_A U^{(z)}(x,dy) U_A^{(z)}(y,A).$$

Letting $z \uparrow 1$ gives for $x \in A$,

$$U(x,A) = 1 + U(x,A),$$

which shows that $U(x,A) = \infty$ for $x \in A$, and hence that A is recurrent.

Suppose now that $\boldsymbol{\Phi}$ is ψ-irreducible. The set $A^\infty = \{x \in \mathsf{X} : L(x,A) = 1\}$ contains A by assumption. Hence we have for any x,

$$\int P(x,dy) L(y,A) = P(x,A) + \int_{A^c} P(x,dy) U_A(y,A) = L(x,A).$$

This shows that A^∞ is absorbing, and hence full by Proposition 4.2.3.

It follows from ψ-irreducibility that $K_{a_{\frac{1}{2}}}(x,A) > 0$ for all $x \in \mathsf{X}$, and we also have for all x that, from (5.47),

$$U(x,A) \geq \int_A K_{a_{\frac{1}{2}}}(x,dy) U(y,A) = \infty$$

as claimed.

(iii) Suppose on the other hand that $L(x,A) \leq \varepsilon < 1, x \in A$. The last exit decomposition again gives

$$U^{(z)}(x,A) = U_A^{(z)}(x,A) + \int_A U^{(z)}(x,dy) U_A^{(z)}(y,A) \leq 1 + \varepsilon U^{(z)}(x,A)$$

and so $U^{(z)}(x,A) \leq [1-\varepsilon]^{-1}$: letting $z \uparrow 1$ shows that A is uniformly transient as claimed.

(iv) Suppose now (8.33) holds. This means that for some fixed $m \in \mathbb{Z}_+$, we have $\varepsilon < 1$ with

$$\mathsf{P}_x(\eta_A \geq m) \leq \varepsilon, \qquad x \in A; \tag{8.35}$$

by induction in (8.35) we find that

$$\begin{aligned} \mathsf{P}_x(\eta_A \geq m(k+1)) &= \textstyle\int_A \mathsf{P}_x(\Phi_{\tau_A(km)} \in dy) \mathsf{P}_y(\eta_A \geq m) \\ &\leq \varepsilon\, \mathsf{P}_x(\tau_A(km) < \infty) \\ &\leq \varepsilon\, \mathsf{P}_x(\eta_A \geq km) \\ &\leq \varepsilon^{k+1}, \end{aligned} \tag{8.36}$$

and so for $x \in A$

$$\begin{aligned} U(x,A) &= \textstyle\sum_{n=1}^{\infty} \mathsf{P}_x(\eta_A \geq n) \\ &\leq m[1+\textstyle\sum_{k=1}^{\infty} \mathsf{P}_x(\eta_A \geq km)] \\ &\leq m/[1-\varepsilon]. \end{aligned} \tag{8.37}$$

We now use (i) to give the required bound over all of X. □

If there is one uniformly transient set then it is easy to identify other such sets, even without irreducibility. We have

Proposition 8.3.2 *If A is uniformly transient, and $B \overset{a}{\rightsquigarrow} A$ for some a, then B is uniformly transient. Hence if A is uniformly transient, there is a countable covering of $\overline{A}$ by uniformly transient sets.*

Proof From Lemma 5.5.2 (iii), we have when $B \overset{a}{\rightsquigarrow} A$ that for some $\delta > 0$,

$$U(x,A) \geq \int U(x,dy)K_a(y,A) \geq \delta U(x,B)$$

so that B is uniformly transient if A is uniformly transient. Since $\overline{A}$ is covered by the sets $\overline{A}(m)$, $m \in \mathbb{Z}_+$, and each $\overline{A}(m) \overset{a}{\rightsquigarrow} A$ for some a, the result follows. □

The next result provides a useful condition under which sets are transient even if not uniformly transient.

Proposition 8.3.3 *Suppose D^c is absorbing and $L(x, D^c) > 0$ for all $x \in D$. Then D is transient.*

Proof Suppose D^c is absorbing and write $B(m) = \{y \in D : P^m(y, D^c) \geq m^{-1}\}$: clearly, the sets $B(m)$ cover D since $L(x, D^c) > 0$ for all $x \in D$, by assumption.

But since D^c is absorbing, for every $y \in B(m)$ we have

$$\mathsf{P}_y(\eta_{B(m)} \geq m) \;\leq\; \mathsf{P}_y(\eta_D \geq m) \;\leq\; [1-m^{-1}]$$

and thus (8.33) holds for $B(m)$; from (8.37) it follows that $B(m)$ is uniformly transient. □

These results have direct application in the ψ-irreducible case. We next give a number of such consequences.

8.3.2 Identifying transient sets for ψ-irreducible chains

We first give an alternative proof that there is a recurrence/transience dichotomy for general state space chains which is an analogue of that in the countable state space case. Although this result has already been shown through the use of the splitting technique in Theorem 8.2.5, the following approach enables us to identify uniformly transient sets without going through the atom.

Theorem 8.3.4 *If $\boldsymbol{\Phi}$ is ψ-irreducible, then $\boldsymbol{\Phi}$ is either recurrent or transient.*

PROOF Suppose Φ is not recurrent: that is, there exists some pair $A \in \mathcal{B}^+(\mathsf{X})$, $x^* \in \mathsf{X}$ with $U(x^*, A) < \infty$. If $A_* = \{y : U(y, A) = \infty\}$, then $\psi(A_*) = 0$: for otherwise we would have $P^m(x^*, A_*) > 0$ for some m, and then

$$\begin{aligned} U(x^*, A) &\geq \int_{\mathsf{X}} P^m(x^*, dw) U(w, A) \\ &\geq \int_{A_*} P^m(x^*, dw) U(w, A) = \infty. \end{aligned} \tag{8.38}$$

Set $A_r = \{y \in A : U(y, A) \leq r\}$. Since $\psi(A) > 0$, and $A_r \uparrow A \cap A_*^c$, there must exist some r such that $\psi(A_r) > 0$, and by Proposition 8.3.1 (i) we have for all y,

$$U(y, A_r) \leq 1 + r. \tag{8.39}$$

Consider now $\overline{A}_r(M) = \{y : \sum_{m=0}^M P^m(y, A_r) > M^{-1}\}$. For any x, from (8.39)

$$\begin{aligned} M(1+r) \geq MU(x, A_r) &\geq \sum_{m=1}^{M} \sum_{n=m}^{\infty} P^n(x, A_r) \\ &= \sum_{n=0}^{\infty} \int_{\mathsf{X}} P^n(x, dw) \sum_{m=1}^{M} P^m(w, A_r) \\ &\geq \sum_{n=0}^{\infty} \int_{\overline{A}_r(M)} P^n(x, dw) \sum_{m=1}^{M} P^m(w, A_r) \\ &\geq M^{-1} \sum_{n=0}^{\infty} P^n(x, \overline{A}_r(M)). \end{aligned} \tag{8.40}$$

Since $\psi(A_r) > 0$ we have $\cup_m \overline{A}_r(m) = \mathsf{X}$, and so the $\{\overline{A}_r(m)\}$ form a partition of X into uniformly transient sets as required. □

The partition of X into uniformly transient sets given in Proposition 8.3.2 and in Theorem 8.3.4 leads immediately to

Theorem 8.3.5 *If Φ is ψ-irreducible and transient then every petite set is uniformly transient.*

PROOF If C is petite then by Proposition 5.5.5 (iii) there exists a sampling distribution a such that $C \overset{a}{\rightsquigarrow} B$ for any $B \in \mathcal{B}^+(\mathsf{X})$. If Φ is transient then there exists at least one $B \in \mathcal{B}^+(\mathsf{X})$ which is uniformly transient, so that C is uniformly transient from Proposition 8.3.2. □

Thus petite sets are also "small" within the transience definitions. This gives us a criterion for recurrence which we shall use in practice for many models; we combine it with a criterion for transience in

Theorem 8.3.6 *Suppose that Φ is ψ-irreducible. Then*

(i) *Φ is recurrent if there exists some petite set $C \in \mathcal{B}(\mathsf{X})$ such that $L(x, C) \equiv 1$ for all $x \in C$.*

(ii) *Φ is transient if and only if there exist two sets D, C in $\mathcal{B}^+(\mathsf{X})$ with $L(x, C) < 1$ for all $x \in D$.*

Proof (i) From Proposition 8.3.1 (ii) C is recurrent. Since C is petite Theorem 8.3.5 shows $\boldsymbol{\Phi}$ is recurrent. Note that we do not assume that C is in $\mathcal{B}^+(\mathsf{X})$, but that this follows also.

(ii) Suppose the sets C, D exist in $\mathcal{B}^+(\mathsf{X})$. There must exist $D_\varepsilon \subset D$ such that $\psi(D_\varepsilon) > 0$ and $L(x, C) \le 1 - \varepsilon$ for all $x \in D_\varepsilon$. If also $\psi(D_\varepsilon \cap C) > 0$ then since $L(x, C) \ge L(D_\varepsilon \cap C)$ we have that $D_\varepsilon \cap C$ is uniformly transient from Proposition 8.3.1 and the chain is transient.

Otherwise we must have $\psi(D_\varepsilon \cap C^c) > 0$. The maximal nature of ψ then implies that for some $\delta > 0$ and some $n \ge 1$ the set $C_\delta := \{y \in C : {}_C P^n(y, D_\varepsilon \cap C^c) > \delta\}$ also has positive ψ-measure. Since, for $x \in C_\delta$,

$$1 - L(x, C_\delta) \ge \int_{D_\varepsilon \cap C^c} {}_C P^n(x, dy)[1 - L(y, C_\delta)] \ge \delta\varepsilon$$

the set C_δ is uniformly transient, and again the chain is transient.

To prove the converse, suppose that $\boldsymbol{\Phi}$ is transient. Then for some petite set $C \in \mathcal{B}^+(\mathsf{X})$ the set $D = \{y \in C^c : L(y, C) < 1\}$ is non-empty; for otherwise by (i) the chain is recurrent. Suppose that $\psi(D) = 0$. Then by Proposition 4.2.3 there exists a full absorbing set $F \subset D^c$. By definition we have $L(x, C) = 1$ for $x \in F \setminus C$, and since F is absorbing it then follows that $L(x, C) = 1$ for every $x \in F$, and hence also that $L(x, C_0) = 1$ for $x \in F$ where $C_0 = C \cap F$ also lies in $\mathcal{B}^+(\mathsf{X})$.

But now from Proposition 8.3.1 (ii), we see that C_0 is recurrent, which is a contradiction of Theorem 8.3.5; and we conclude that $D \in \mathcal{B}^+(\mathsf{X})$ as required. □

We would hope that ψ-null sets would also have some transience property, and indeed they do.

Proposition 8.3.7 *If $\boldsymbol{\Phi}$ is ψ-irreducible then every ψ-null set is transient.*

Proof Suppose that $\boldsymbol{\Phi}$ is ψ-irreducible, and D is ψ-null. By Proposition 4.2.3, D^c contains an absorbing set, whose complement can be covered by uniformly transient sets as in Proposition 8.3.3: clearly, these uniformly transient sets cover D itself, and we are finished. □

As a direct application of Proposition 8.3.7 we extend the description of the cyclic decomposition for ψ-irreducible chains to give

Proposition 8.3.8 *Suppose that $\boldsymbol{\Phi}$ is a ψ-irreducible Markov chain on $(\mathsf{X}, \mathcal{B}(\mathsf{X}))$. Then there exist sets $D_1 \dots D_d \in \mathcal{B}(\mathsf{X})$ such that*

(i) *for $x \in D_i$, $P(x, D_{i+1}) = 1$, $i = 0, \dots, d-1 \pmod d$*

(ii) *the set $N = [\bigcup_{i=1}^d D_i]^c$ is ψ-null and transient.*

Proof The existence of the periodic sets D_i is guaranteed by Theorem 5.4.4, and the fact that the set N is transient is then a consequence of Proposition 8.3.3, since $\bigcup_{i=1}^d D_i$ is itself absorbing. □

In the main, transient sets and chains are ones we wish to exclude in practice. The results of this section have formalized the situation we would hope would hold: sets which appear to be irrelevant to the main dynamics of the chain are indeed so, in many different ways. But one cannot exclude them all, and for all of the statements where ψ-null (and hence transient) exceptional sets occur, one can construct examples to show that the "bad" sets need not be empty.

8.4 Classification using drift criteria

Identifying whether any particular model is recurrent or transient is not trivial from what we have done so far, and indeed, the calculation of the matrix U or the hitting time probabilities L involves in principle the calculation and analysis of all of the P^n, a daunting task in all but the most simple cases such as those addressed in Section 8.1.2.

Fortunately, it is possible to give practical criteria for both recurrence and transience, couched purely in terms of the drift of the one-step transition matrix P towards individual sets, based on Theorem 8.3.6.

8.4.1 A drift criterion for transience

We first give a criterion for transience of chains on general spaces, which rests on finding the minimal solution to a class of inequalities.

Recall that σ_C, the hitting time on a set C, is identical to τ_C on C^c and $\sigma_C = 0$ on C.

Proposition 8.4.1 *For any $C \in \mathcal{B}(\mathsf{X})$, the pointwise minimal non-negative solution to the set of inequalities*

$$\begin{aligned} \int P(x,dy)h(y) \le h(x), &\qquad x \in C^c \\ h(x) \ge 1, &\qquad x \in C, \end{aligned} \tag{8.41}$$

is given by the function

$$h^*(x) = \mathsf{P}_x(\sigma_C < \infty), \qquad x \in \mathsf{X};$$

and h^* satisfies (8.41) with equality.

Proof Since for $x \in C^c$

$$\mathsf{P}_x(\sigma_C < \infty) = P(x,C) + \int_{C^c} P(x,dy)\mathsf{P}_y(\sigma_C < \infty) = Ph^*(x)$$

it is clear that h^* satisfies (8.41) with equality.

Now let h be any solution to (8.41). By iterating (8.41) we have

$$\begin{aligned} h(x) &\ge \int_C P(x,dy)h(y) + \int_{C^c} P(x,dy)h(y) \\ &\ge \int_C P(x,dy)h(y) + \int_{C^c} P(x,dy)[\int_C P(y,dz)h(z) + \int_{C^c} P(x,dz)h(z)] \\ &\vdots \\ &\ge \sum_{j=1}^N \int_C {}_CP^j(x,dy)h(y) + \int_{C^c} {}_CP^N(x,dy)h(y). \end{aligned} \tag{8.42}$$

Letting $N \to \infty$ shows that $h(x) \geq h^*(x)$ for all x. □

This gives the required drift criterion for transience. Recall the definition of the drift operator as $\Delta V(x) = \int P(x, dy)V(y) - V(x)$; obviously Δ is well-defined if V is bounded. We define the sublevel set $C_V(r)$ of any function V for $r \geq 0$ by

$$C_V(r) := \{x : V(x) \leq r\}.$$

Theorem 8.4.2 *Suppose $\boldsymbol{\Phi}$ is a ψ-irreducible chain. Then $\boldsymbol{\Phi}$ is transient if and only if there exists a bounded function $V : \mathsf{X} \to \mathbb{R}_+$ and $r \geq 0$ such that*

(i) *both $C_V(r)$ and $C_V(r)^c$ lie in $\mathcal{B}^+(\mathsf{X})$;*

(ii) *whenever $x \in C_V(r)^c$,*

$$\Delta V(x) > 0. \tag{8.43}$$

Proof Suppose that V is an arbitrary bounded solution of (i) and (ii), and let M be a bound for V over X. Clearly $M > r$. Set $C = C_V(r)$, $D = C^c$, and

$$h_V(x) = \begin{cases} [M - V(x)]/[M - r] & x \in D \\ 1 & x \in C \end{cases}$$

so that h_V is a solution of (8.41). Then from the minimality of h^* in Proposition 8.4.1, h_V is an upper bound on h^*, and since for $x \in D, h_V(x) < 1$ we must have $L(x, C) < 1$ also for $x \in D$.

Hence $\boldsymbol{\Phi}$ is transient as claimed, from Theorem 8.3.6.

Conversely, if $\boldsymbol{\Phi}$ is transient, there exists a bounded function V satisfying (i) and (ii). For from Theorem 8.3.6 we can always find $\varepsilon < 1$ and a petite set $C \in \mathcal{B}^+(\mathsf{X})$ such that $\{y \in C^c : L(y, C) < \varepsilon\}$ is also in $\mathcal{B}^+(\mathsf{X})$. Thus from Proposition 8.4.1, the function $V(x) = 1 - \mathsf{P}_x(\sigma_C < \infty)$ has the required properties. □

8.4.2 A drift criterion for recurrence

Theorem 8.4.2 essentially asserts that if $\boldsymbol{\Phi}$ "drifts away" in expectation from a set in $\mathcal{B}^+(\mathsf{X})$, as indicated in (8.43), then $\boldsymbol{\Phi}$ is transient. Of even more value in assessing stability are conditions which show that "drift toward" a set implies recurrence, and we provide the first of these now. The condition we will use is

Drift criterion for recurrence

(V1) There exists a positive function V and a set $C \in \mathcal{B}(\mathsf{X})$ satisfying

$$\Delta V(x) \leq 0, \qquad x \in C^c \tag{8.44}$$

We will find frequently that, in order to test such drift for the process $\boldsymbol{\Phi}$, we need to consider functions $V : \mathsf{X} \to \mathbb{R}$ such that the set $C_V(M) = \{y \in \mathsf{X} : V(y) \leq M\}$ is "finite" for each M. Such a function on a countable space or topological space is easy to define: in this abstract setting we first need to define a class of functions with this property, and we will find that they recur frequently, giving further meaning to the intuitive meaning of petite sets.

Functions unbounded off petite sets

We will call a measurable function $V : \mathsf{X} \to \mathbb{R}_+$ *unbounded off petite sets* for $\boldsymbol{\Phi}$ if for any $n < \infty$, the sublevel set

$$C_V(n) = \{y : V(y) \leq n\}$$

is petite.

Note that since, for an irreducible chain, a finite union of petite sets is petite, and since any subset of a petite set is itself petite, a function $V : \mathsf{X} \to \mathbb{R}_+$ will be unbounded off petite sets for $\boldsymbol{\Phi}$ if there merely exists a sequence $\{C_j\}$ of petite sets such that, for any $n < \infty$

$$C_V(n) \subseteq \bigcup_{j=1}^{N} C_j \tag{8.45}$$

for some $N < \infty$. In practice this may be easier to verify directly.

We now have a drift condition which provides a test for recurrence.

Theorem 8.4.3 *Suppose $\boldsymbol{\Phi}$ is ψ-irreducible. If there exists a petite set $C \subset \mathsf{X}$, and a function V which is unbounded off petite sets such that (V1) holds then $L(x, C) \equiv 1$ and $\boldsymbol{\Phi}$ is recurrent.*

Proof We will show that $L(x, C) \equiv 1$ which will give recurrence from Theorem 8.3.6. Note that by replacing the set C by $C \cup C_V(n)$ for n suitably large, we can assume without loss of generality that $C \in \mathcal{B}^+(\mathsf{X})$.

Suppose by way of contradiction that the chain is transient, and thus that there exists some $x^* \in C^c$ with $L(x^*, C) < 1$.

Set $C_V(n) = \{y \in \mathsf{X} : V(y) \leq n\}$: we know this is petite, by definition of V, and hence it follows from Theorem 8.3.5 that $C_V(n)$ is uniformly transient for any n. Now fix M large enough that

$$M > V(x^*)/[1 - L(x^*, C)]. \tag{8.46}$$

Let us modify P to define a kernel $\widehat{P}$ with entries $\widehat{P}(x,A) = P(x,A)$ for $x \in C^c$ and $\widehat{P}(x,x) = 1, x \in C$. This defines a chain $\widehat{\boldsymbol{\Phi}}$ with C as an absorbing set, and with the property that for all $x \in \mathsf{X}$

$$\int \widehat{P}(x,dy)V(y) \leq V(x). \tag{8.47}$$

Since P is unmodified outside C, but $\widehat{\boldsymbol{\Phi}}$ is absorbed in C, we also have

$$\widehat{P}^n(x,C) = \mathsf{P}_x(\tau_C \leq n) \uparrow L(x,C), \qquad x \in C^c, \tag{8.48}$$

whilst for $A \subseteq C^c$

$$\widehat{P}^n(x,A) \leq P^n(x,A), \qquad x \in C^c. \tag{8.49}$$

By iterating (8.47) we thus get, for fixed $x \in C^c$

$$\begin{aligned} V(x) &\geq \textstyle\int \widehat{P}^n(x,dy)V(y) \\ &\geq \int_{C^c \cap [C_V(M)]^c} \widehat{P}^n(x,dy)V(y) \\ &\geq M\Big[1 - \widehat{P}^n(x, C_V(M) \cup C)\Big]. \end{aligned} \tag{8.50}$$

Since $C_V(M)$ is uniformly transient, from (8.49) we have

$$\widehat{P}^n(x^*, C_V(M) \cap C^c) \leq P^n(x^*, C_V(M) \cap C^c) \to 0, \qquad n \to \infty. \tag{8.51}$$

Combining this with (8.48) gives

$$[1 - \widehat{P}^n(x^*, C_V(M) \cup C)] \to [1 - L(x^*, C)], \qquad n \to \infty. \tag{8.52}$$

Letting $n \to \infty$ in (8.50) for $x = x^*$ provides a contradiction with (8.52) and our choice of M. Hence we must have $L(x,C) \equiv 1$, and $\boldsymbol{\Phi}$ is recurrent, as required. □

8.4.3 Random walks with bounded range

The drift condition on the function V in Theorem 8.4.3 basically says that, whenever the chain is outside C, it "moves down" towards that part of the space described by the petite sets outside which V tends to infinity.

This condition implies that we know where the petite sets for $\boldsymbol{\Phi}$ lie, and can identify those functions which are unbounded off the petite sets. This provides very substantial motivation for the identification of petite sets in a manner independent of $\boldsymbol{\Phi}$; and for many chains we can use the results in Chapter 6 to give such form to the results.

On a countable space, of course, finite sets are petite. Our problem is then to identify the correct test function to use in the criteria.

In order to illustrate the use of the drift criteria we will first consider the simplest case of a random walk on $\mathbb{Z}$ with finite range r. Thus we assume the increment distribution Γ is concentrated on the integers and is such that $\Gamma(x) = 0$ for $|x| > r$. We then have a relatively simple proof of the result in Theorem 8.1.5.

Proposition 8.4.4 *Suppose that $\boldsymbol{\Phi}$ is an irreducible random walk on the integers. If the increment distribution Γ has a bounded range and the mean of Γ is zero, then $\boldsymbol{\Phi}$ is recurrent.*

Proof In Theorem 8.4.3 choose the test function $V(x) = |x|$. Then for $x > r$ we have that

$$\sum_y P(x,y)[V(y) - V(x)] = \sum_y \Gamma(w)w,$$

whilst for $x < -r$ we have that

$$\sum_y P(x,y)[V(y) - V(x)] = -\sum_w \Gamma(w)w.$$

Suppose the "mean drift"

$$\beta = \sum_w \Gamma(w)w = 0.$$

Then the conditions of Theorem 8.4.3 are satisfied with $C = \{-r, \dots, r\}$ and with (8.44) holding for $x \in C^c$, and so the chain is recurrent. □

Proposition 8.4.5 *Suppose that Φ is an irreducible random walk on the integers. If the increment distribution Γ has a bounded range and the mean of Γ is non-zero, then Φ is transient.*

Proof Suppose Γ has non-zero mean $\beta > 0$. We will establish for some bounded monotone increasing V that

$$\sum_y P(x,y)V(y) = V(x) \tag{8.53}$$

for $x \geq r$.

This time choose the test function $V(x) = 1 - \rho^x$ for $x \geq 0$, and $V(x) = 0$ elsewhere. The sublevel sets of V are of the form $(-\infty, r]$ with $r \geq 0$. This function satisfies (8.53) if and only if for $x \geq r$

$$\sum_y P(x,y)[\rho^y/\rho^x] = 1 \tag{8.54}$$

so that this V can be constructed as a valid test function if (and only if) there is a $\rho < 1$ with

$$\sum_w \Gamma(w)\rho^w = 1. \tag{8.55}$$

Therefore the existence of a solution to (8.55) will imply that the chain is transient, since return to the whole half line $(-\infty, r]$ is less than sure from Proposition 8.4.2. Write $\beta(s) = \sum_w \Gamma(w)s^w$: then β is well defined for $s \in (0,1]$ by the bounded range assumption. By irreducibility, we must have $\Gamma(w) > 0$ for some $w < 0$, so that $\beta(s) \to \infty$ as $s \to 0$. Since $\beta(1) = 1$, and $\beta'(1) = \sum_w w\Gamma(w) = \beta > 0$ it follows that such a ρ exists, and hence the chain is transient.

Similarly, if the mean of Γ is negative, we can by symmetry prove transience because the chain fails to return to the half line $[-r, \infty)$. □

For random walk on the half line $\mathbb{Z}_+$ with bounded range, as defined by (RWHL1) we find

Proposition 8.4.6 *If the random walk increment distribution Γ on the integers has mean β and a bounded range, then the random walk on $\mathbb{Z}_+$ is recurrent if and only if $\beta \leq 0$.*

PROOF If β is positive, then the probability of return of the unrestricted random walk to $(-\infty, r]$ is less than one, for starting points above r, and since the probability of return of the random walk on a half line to $[0, r]$ is identical to the return to $(-\infty, r]$ for the unrestricted random walk, the chain is transient.

If $\beta \leq 0$, then we have as for the unrestricted random walk that, for the test function $V(x) = x$ and all $x \geq r$

$$\sum_y P(x,y)[V(y) - V(x)] = \sum_w \Gamma(w) w \leq 0;$$

but since, in this case, the set $\{x \leq r\}$ is finite, we have (8.44) holding and the chain is recurrent. □

The first part of this proof involves a so-called "stochastic comparison" argument: we use the return time probabilities for one chain to bound the same probabilities for another chain. This is simple but extremely effective, and we shall use it a number of times in classifying random walk. A more general formulation will be given in Section 9.5.1.

Varying the condition that the range of the increment is bounded requires a much more delicate argument, and indeed the known result of Theorem 8.1.5 for a general random walk on $\mathbb{Z}$, that recurrence is equivalent to the mean $\beta = 0$, appears difficult if not impossible to prove by drift methods without some bounds on the spread of Γ.

8.5 Classifying random walk on $\mathbb{R}_+$

In order to give further exposure to the use of drift conditions, we will conclude this chapter with a detailed examination of random walk on $\mathbb{R}_+$.

The analysis here is obviously immediately applicable to the various queueing and storage models introduced in Chapter 2 and Chapter 3, although we do not fill in the details explicitly. The interested reader will find, for example, that the conditions on the increment do translate easily into intuitively appealing statements on the mean input rate to such systems being no larger than the mean service or output rate if recurrence is to hold.

These results are intended to illustrate a variety of approaches to the use of the stability criteria above. Different test functions are utilized, and a number of different methods of ensuring they are applicable are developed. Many of these are used in the sequel where we classify more general models.

As in (RW1) and (RWHL1) we let $\boldsymbol{\Phi}$ denote a chain with

$$\Phi_n = [\Phi_{n-1} + W_n]^+$$

where as usual W_n is a noise variable with distribution Γ and mean β which we shall assume in this section is well-defined and finite.

Clearly we would expect from the bounded increments results above that $\beta \leq 0$ is the appropriate necessary and sufficient condition for recurrence of $\boldsymbol{\Phi}$. We now address the three separate cases in different ways.

8.5.1 Recurrence when β is negative

When the inequality is strict it is not hard to show that the chain is recurrent.

Proposition 8.5.1 *If Φ is random walk on a half line and if*

$$\beta = \int w\,\Gamma(dw) < 0$$

then Φ is recurrent.

PROOF Clearly the chain is φ-irreducible when $\beta < 0$ with $\varphi = \delta_0$, and all compact sets are small as in Chapter 5. To prove recurrence we use Theorem 8.4.3, and show that we can in fact find a suitably unbounded function V and a compact set C satisfying

$$\int P(x, dy)V(y) \le V(x) - \varepsilon, \qquad x \in C^c, \tag{8.56}$$

for some $\varepsilon > 0$. As in the countable case we note that since $\beta < 0$ there exists $x_0 < \infty$ such that

$$\int_{-x_0}^{\infty} w\,\Gamma(dw) < \beta/2 < 0,$$

and thus if $V(x) = x$, for $x > x_0$

$$\int P(x, dy)[V(y) - V(x)] \le \int_{-x_0}^{\infty} w\,\Gamma(dw). \tag{8.57}$$

Hence taking $\varepsilon = \beta/2$ and $C = [0, x_0]$ we have the required result. □

8.5.2 Recurrence when β is zero

When the mean increment $\beta = 0$ the situation is much less simple, and in general the drift conditions can be verified simply only under somewhat stronger conditions on the increment distribution Γ, such as an assumption of a finite variance of the increments.

We will find it convenient to develop prior to our calculations some detailed bounds on the moments of Γ, which will become relevant when we consider test functions of the form $V(x) = \log(1 + |x|)$.

Lemma 8.5.2 *Let W be a random variable with law Γ, s a positive number and t any real number. Then for any $A \subseteq \{w \in \mathbb{R} : s + tw > 0\}$,*

$$\begin{aligned} \mathsf{E}[\log(s + tW)\mathbb{1}\{W \in A\}] &\le \Gamma(A)\log(s) + (t/s)\mathsf{E}[W\mathbb{1}\{W \in A\}] \\ &\quad -(t^2/(2s^2))\mathsf{E}[W^2\mathbb{1}\{W \in A, tW < 0\}] \end{aligned}$$

PROOF For all $x > -1$, $\log(1 + x) \le x - (x^2/2)\mathbb{1}\{x < 0\}$. Thus

$$\begin{aligned} \log(s + tW)\mathbb{1}\{W \in A\} &= [\log(s) + \log(1 + tW/s)]\mathbb{1}\{W \in A\} \\ &\le [\log(s) + tW/s]\mathbb{1}\{W \in A\} \\ &\quad -((tW)^2/(2s^2))\mathbb{1}\{tW < 0, W \in A\}] \end{aligned}$$

and taking expectations gives the result. □

Lemma 8.5.3 *Let W be a random variable with law Γ and finite variance. Let s be a positive number and t a real number. Then*

$$\lim_{x\to\infty} -x\mathsf{E}[W\mathbb{1}\{W < t - sx\}] = \lim_{x\to\infty} x\mathsf{E}[W\mathbb{1}\{W > t + sx\}] = 0. \tag{8.58}$$

Furthermore, if $\mathsf{E}[W] = 0$, then

$$\lim_{x\to\infty} -x\mathsf{E}[W\mathbb{1}\{W > t - sx\}] = \lim_{x\to\infty} x\mathsf{E}[W\mathbb{1}\{W < t + sx\}] = 0. \tag{8.59}$$

Proof This is a consequence of

$$0 \le \lim_{x\to\infty} (t+sx) \int_{t+sx}^{\infty} w\Gamma(dw) \le \lim_{x\to\infty} \int_{t+sx}^{\infty} w^2\Gamma(dw) = 0,$$

and

$$0 \le \lim_{x\to-\infty} (t+sx) \int_{-\infty}^{t+sx} w\Gamma(dw) \le \lim_{x\to-\infty} \int_{-\infty}^{t+sx} w^2\Gamma(dw) = 0.$$

If $\mathsf{E}[W] = 0$, then $\mathsf{E}[W\mathbb{1}\{W > t+sx\}] = -\mathsf{E}[W\mathbb{1}\{W < t+sx\}]$, giving the second result. □

We now prove

Proposition 8.5.4 *If W is an increment variable on $\mathbb{R}$ with $\beta = 0$ and*

$$0 < \mathsf{E}[W^2] = \int w^2\,\Gamma(dw) < \infty$$

then the random walk on $\mathbb{R}_+$ with increment W is recurrent.

Proof We use the test function

$$V(x) = \begin{cases} \log(1+x) & x > R \\ 0 & 0 \le x \le R \end{cases} \tag{8.60}$$

where R is a positive constant to be chosen. Since $\beta = 0$ and $0 < \mathsf{E}[W^2]$ the chain is δ_0-irreducible, and we have seen that all compact sets are small as in Chapter 5. Hence V is unbounded off petite sets.

For $x > R, 1+x > 0$, and thus by Lemma 8.5.2,

$$\begin{aligned} \mathsf{E}_x[V(X_1)] &= \mathsf{E}[\log(1+x+W)\mathbb{1}\{x+W > R\}] \\ &\le (1-\Gamma(-\infty, R-x))\log(1+x) + U_1(x) - U_2(x), \end{aligned} \tag{8.61}$$

where in order to bound the terms in the expansion of the logarithms in V, we consider separately

$$\begin{aligned} U_1(x) &= (1/(1+x))\mathsf{E}[W\mathbb{1}\{W > R-x\}] \\ U_2(x) &= (1/(2(1+x)^2))\mathsf{E}[W^2\mathbb{1}\{R-x < W < 0\}] \end{aligned} \tag{8.62}$$

Since $\mathsf{E}[W^2] < \infty$

$$U_2(x) = (1/(2(1+x)^2))\mathsf{E}[W^2\mathbb{1}\{W < 0\}] - o(x^{-2}),$$

and by Lemma 8.5.3, U_1 is also $o(x^{-2})$.

Thus by choosing R large enough

$$\begin{aligned} \mathsf{E}_x[V(X_1)] &\le V(x) - (1/(2(1+x)^2))\mathsf{E}[W^2\mathbb{1}\{W < 0\}] + o(x^{-2}) \\ &\le V(x), \qquad x > R. \end{aligned} \tag{8.63}$$

Hence the conditions of Theorem 8.4.3 hold, and chain is recurrent. □

8.5.3 Transience of skip-free random walk when β is positive

It is possible to verify transience when $\beta > 0$, without any restrictions on the range of the increments of the distribution Γ, thus extending Proposition 8.4.5; but the argument (in Proposition 9.1.2) is a somewhat different one which is based on the Strong Law of Large Numbers and must wait some stronger results on the meaning of recurrence in the next chapter.

Proving transience for random walk without bounded range using drift conditions is difficult in general. There is however one model for which some exact calculations can be made: this is the random walk which is "skip-free to the right" and which models the GI/M/1 queue as in Theorem 3.3.1.

Proposition 8.5.5 *If Φ denotes random walk on a half line $\mathbb{Z}_+$ which is skip-free to the right (so $\Gamma(x) = 0$ for $x > 1$), and if*

$$\beta = \sum w\, \Gamma(w) > 0$$

then Φ is transient.

PROOF We can assume without loss of generality that $\Gamma(-\infty, 0) > 0$: for clearly, if $\Gamma[0, \infty) = 1$ then $\mathsf{P}_x(\tau_0 < \infty) = 0, x > 0$ and the chain moves inexorably to infinity; hence it is not irreducible, and it is transient in every meaning of the word.

We will show that for a chain which is skip-free to the right the condition $\beta > 0$ is sufficient for transience, by examining the solutions of the equations

$$\sum P(x,y)V(y) = V(x), \qquad x \geq 1 \tag{8.64}$$

and actually constructing a bounded non-constant positive solution if β is positive. The result will then follow from Theorem 8.4.2.

First note that we can assume $V(0) = 0$ by linearity, and write out the equation (8.64) in this case as

$$V(x) = \Gamma(-x+1)V(1) + \Gamma(-x+2)V(2) + \ldots + \Gamma(1)V(1+x). \tag{8.65}$$

Once the first value in the $V(x)$ sequence is chosen, we therefore have the remaining values given by an iterative process. Our goal is to show that we can define the sequence in a way that gives us a non-constant positive bounded solution to (8.65).

In order to do this we first write

$$V^*(z) = \sum_0^\infty V(x)z^x, \qquad \Gamma^*(z) = \sum_{-\infty}^\infty \Gamma(x)z^x,$$

where $V^*(z)$ has yet to be shown to be defined for any z and $\Gamma^*(z)$ is clearly defined at least for $|z| \geq 1$. Multiplying by z^x in (8.65) and summing we have that

$$V^*(z) = \Gamma^*(z^{-1})V^*(z) - \Gamma(1)V(1) \tag{8.66}$$

Now suppose that we can show (as we do below) that there is an analytic expansion of the function

$$z^{-1}[1-z]/[\Gamma^*(z^{-1}) - 1] = \sum_0^\infty b_n z^n \tag{8.67}$$

in the region $0 < z < 1$ with $b_n \geq 0$. Then we will have the identity

$$\begin{aligned} V^*(z) &= z\Gamma(1)V(1)z^{-1}/[\Gamma^*(z^{-1})-1] \\ &= z\Gamma(1)V(1)(\textstyle\sum_0^\infty z^n)z^{-1}[1-z]/[\Gamma^*(z^{-1})-1] \qquad (8.68) \\ &= z\Gamma(1)V(1)(\textstyle\sum_0^\infty z^n)(\sum_0^\infty b_m z^m). \end{aligned}$$

From this, we will be able to identify the form of the solution V. Explicitly, from (8.68) we have

$$V^*(z) = z\Gamma(1)V(1)\textstyle\sum_{n=0}^\infty z^n \sum_{m=0}^n b_m \qquad (8.69)$$

so that equating coefficients of z^n in (8.69) gives

$$V(x) = \Gamma(1)V(1)\sum_{m=0}^{x-1} b_m.$$

Clearly then the solution V is bounded and non-constant if

$$\sum_m b_m < \infty. \qquad (8.70)$$

Thus we have reduced the question of transience to identifying conditions under which the expansion in (8.67) holds with the coefficients b_j positive and summable.

Let us write $a_j = \Gamma(1-j)$ so that

$$A(z) := \sum_0^\infty a_j z^j = z\Gamma^*(z^{-1})$$

and for $0 < z < 1$ we have

$$\begin{aligned} B(z) := z[\Gamma^*(z^{-1})-1]/[1-z] &= [A(z)-z]/[1-z] \\ &= 1-[1-A(z)]/[1-z] \qquad (8.71) \\ &= 1-\textstyle\sum_0^\infty z^j \sum_{n=j+1}^\infty a_n. \end{aligned}$$

Now if we have a positive mean for the increment distribution,

$$|\sum_0^\infty z^j \sum_{n=j+1}^\infty a_n| \leq \sum_n n a_n < 1$$

and so $B(z)^{-1}$ is well defined for $|z| < 1$; moreover, by the expansion in (8.71)

$$B(z)^{-1} = \sum b_j z^j$$

with all with all $b_j \geq 0$, and hence by Abel's Theorem,

$$\sum b_j = [1 - \sum_n n a_n]^{-1} = \beta^{-1}$$

which is finite as required. □

8.6 Commentary

On countable spaces the solidarity results we generalize here are classical, and thorough expositions are in Feller [76], Chung [49], Çinlar [40] and many more places. Recurrence is called persistence by Feller, but the terminology we use here seems to have become the more standard. The first entrance, and particularly the last exit, decomposition are vital tools introduced and exploited in a number of ways by Chung [49].

There are several approaches to the transience/recurrence dichotomy. A common one which can be shown to be virtually identical with that we present here uses the concept of inessential sets (sets for which η_A is almost surely finite). These play the role of transient parts of the space, with recurrent parts of the space being sets which are not inessential. This is the approach in Orey [208], based on the original methods of Doeblin [67] and Doob [68].

Our presentation of transience, stressing the role of uniformly transient sets, is new, although it is implicit in many places. Most of the individual calculations are in Nummelin [202], and a number are based on the more general approach in Tweedie [272]. Equivalences between properties of the kernel $U(x, A)$, which we have called recurrence and transience properties, and the properties of essential and inessential sets are studied in Tuominen [268].

The uniform transience property is inherently stronger than the inessential property, and it certainly aids in showing that the skeletons and the original chain share the dichotomy between recurrence and transience. For use of the properties of skeleton chains in direct application, see Tjøstheim [265].

The drift conditions we give here are due in the countable case to Foster [82], and the versions for more general spaces were introduced in Tweedie [275, 276] and in Kalashnikov [117]. We shall revisit these drift conditions, and expand somewhat on their implications in the next chapter. Stronger versions of (V1) will play a central role in classifying chains as yet more stable in due course.

The test functions for classifying random walk in the bounded range case are directly based on those introduced by Foster [82]. The evaluation of the transience condition for skip-free walks, given in Proposition 8.5.5, is also due to Foster. The approximations in the case of zero drift are taken from Guo and Petrucelli [92] and are reused in analyzing SETAR models in Section 9.5.2.

The proof of recurrence of random walk in Theorem 8.1.5, using the weak law of large numbers, is due to Chung and Ornstein [51]. It appears difficult to prove this using the elementary drift methods.

The drift condition in the case of negative mean gives, as is well known, a stronger form of recurrence: the concerned reader will find that this is taken up in detail in Chapter 11, where it is a central part of our analysis.