Mixing Times for Uniformly Ergodic Markov Chains

David Aldous[1] and László Lovász[2] and Peter Winkler[3]

Abstract

Consider the class of discrete time, general state space Markov chains which satisfy a "uniform ergodicity under sampling" condition. There are many ways to quantify the notion of "mixing time", that is time to approach stationarity from a worst initial state. We prove results asserting equivalence (up to universal constants) of different quantifications of mixing time. This work combines three areas of Markov theory which are rarely connected: the potential-theoretical characterization of optimal stopping times, the theory of stability and convergence to stationarity for general-state chains, and the theory surrounding mixing times for finite-state chains.

Keywords. Markov chain, minorization, mixing time, randomized algorithm, stopping time.

1 Introduction

Our topic lies near the intersection of three different areas of the theory of (discrete time, general state space) Markov chains.

(a) Potential theory, as treated in e.g. Revuz [29] or Dellacherie and Meyer [6]. This theory classically focused on transient chains, but does include results on *recurrent potential* and its relation to hitting times for recurrent chains, which are our concern (see also Syski [32]).

(b) The theory of convergence to stationarity for general state space chains, treated in Orey [26] and Nummelin [25] and in particular given a recent very clear exposition by Meyn and Tweedie [23].

(c) The theory surrounding *mixing times*, i.e. quantitative measures of times to approach stationarity, for finite-state chains. This is treated (in the reversible setting) in the forthcoming book Aldous and Fill [2]. See also Diaconis [7] for the case of random walks on groups, and Sinclair [31] and Motwani and Raghavan [24] for uses in the theory of algorithms.

These areas have developed rather independently, and the connections are not easy to find in the monographs above. The purpose of this paper is to record two related results which explicitly use aspects of all three areas. These results (Theorems 1 and 3) assert the equivalence (up to universal constants) of different formalizations of "mixing time" in (essentially) the context of uniformly ergodic general state-space chains.

In section 1.1 we recall the underlying algorithmic motivation for studying mixing times. In section 1.2 we describe, as mathematical background, known results from each of the three

[1]Statistics Dept., U.C. Berkeley CA 94720: aldous@stat.berkeley.edu
[2]Dept. of Computer Science, Yale University, New Haven CT 06510: lovasz@cs.yale.edu
[3]Bell Laboratories 2C-379, Murray Hill NJ 07960: pw@bell-labs.com

areas (a,b,c) above. Section 1.3 states our new results, and section 1.4 interprets the conceptual significance of the new results.

1.1 Mixing times and randomized algorithms

One motivation for the study of mixing times comes from computer science, more exactly from the analysis of sampling algorithms, which has been an active area over the last ten years. In randomized algorithms solving a variety of computational tasks (approximate enumeration, volume computation, integration, simulated annealing, generation of contingency tables etc.) the key element is to sample from a given distribution π over a known but large and complicated set. The basic method is to construct an ergodic Markov chain with stationary distribution π, and then run the chain for an appropriately large number of steps. The details vary according to the goal of the algorithm, which might be to estimate an average $\int f d\pi$, or to bound the π-probability of some set of unlikely states, or to generate typical realizations from π for illustrative purposes. The number of steps required by a particular algorithm (as a function of the Markov chain) will depend on some algorithm-specific notion of "mixing time", i.e. the number of steps until the distribution approaches stationarity. Three such notions are mentioned below. Even for the more restricted issue of quantifying the distance between the time-t distribution and π there are several answers: total variation distance (i.e. l_1 distance for densities), the analogous l_2 or l_∞ distances, Kullback-Leibler distance, etc.

In a sampling algorithm, we may want to generate a single state from the stationary distribution, starting from some fixed state (determined by the rest of the algorithm). The minimum mean time to do so is a definition of a mixing time from a given state. If we don't have more information about the starting state, we have to use the maximum over all starting states, which we'll call the *mixing time* $\mathcal{T}_{\text{mix}}$ (precise definitions will be given later). But it may be the case that we need to generate several independent samples from the stationary distribution. In this case we might start the second run of the Markov chain where the first one stopped, and so the expected time needed for this will be the average, rather than the maximum, of mixing times from individual states. This leads us to the definition of the *reset time* $\mathcal{T}_{\text{reset}}$. Alternatively, we may use the Markov chain to find an element from a specified, but not directly accessible subset of the state space. The worst expected time needed for this (normalized by the measure of the subset) is the *set hitting time* $\mathcal{T}_{\text{set}}$.

This paper is motivated by a foundational question.

> Does it make sense to undertake a mathematical analysis of a given chain being used in some sampling algorithm, without paying attention to the algorithmic use of the samples?

If the mixing times for the given chain associated with different algorithms were incomparable, then it wouldn't make sense. Fortunately this is not the case. For reversible chains, it has been known for a long time [1] that various mixing times (including the three mixing times above) are within absolute constant factors of each other, assuming that they are finite at all. The case of non-reversible chains is a bit more complicated, but results of this paper show that many mixing measures fall into three groups only, where measures in the same group are within absolute constant factors of each other, one group is always "above" the other two, and these two are related in an interesting way through time-reversal.

To discuss a celebrated example, consider a convex body K in R^n and suppose that we want to generate a uniformly distributed point in it. We assume that the body is in isotropic position (i.e. a uniform random point $(X_1, \ldots, X_n)$ of K has $EX_i = 0$ and $EX_iX_j = 1_{(i=j)}$). Choose an appropriately small $\delta > 0$, say $\delta = 1/\sqrt{n}$, and start a random walk from a point s by stepping distance δ in a uniformly chosen random direction. (If this step takes us outside the body, we choose another direction, until we finally are able to make a step.) The stationary distribution of this walk is close to uniform. Mixing properties of this walk were analyzed in several papers [19, 16]. It turns out that the mixing time of the walk is $O(n^3)$ independently of the body (it may even be $O(n^2)$; this is an open question). Our results than say that a number of other mixing measures have the same order of magnitude.

Other algorithmic contexts where mixing times have been studied include sampling from log-concave distributions [13], matchings in graphs [15, 24], and Metropolis-type algorithms [11].

1.2 Background mathematical results

We set the stage by first quoting one standard theorem from each of the three areas mentioned initially. None of these theorems is recent. Theorem A, in explicit form, is due to Baxter and Chacon [3] in 1976, though seems implicit in earlier work of Dinges [12] and Rost [30] (see also Pitman [27]): extensions can be found in Revuz [28] and finite-state applications in Lovász and Winkler [20]. Theorem B is part of Theorem 16.0.2 of [23], who describe its history, tracing the various parts of the cycle of equivalences to dates between 1941 and 1980. Theorem C is from Aldous [1] in 1981.

Write $(X(t); t = 0, 1, 2, \ldots)$ for a Markov chain with transition kernel $P(x, A)$ on a measurable state space $\mathcal{X}$. Suppose an invariant probability measure π exists. Write $||\nu||$ for the total variation norm on signed measures on $\mathcal{X}$, so that for probability measures μ_1, μ_2 we have

$$||\mu_1 - \mu_2|| = 2 \sup_A |\mu_1(A) - \mu_2(A)|.$$

Here we have used the "functional analysis" normalization, even though in the "mixing time" literature it is common to divide the right sides by 2. Write $H_A \geq 0$ for the first hitting time on A. Write $E_\mu(\cdot)$ and $\mathcal{L}_\mu(\cdot)$ for expectation and distribution w.r.t. the initial distribution μ. For probability measures μ, θ define

$$h(\mu, \theta) = \inf\{E_\mu T : T \text{ is a randomized stopping time, } \mathcal{L}_\mu(X(T)) = \theta\}.$$

Abusing notation slightly, write $h(x, \theta)$ instead of $h(\delta_x, \theta)$ for the case of an initial distribution δ_x concentrated at x.

Consider the hypothesis

$$G(x, \cdot) \equiv \lim_{t \to \infty} \sum_{s=0}^{t} (P^s(x, \cdot) - \pi(\cdot)) \text{ exists } \forall x, \tag{1}$$

where the limit is w.r.t. total variation. If (1) holds, then $G(x, \cdot)$ is a signed measure with $G(x, \mathcal{X}) = 0$. It may not be true that $G(x, \cdot) \ll \pi$, but it is easy to see that $G(x, \cdot)$ decomposes as the sum of a positive measure singular w.r.t. π, and a signed measure with some density $g(x, \cdot)$ w.r.t. π. We call $g = g(x, y)$ the *recurrent potential density.*

Theorem A *Suppose (1). If the negative part of $(\mu - \sigma)G$ has a density $-\phi(y)$ w.r.t. π then*

$$h(\mu, \sigma) = \text{ess sup } \phi.$$

Otherwise, $h(\mu, \sigma) = \infty$.
To state the second theorem, define

$$d(t) = \sup_x ||P^t(x, \cdot) - \pi||. \tag{2}$$

If $d(t) \to 0$ the chain is called *uniformly ergodic*. It is well known that $d(t)$ is submultiplicative, so if $d(t) \to 0$ then the convergence is geometrically fast. Next, a *petite structure* is a collection $\{C, \mu, m, \delta\}$ where C is a subset of $\mathcal{X}$, μ is a probability distribution on $\mathcal{X}$, $m \geq 1$, $\delta > 0$ and

$$K_m(x, \cdot) \equiv (m+1)^{-1} \sum_{t=0}^{m} P^t(x, \cdot) \geq \delta\mu(\cdot) \ \ \forall x \in C. \tag{3}$$

Call C a *petite set* if it is part of some petite structure $\{C, \mu, m, \delta\}$.
Theorem B *The following are equivalent.*
(i) The chain is uniformly ergodic.
(ii) There exist $m < \infty$, $\delta > 0$ *and a probability measure* μ *such that*

$$P^m(x, \cdot) \geq \delta\mu(\cdot) \ \forall x.$$

(iii) The chain is aperiodic and there exists a petite set C *such that*

$$\sup_x E_x H_C < \infty.$$

(iv) The chain is aperiodic and there exist a petite set C*, a bounded function* $V(x) \geq 0$ *and constants* $\beta > 0$, $b < \infty$ *such that*

$$\begin{aligned} E_x V(X_1) - V(x) &\leq -\beta, \ x \notin C \\ &\leq b, \ x \in C. \end{aligned}$$

To state the third theorem, define

$$\mathcal{T}_{\text{mix}} = \sup_x h(x, \pi).$$

$$\mathcal{T}_{\text{set}} = \sup_{x, A: \pi(A) > 0} \pi(A) E_x H_A.$$

$$\mathcal{T}_{\text{continuize}} = \inf\{t : ||\tilde{P}^t(x, \cdot) - \pi|| \leq e^{-1} \ \forall x\}$$

where

$$\tilde{P}^t(x, \cdot) = \sum_{i=0}^{\infty} \frac{e^{-t} t^i}{i!} P^i(x, \cdot) \tag{4}$$

is the transition kernel for the associated continuous-time chain. Finally, in the finite-state case, where $h(x, y)$ and $h(\pi, y)$ are finite for all x, y, we may define

$$\mathcal{T}_{\text{G}} = \sup_x \sum_y |h(x, y) - h(\pi, y)| \ \pi(y).$$

So $\mathcal{T}_{\text{G}}$ measures variability of mean hitting times as a function of starting state. (See (9) for a more generally-applicable redefinition of $\mathcal{T}_{\text{G}}$).

Theorem C *For each pair* $(\mathcal{T}_i, \mathcal{T}_j)$ *from* $\{\mathcal{T}_{\text{mix}}, \mathcal{T}_{\text{set}}, \mathcal{T}_{\text{G}}, \mathcal{T}_{\text{continuize}}\}$ *there is a constant* $K_{i,j} < \infty$ *such that for every irreducible* reversible *chain on every* finite *state space* $\mathcal{X}$,

$$\mathcal{T}_i \leq K_{i,j}\mathcal{T}_j.$$

Though the hypotheses and conclusions of Theorems A-C are somewhat different, it seems intuitively clear that they refer in part to the same idea: the relation between means of stopping times and convergence to stationarity. Means of stopping times are explicit in Theorem A, in Theorem B(iii) and the definitions of $\mathcal{T}_{\text{mix}}$ and $\mathcal{T}_{\text{set}}$ in Theorem C. And as regards convergence, the parameters $\mathcal{T}_{\text{mix}}$ and $\mathcal{T}_{\text{continuize}}$ in Theorem C provide quantifications of the uniform ergodicity assertion in Theorem B (i), while in Theorem A one expects the size of the measures $G(x,\cdot)$ to be related to the speed of convergence of the sum in (1).

1.3 Statement of new results

The goal of our paper, in brief, is to establish quantitative bounds like those in Theorem C in the continuous-space setting of Theorem B.

The setting we shall adopt is best described as "uniform ergodicity, but without assuming aperiodicity". More precisely, define $\bar{d}(t)$ as "$d(t)$ for the uniformly-sampled chain", that is

$$\bar{d}(t) = \sup_x ||P_x(X(U_t) \in \cdot) - \pi(\cdot)|| = \sup_x ||K_t(x,\cdot) - \pi(\cdot)||$$

where K_t was defined at (3) and where U_t denotes a random variable distributed uniformly on $\{0,1,2,\ldots,t\}$, independent of the chain. If $\bar{d}(t) \to 0$, call the chain *uniformly ergodic under sampling* (UES). Minor modifications to the proof of Theorem B would establish the parallel result

Theorem B* *The following are equivalent.*
(i) The chain is UES.
(ii) There exist $m \geq 1$, $\delta > 0$ *and a probability measure* μ *such that*

$$K_m(x,\cdot) \geq \delta\mu(\cdot)\ \forall x.$$

(iii,iv): the corresponding statements in Theorem B, without the "aperiodic" assertion.

Our goal is to give a "quantitative" version of Theorem B^*. That is, we replace assertions of the form

> there exist objects $\{a, b, \ldots\}$ satisfying requirements $\{R, S, T \ldots\}$

by parameters $\mathcal{T}$ defined via

> $\mathcal{T}$ is the minimum, over all choices of objects $\{a, b, \ldots\}$ satisfying requirements $\{R, S, T \ldots\}$, of a certain numerical function of $\{a, b, \ldots\}$.

Applying this procedure to the four parts of Theorem B^* leads to the following four definitions.

$$\mathcal{T}_{\text{uniform}}(c) = \min\{t : \bar{d}(t) \leq c\},\ 0 < c < 1. \tag{5}$$

$$\mathcal{T}_{\text{minorize}} \text{ is the infimum of } \delta^{-1}m \text{ over all } \{m, \delta, \mu\} \text{ in Theorem } B^* \text{ (ii).} \tag{6}$$

$$\mathcal{T}_{\text{petite}} \text{ is the infimum of } \delta^{-1}(m + \sup_x E_x H_C) \text{ over all petite structures } \{C, \mu, m, \delta\}. \tag{7}$$

$\mathcal{T}_{\text{drift}}$ is the infimum of $\delta^{-1}(m + \max(b, \beta^{-1} \sup_x V(x)))$ over all petite structures

$$\{C, \mu, m, \delta\} \text{ and all } \{V, \beta, b\} \text{ satisfying the inequality in Theorem B (iv).} \tag{8}$$

But it is almost obvious (see section 6.2) that in fact $\mathcal{T}_{\text{drift}} = \mathcal{T}_{\text{petite}}$, so we need not consider $\mathcal{T}_{\text{drift}}$ separately. We shall also consider parameters equal or similar to those in Theorem C. Redefine $\mathcal{T}_{\text{G}}$ as

$$\mathcal{T}_{\text{G}} = \sup_x ||G(x, \cdot)||. \tag{9}$$

This is consistent with the previous definition in the finite-state case, where it is classical (see the discussion of the fundamental matrix in [17]) that $h(i,j) = (G(j,j) - G(i,j))/\pi(j)$, $h(\pi, j) = G(j,j)/\pi(j)$ and so $\pi(j)(h(i,j) - h(\pi, j)) = -G(i,j)$. Next, we give two weaker variants of $\mathcal{T}_{\text{mix}}$. The first requires only approximately attaining the target distribution π:

$$\mathcal{T}_{\text{stop}}(c) = \sup_x \inf\{E_x T : ||\mathcal{L}_x X(T) - \pi|| \leq c\},\ 0 < c < 1. \tag{10}$$

The second replaces π by some target distribution σ of our choice.

$$\mathcal{T}_{\text{forget}} = \inf_\sigma \sup_\mu h(\mu, \sigma) = \inf_\sigma \sup_x h(x, \sigma). \tag{11}$$

Theorem 1 *A chain is UES if and only if one of the parameters* $\{\mathcal{T}_{\text{G}}, \mathcal{T}_{\text{set}}, \mathcal{T}_{\text{forget}}, \mathcal{T}_{\text{minorize}}, \mathcal{T}_{\text{petite}}, \mathcal{T}_{\text{uniform}}(c), 0 < c < 1, \mathcal{T}_{\text{stop}}(c), 0 < c < 1\}$ *is finite, in which case all of these parameters are finite. For each pair* $(\mathcal{T}_i, \mathcal{T}_j)$ *of parameters in that set, there is a constant* $K_{i,j} < \infty$ *such that* $\mathcal{T}_i \leq K_{i,j} \mathcal{T}_j$ *for every UES chain.*

More concisely, call these parameters *equivalent*. In addition to quantifying Theorem B^*, Theorem 1 shows that part of Theorem C remains true in the non-reversible setting. One might hope that $\mathcal{T}_{\text{mix}}$ remained equivalent to these parameters in the non-reversible setting, but this hope is dashed by

Example 2 *The winning streak chain.*

Take $\mathcal{X} = \{0, 1, 2, \ldots\}$ and $P(x, x+1) = p, P(x, 0) = 1 - p$ for fixed $0 < p < 1$. So $\pi(x) = (1-p)p^x$. By considering $\sigma = \delta_0$ we have $\mathcal{T}_{\text{forget}} = 1/(1-p)$. But an elementary calculation gives $E_0 H_x = \frac{1}{\pi(x)} - \frac{1}{1-p}$ and so $\mathcal{T}_{\text{mix}} \geq h(0, \pi) = \sum_x \pi(x) E_0 H_x = \infty$. □

It turns out that $\mathcal{T}_{\text{mix}}$ is related instead to yet another parameter. Define, for $0 < c < 1$,

$$\mathcal{T}_{\text{separate}}(c) = \min\{t : P_x(X(U_t) \in \cdot) \geq (1-c)\pi(\cdot)\ \forall x\} \tag{12}$$

For a UES chain the parameters $\mathcal{T}_{\text{mix}}$ and $\mathcal{T}_{\text{separate}}(c)$ may be infinite, but they are equivalent.

Theorem 3 $\mathcal{T}_{\text{mix}} \leq \frac{1}{2(1-c)} \mathcal{T}_{\text{separate}}(c)$; $0 < c < 1$. *Conversely, if* $1/c$ *is an integer then* $\mathcal{T}_{\text{separate}}(c) \leq \frac{4}{c^2} \mathcal{T}_{\text{mix}}$.

To connect this with recurrent potential, note that Theorem A gives

$$h(x, \pi) = \text{ess sup}_y\ (-g(x,y)). \tag{13}$$

Thus $\mathcal{T}_{\text{mix}}$ can be defined directly in terms of the recurrent potential density g as

$$\mathcal{T}_{\text{mix}} = \sup_x \text{ess sup}_y\ (-g(x,y)).$$

We should emphasize that Theorems 1 and 3 are not really difficult or deep. Our proofs use the same mix of ingredients as the proof of Theorem C, with occasional modifications which use Theorem A in place of considering mean hitting times on single states. Textbooks sometimes leave the impression that general-state chains require different techniques than finite-state chains, but from the quantitative viewpoint this isn't so: our proofs were originally written for finite-state chains but then extended to the UES setting with only minor rephrasing.

Some further results dealing with time-reversals (and requiring some measure-theoretic technicalities) will be given in section 5.

1.4 Interpretation of results

In the setting of section 1.1, there is a specific Markov chain which we use to obtain samples for some ultimate algorithmic use. For an analysis of the number of steps needed, the ultimate use affects the notion of "mixing time" needed. The significance of our results is that one can to some extent "decouple" mathematical analysis of the chain from the ultimate algorithmic use of the samples, because many different mixing times are equivalent up to constants. In other words, for a sequence of Markov chains with size-parameter n, Theorem 1 says there is a well-defined "order of magnitude of mixing times" $t(n)$ such that each parameter in Theorem 1 is $\Theta(t(n))$. In contrast, Theorems B and B^* are typically uninformative in this context.

Of course, to actually bound mixing times for specific chains is a more interesting and important problem. Our results do not directly help, beyond providing flexibility in what one needs to prove to obtain an order-of-magnitude bound. (For instance, in obtaining upper bounds the freedom of choice of σ in $\mathcal{T}_{\mathrm{forget}}$ may be helpful; in obtaining lower bounds the freedom of choice of x and A in $\mathcal{T}_{\mathrm{set}}$ may be helpful.)

We remark that most of the algorithmic problems of section 1.1 are so hard that one cannot get the correct order of magnitude bound for mixing times. On the other hand, in the more highly-structured setting of card-shuffling and random walks on groups, one can often do rather precise calculations of mixing times: see for instance the analysis [4] of the riffle shuffle. Our work is perhaps most relevant to examples whose complexity is such that one can get only the correct order of magnitude. Here are two recent examples. Chung and Graham [5] analyze the chain on states $\{0,1\}^n$ in which two coordinates i, j are chosen at random, and the parity of x_i is changed if $x_j = 1$. They show the mixing time is $\Theta(n \log n)$. Diaconis and Saloff-Coste [10] study simple symmetric random walk on a convex subset of the two-dimensional lattice, and show that the mixing time is $\Theta(\text{diameter}^2)$.

In the setting of random walks on groups, the main focus of study has been

$$\mathcal{T}(c) = \min\{t : d(t) \leq c\} \tag{14}$$

and the *cut-off phenomenon* [8], in place of the time-averaged analog $\mathcal{T}_{\mathrm{uniform}}(c)$. While this is natural in examples, there seems no elegant "equivalence theory" analogous to Theorem 1 for $\mathcal{T}(c)$, and indeed Corollary 9 later indicates how $\mathcal{T}(c)$ may behave undesirably. The underlying difficulty is to quantify aperiodicity. Since periodicity is irrelevant for algorithmic sampling purposes, the Theorem 1 mixing times are more natural in that context.

2 Some technical tools

The minorization construction. Let $(Y(t))$ have kernel Q satisfying the minorization condition $Q(x, \cdot) \geq \delta\mu(\cdot)\, \forall x$, for some $\delta > 0$ and some probability measure μ. Obviously we can construct a randomized stopping time T with geometric(δ) distribution such that $Y(T)$ has distribution μ and is independent of both the starting state and the value of T: in particular

$$\mathcal{L}Y(T) = \mu, \quad ET = 1/\delta. \tag{15}$$

Elaborations of Theorem A. We need to use some ingredients of the proof of Theorem A, so we shall outline parts of the proof. See [3],[21],[2] for more details.

Fix μ, σ and consider a stopping time T with $E_\mu T < \infty$ and $\mathcal{L}_\mu X(T) = \sigma$. Write $\psi(\cdot) = E_\mu \sum_{t=0}^{T-1} 1_{(X(t)\in\cdot)}$. Then ψ is one solution of the identity $\psi - \psi P = \mu - \sigma$. Assuming (1), a particular solution of this identity is $\psi_0 = (\mu - \sigma)G$ and then the general solution is $\psi = (\mu - \sigma)G + c\pi$ for some constant c. Since $\psi(\mathcal{X}) = E_\mu T$ we have $c = E_\mu T$. To summarize:

$$\psi(\cdot) \equiv E_\mu \sum_{t=0}^{T-1} 1_{(X(t)\in\cdot)} = (\mu - \sigma)G + (E_\mu T)\pi. \tag{16}$$

Since $\psi \geq 0$, (16) implies

$$E_\mu T \geq \text{ess sup } \frac{d(\sigma - \mu)G}{d\pi}. \tag{17}$$

The proof of Theorem A is completed via a *filling scheme* construction, which defines inductively a certain decreasing sequence A_t of random subsets such that

$$T = \min\{t : X(t) \in A_t\} \tag{18}$$

achieves equality in (17).

3 Proof of Theorem 1

The proof is structured as three cycles of inequalities, in which $0 < c < 1$ is arbitrary. The first cycle is

$$\mathcal{T}_{\text{stop}}(c) \leq 4\mathcal{T}_{\text{set}}/c^2. \tag{19}$$

$$\mathcal{T}_{\text{set}} \leq \mathcal{T}_{\text{G}} \leq \frac{2}{1-c}\mathcal{T}_{\text{stop}}(c). \tag{20}$$

These imply that $\mathcal{T}_{\text{stop}}(c) \leq \frac{8}{c^2(1-c')}\mathcal{T}_{\text{stop}}(c')$. So the parameters $\{\mathcal{T}_{\text{set}}, \mathcal{T}_{\text{G}}, \mathcal{T}_{\text{stop}}(c), 0 < c < 1\}$ are all *equivalent*, i.e. ratios are bounded by constants. The second cycle is

$$\mathcal{T}_{\text{stop}}(c) \leq \frac{1}{2}\mathcal{T}_{\text{uniform}}(c). \tag{21}$$

$$\mathcal{T}_{\text{uniform}}(c') \leq \frac{2\mathcal{T}_{\text{stop}}(c)}{c' - c}, \ c' > c. \tag{22}$$

These show that the parameters $\{\mathcal{T}_{\text{uniform}}(c), 0 < c < 1\}$ are equivalent to the parameters above. The third cycle is

$$\mathcal{T}_{\text{forget}} \leq \mathcal{T}_{\text{petite}} \leq \mathcal{T}_{\text{minorize}} \leq 43\mathcal{T}_{\text{uniform}}(1/4). \tag{23}$$

$$\mathcal{T}_{\text{uniform}}(c) \leq \frac{4}{c}\mathcal{T}_{\text{forget}}. \tag{24}$$

These imply equivalence of the remaining parameters $\{\mathcal{T}_{\text{forget}}, \mathcal{T}_{\text{petite}}, \mathcal{T}_{\text{minorize}}\}$.

The first cycle. Fix some initial distribution. The fact that a minimal-mean stopping time T with $\mathcal{L}X(T) = \pi$ can be constructed via a filling scheme (18) implies

$$P(T \geq t) \leq P(H_{A_t} \geq t), \;\; P(T \geq t) \leq \pi(A_t).$$

Using the definition of $\mathcal{T}_{\rm set}$ and the inequalities above,

$$P(T \geq t) \leq P(H_{A_t} \geq t) \leq t^{-1} E H_{A_t} \leq \frac{\mathcal{T}_{\rm set}}{t\pi(A_t)} \leq \frac{\mathcal{T}_{\rm set}}{tP(T \geq t)}$$

and so $P(T \geq t) \leq \sqrt{\mathcal{T}_{\rm set}/t}$, in particular

$$P(T > \lfloor 4\mathcal{T}_{\rm set}/c^2 \rfloor) \leq c/2.$$

But $||\mathcal{L}X(\min(T,t)) - \pi|| \leq 2P(T > t)$ and $\min(T,t)$ is a stopping time with mean at most t, so by definition of $\mathcal{T}_{\rm stop}(c)$ we have $\mathcal{T}_{\rm stop}(c) \leq 4\mathcal{T}_{\rm set}/c^2$, which is (19).

Fix μ and A and write $\sigma = \mathcal{L}_\mu X(H_A)$. Consider (16) with $T = H_A$: since $\psi = 0$ on A we have $\frac{d(\sigma-\mu)G}{d\pi} = E_x H_A$ on A. Then

$$\pi(A) E_\mu H_A = (\sigma - \mu)G(A) \leq \tfrac{1}{2}||(\sigma - \mu)G|| \leq \mathcal{T}_{\rm G}.$$

So by definition of $\mathcal{T}_{\rm set}$ we have $\mathcal{T}_{\rm set} \leq \mathcal{T}_{\rm G}$, which is the first inequality of (20).

Fix x. By Theorem A, for any distribution σ and any set A,

$$h(x,\sigma) \geq \frac{(\sigma - \delta_x)G(A)}{\pi(A)}.$$

Rearranging, and using the fact $\pi G(\cdot) = 0$ (a simple consequence of (1)),

$$-G(x,A) \leq h(x,\sigma) - \sigma G(A) = h(x,\sigma) + (\pi - \sigma)G(A) \leq h(x,\sigma) + \tfrac{1}{2}||\pi - \sigma||\mathcal{T}_{\rm G}.$$

By definition of $\mathcal{T}_{\rm stop}(c)$, minimizing over $\{\sigma : ||\sigma - \pi|| \leq c\}$ gives

$$-G(x,A) \leq \mathcal{T}_{\rm stop}(c) + \tfrac{c}{2}\mathcal{T}_{\rm G}.$$

Applying this to $A = \{y : g(x,y) < 0\}$ and maximizing over x gives, by definition of $\mathcal{T}_{\rm G}$,

$$\tfrac{1}{2}\mathcal{T}_{\rm G} \leq \mathcal{T}_{\rm stop}(c) + \tfrac{c}{2}\mathcal{T}_{\rm G}.$$

In other words, $\mathcal{T}_{\rm G} \leq \frac{2}{1-c}\mathcal{T}_{\rm stop}(c)$, which is the second inequality of (20).

The second cycle. Inequality (21) follows from the definitions of $\mathcal{T}_{\rm stop}(c)$ and $\mathcal{T}_{\rm uniform}(c)$, because $EU_t = t/2$. Now fix x, consider T as in the definition of $\mathcal{T}_{\rm stop}(c)$, so that $E_x T \leq \mathcal{T}_{\rm stop}(c)$ and $||\mathcal{L}_x X(T) - \pi|| \leq c$. The latter implies $||\mathcal{L}_x X(T + U_t) - \pi|| \leq c$. And

$$\begin{aligned} ||\mathcal{L}_x X(U_t) - \mathcal{L}_x X(T+U_t)|| &\leq ||\mathcal{L}_x U_t - \mathcal{L}_x(T+U_t)|| \\ &\leq 2t^{-1}E_x T \text{ because } ||\mathcal{L}U_t - \mathcal{L}(a+U_t)|| \leq 2t^{-1}a \\ &\leq 2t^{-1}\mathcal{T}_{\rm stop}(c). \end{aligned}$$

By the triangle inequality

$$||\mathcal{L}_x X(U_t) - \pi|| \leq c + 2t^{-1}\mathcal{T}_{\rm stop}(c). \tag{25}$$

In other words, $\mathcal{T}_{\text{uniform}}(c') \leq t$ whenever $c + 2t^{-1}\mathcal{T}_{\text{stop}}(c) \leq c'$. Rearranging gives $\mathcal{T}_{\text{uniform}}(c') \leq \frac{2\mathcal{T}_{\text{stop}}(c)}{c'-c}$, $c' > c$, which is (22).

The third cycle. We start by proving (24), the proof being similar to the proof of (22). Let σ attain the *inf* in the definition of $\mathcal{T}_{\text{forget}}$. So given an initial state x, we can choose S_x and S_π such that

$$\mathcal{L}_x X(S_x) = \sigma, E_x S_x \leq \mathcal{T}_{\text{forget}}, \mathcal{L}_\pi X(S_\pi) = \sigma, E_\pi S_\pi \leq \mathcal{T}_{\text{forget}}.$$

Then $\mathcal{L}_x X(S_x + U_t) = \mathcal{L}_\pi X(S_\pi + U_t) = \mathcal{L}_\sigma X(U_t)$, and so

$$\begin{aligned} ||\mathcal{L}_x X(U_t) - \mathcal{L}_\pi X(U_t)|| &\leq ||\mathcal{L}_x X(U_t) - \mathcal{L}_x X(S_x + U_t)|| + ||\mathcal{L}_\pi X(U_t) - \mathcal{L}_\pi X(S_\pi + U_t)|| \\ &\leq ||\mathcal{L}_x U_t - \mathcal{L}_x(S_x + U_t)|| + ||\mathcal{L}_\pi(U_t) - \mathcal{L}_\pi(S_\pi + U_t)|| \\ &\leq 2t^{-1}(E_x S_x + E_\pi S_\pi) \\ &\leq 4t^{-1}\mathcal{T}_{\text{forget}}. \end{aligned}$$

Since $\mathcal{L}_\pi X(U_t) = \pi$, we have established (24).

Given a petite structure $\{C, \mu, m, \delta\}$, we have

$$Q(x, \cdot) \equiv P_x(X(H_C + U_m) \in \cdot) \geq \delta\mu(\cdot) \quad \forall x \in \mathcal{X}.$$

Use (15) to construct a stopping time attaining distribution μ with mean $\leq \delta^{-1}\sup_x E_x(H_C + U_m) = \delta^{-1}(m/2 + \sup_x E_x H_C)$. This implies $\mathcal{T}_{\text{forget}} \leq \mathcal{T}_{\text{petite}}$, the first inequality in (23). The second inequality, $\mathcal{T}_{\text{petite}} \leq \mathcal{T}_{\text{minorize}}$, is immediate by taking $C = \mathcal{X}$ in the definition of $\mathcal{T}_{\text{petite}}$. The third inequality requires a preliminary lemma.

Lemma 4 *Let Q be a transition kernel for which π is invariant, and suppose*

$$\sup_x ||Q(x, \cdot) - \pi(\cdot)|| \leq 1/4. \tag{26}$$

Then there exists a set A with $\pi(A) \geq 1/2$ such that, for all probability distributions μ,

$$\frac{d\mu Q}{d\pi} \geq \tfrac{1}{4}(1 - ||\mu - \pi||) \text{ on } A.$$

Proof. Let B_x be the set where $\frac{dQ(x,\cdot)}{d\pi} \leq 1/2$, or more precisely where the measure $Q(x, \cdot) - \frac{1}{2}\pi(\cdot)$ is negative. Write $B = \{(x, y) : y \in B_x\} \subset \mathcal{X} \times \mathcal{X}$. For each x we have $Q(x, B_x) \leq \frac{1}{2}\pi(B_x)$, and so

$$\tfrac{1}{2}\pi(B_x) \leq \pi(B_x) - Q(x, B_x) \leq \tfrac{1}{2}||Q(x, \cdot) - \pi(\cdot)||.$$

Using (26), $\pi(B_x) \leq 1/4$ and hence $\pi \times \pi(B) \leq 1/4$. Write $B^y = \{x : (x, y) \in B\}$ and define $A = \{y : \pi(B^y) \leq 1/2\}$. Then $\pi(A^c) \leq \frac{\pi\times\pi(B)}{1/2} \leq 1/2$. By definition we have $\frac{dQ(x,\cdot)}{d\pi}(y) \geq 1/2$ for $x \in \mathcal{X} \setminus B^y$, and so

$$\frac{d\mu Q}{d\pi}(y) \geq \tfrac{1}{2}\mu(\mathcal{X} \setminus B^y) \geq \tfrac{1}{2}(\pi(\mathcal{X} \setminus B^y) - \tfrac{1}{2}||\mu - \pi||) \geq \tfrac{1}{2}(\tfrac{1}{2} - \tfrac{1}{2}||\mu - \pi||) \text{ on } A,$$

the final inequality by definition of A. $\square$

Now set $t = \mathcal{T}_{\text{uniform}}(1/4)$ and let U and U' be independent, uniform on $\{0, 1, \ldots, t\}$. Let Q be the kernel associated with $X(U')$, so that (26) holds by definition. Let A be the set guaranteed by the lemma. For fixed x, write $\mu = \mathcal{L}_x X(U)$, and then

$$\frac{d\mathcal{L}_x X(U + U')}{d\pi} = \frac{d\mu Q}{d\pi} \geq \frac{1}{4}(1 - \frac{1}{4}) = \frac{3}{16} \text{ on } A$$

by the lemma. In other words, if $R(x,\cdot)$ is the kernel associated with $X(U+U')$ and if we set $\mu = \pi(\cdot|A)$ then $R(x,\cdot) \geq \frac{3}{32}\mu(\cdot)\ \forall x$, because $\pi(A) \geq 1/2$ by the lemma. It is elementary that $P(U_{2t}=i) \geq \frac{1}{2}P(U+U'=i)\ \forall i$, and so $\mathcal{L}_x X(U_{2t}) \geq \frac{3}{64}\mu\ \forall x$. So by definition of $\mathcal{T}_{\text{minorize}}$ we have $\mathcal{T}_{\text{minorize}} \leq \frac{64}{3}2t \leq 43\mathcal{T}_{\text{uniform}}(1/4)$.

4 Proof of Theorem 3

Fix t, write

$$s(t) = \inf\{c : P_x(X(U_t) \in \cdot) \geq (1-c)\pi(\cdot)\ \forall x\}$$

and consider the chain Y with transition kernel $K_t(x,\cdot) = P_x(X(U_t) \in \cdot)$. Construction (15) gives a stopping time S for Y satisfying $\mathcal{L}Y(S) = \pi$ and $ES = 1/(1-s(t))$. This in turn specifies a stopping time $T = U_t^{(1)} + \ldots + U_t^{(S)}$ for X satisfying $\mathcal{L}X(T) = \pi$ and $ET = (ES)(EU_t) = \frac{t}{2(1-s(t))}$. Putting $t = \mathcal{T}_{\text{separate}}(c)$ gives

$$\mathcal{T}_{\text{mix}} \leq \frac{\mathcal{T}_{\text{separate}}(c)}{2(1-c)}.$$

For a reverse inequality, the central idea is contained in the following lemma, analogous to Lemma 4. Write $N(t,\cdot) = \sum_{n=0}^{t-1} 1_{(X(n)\in\cdot)}$.

Lemma 5 $\frac{d}{d\pi}E_\sigma N(t,\cdot) \geq \left(1 - \frac{||\sigma-\pi||}{2}\right) t \ - \mathcal{T}_{\text{mix}}$.

Proof. Write $\delta = ||\sigma-\pi||/2$. Decompose the initial distribution σ as $\sigma = \pi + \delta(\alpha-\beta)$, where $\alpha = (\sigma-\pi)^+/\delta$ and $\beta = (\pi-\sigma)^+/\delta$. Note that

$$\text{ess sup } \frac{d\beta}{d\pi} \leq 1/\delta \tag{27}$$

and that

$$E_\sigma N(t,\cdot) = E_\pi N(t,\cdot) + \delta(E_\alpha N(t,\cdot) - E_\beta N(t,\cdot)) \geq t\pi(\cdot) - \delta E_\beta N(t,\cdot). \tag{28}$$

Next observe that for any stopping time S with $\mathcal{L}_\pi X(S) = \pi$ (note this refers to the *stationary* chain) we have by (16)

$$E_\pi N(S,\cdot) = (E_\pi S)\ \pi(\cdot). \tag{29}$$

By definition of $\mathcal{T}_{\text{mix}}$ there exists a stopping time T with $E_x T \leq \mathcal{T}_{\text{mix}}$ and $\mathcal{L}_x X(T) = \pi$ for all x. So

$$\begin{aligned}
\int \beta(dx) E_x N(T,\cdot) &\leq \frac{1}{\delta}\int \pi(dx) E_x N(T,\cdot) \text{ by (27)} \\
&= \frac{1}{\delta}\left(\int \pi(dx) E_x T\right)\ \pi(\cdot) \text{ by (29)} \\
&\leq \frac{1}{\delta}\mathcal{T}_{\text{mix}}\ \pi(\cdot)
\end{aligned}$$

and then

$$\begin{aligned}
E_\beta N(t,\cdot) &\leq \int \beta(dx) E_x N(T+t,\cdot) \\
&= \int \beta(dx) E_x N(T,\cdot)\ + E_\pi N(t,\cdot) \\
&\leq \frac{1}{\delta}\mathcal{T}_{\text{mix}}\ \pi(\cdot)\ + t\pi(\cdot).
\end{aligned}$$

The lemma now follows, using (28). $\Box$

To deduce the reverse inequality in Theorem 3, fix an initial state x and a time s. Then

$$E_x N(s+t,\cdot) \geq E_x N(U_s+t,\cdot) = E_\sigma N(t,\cdot)$$

for $\sigma = \mathcal{L}_x X(U_s)$. By (25) with $c=0$ we have $||\sigma - \pi|| \leq 2\mathcal{T}_{\mathrm{mix}}/s$ and so using Lemma 5

$$\frac{d}{d\pi} E_x N(s+t,\cdot) \geq t - \mathcal{T}_{\mathrm{mix}} t/s - \mathcal{T}_{\mathrm{mix}}.$$

In other words,

$$\frac{d}{d\pi} K_{s+t-1}(x,\cdot) \geq \frac{t - \mathcal{T}_{\mathrm{mix}} t/s - \mathcal{T}_{\mathrm{mix}}}{s+t-1} \geq \frac{t - (1-\frac{t}{s})\mathcal{T}_{\mathrm{mix}}}{s+t}.$$

Provided $1/c$ is an integer we may set $s = \frac{2}{c}\ \mathcal{T}_{\mathrm{mix}}$ and $t = \left(\frac{4}{c^2} - \frac{2}{c}\right) \mathcal{T}_{\mathrm{mix}}$. The right side becomes $1-c$, so we have proved

$$\mathcal{T}_{\mathrm{separate}}(c) \leq s+t = \frac{4\mathcal{T}_{\mathrm{mix}}}{c^2}. \tag{30}$$

5 Time-reversals

A kernel P with stationary distribution π and the time-reversed kernel P^* are related by the following identity for measures on $\mathcal{X} \times \mathcal{X}$.

$$\pi(dx)P^*(x,dy) = \pi(dy)P(y,dx). \tag{31}$$

It is perhaps surprising that for a UES chain the time-reversed (or *dual*) chain need not be UES. For instance, the time-reversal of Example 2 is the chain with $p^*(x,x-1) = 1, x \geq 1$ and $p(0,x) = (1-p)p^x, x \geq 0$, which is plainly not UES. This lack of symmetry suggests study of the class of processes whose time-reversals are UES. For each parameter $\mathcal{T}$ we may define a parameter $\mathcal{T}^*$ as "$\mathcal{T}$ for P^*". For instance, $\mathcal{T}^*_{\mathrm{G}} = \sup_x ||G^*(x,\cdot)||$ where G^* is defined in terms of P^* as in (1). Theorem 1 implies equivalence of the "starred" parameters therein, but what does this mean for the P-chain? The parameters involving stopping times for the P^*-chain have no very clear interpretation as parameters for the P-chain, but it turns out (Lemma 6) that the remaining parameters $\{\mathcal{T}^*_{\mathrm{G}}, \mathcal{T}^*_{\mathrm{minorize}}, \mathcal{T}^*_{\mathrm{uniform}}(c)\}$ can be expressed directly in terms of the P-chain. But first we must deal with a technical issue. If we use (31) as a *definition* of the time-reversed kernel $P^*(x,\cdot)$ of a given kernel P, then $P^*(x,\cdot)$ is defined uniquely only up to π-null sets of x-values. This matters because the parameters $\mathcal{T}$ were defined as *sups*, rather than *ess sups*, over x. Issues like this are frequently resolved by imposing topological assumptions, but for our purposes we may just adopt the following simple though inelegant assumption.

Assumption. P and P^* are related by (31). Furthermore, in the definition of each parameter $\mathcal{T}$ for both P and P^*, using $\sup_x$ and ess $\sup_x$ give the same value.

It is straightforward to check that, given P and P^* related by (31), we can delete a single π-null set from $\mathcal{X}$ so that the second requirement holds on the remaining space. In this sense, the assumption involves "no loss of generality". To see the need for some such assumption, consider a finite state space chain containing both transient states and a single recurrent class $\mathcal{R}$. Then there exists a (unique) stationary distribution supported on $\mathcal{R}$, the chain is UES, so

Theorems 1 and 3 are meaningful (and true), for parameters defined as maxima over the whole state space. But there is no natural way to define P^* outside $\mathcal{R}$, and hence no natural way of making results like Theorem D true. Our technical assumption has the effect of pruning the state space down to $\mathcal{R}$.

Relation (31) easily implies that the density components $k_t(x,y)\pi(dy)$ and $g(x,y)\pi(dy)$ of $K_t(x,\cdot)$ and $G(x,\cdot)$ are related to their duals by symmetry:

$$k_t^*(x,y) = k_t(y,x), \quad g^*(x,y) = g(y,x)\ , \text{ a.e. } (\pi\times\pi). \tag{32}$$

Now using this symmetry and our technical assumption, it is easy to relate certain "starred" parameters to their unstarred versions. The parameters in Theorem 3 can be written as

$$\mathcal{T}_{\text{mix}} = \text{ ess sup}_{(x,y)}\ (-g(x,y))$$

$$\mathcal{T}_{\text{separate}}(c) = \min\{t:\ \text{ess inf}_{(x,y)}\ t^{-1}k_t(x,y) \geq (1-c)\}.$$

So by symmetry they are unchanged by time-reversal. The next lemma (proved in section 5.1) expresses the "starred" versions of certain parameters in Theorem 1 in terms of the P-chain.

Lemma 6 *(a)* $\mathcal{T}^*_{\text{minorize}}$ *is the infimum of* $\delta^{-1}m$ *over all triples* $\{m,\delta,V\}$ *such that*

$$V\geq 0;\ \int V\,d\pi = 1;\ K_m(x,\cdot)\geq \delta V(x)\ \pi(\cdot)\ \ \forall x.$$

(b)

$$\mathcal{T}^*_{\text{uniform}}(c) = \min\{t: d^*(t)\leq c\},\ \ 0<c<1$$

where

$$d^*(t) = 2\ \text{ess sup}_y \int (k_t(x,y)-1)^-\ \pi(dx)$$

where $k_t(x,y)\pi(dy)$ *is the density component of* $K_t(x,\cdot)$.
(c) Write $G = G^+ - G^-$ *for the Hahn-Jordan decomposition of* G *as a difference of positive kernels, and write* $|G| = G^+ + G^-$. *Then*

$$\mathcal{T}^*_{\text{G}} = ess\ sup \frac{d\pi|G|}{d\pi}.$$

Whereas the original parameters in Theorem 1 all explicitly involved *sup*s over initial states, the starred parameters in Lemma 6 have a different flavor: roughly, they involve approximate minorization at a terminal time. A more natural parameter with that flavor is

$$\mathcal{T}_{\text{reset}} = \int\ h(x,\pi)\ \pi(dx).$$

In section 5.1 we prove the simple bounds

Lemma 7 $\frac{1}{2}\mathcal{T}^*_{\text{G}} \leq \mathcal{T}_{\text{reset}} \leq \frac{1}{2}\mathcal{T}^*_{\text{minorize}}$.

Combining with Theorem 1 (applied to time-reversed chains) gives

Corollary 8 *The parameters* $\{\mathcal{T}_{\text{reset}}, \mathcal{T}^*_{\text{G}}, \mathcal{T}^*_{\text{minorize}}, \mathcal{T}^*_{\text{uniform}}(c), 0<c<1\}$ *are equivalent. The time-reversed chain is UES iff one (all) of these parameters are finite.*

We observed above that the parameters in Theorem 3 are unchanged by time-reversal. So (e.g. because $\mathcal{T}_{\mathrm{stop}}(c) \leq \mathcal{T}_{\mathrm{mix}}$) if $\mathcal{T}_{\mathrm{mix}}$ is finite then both the chain and its time-reversal are UES. The converse also holds: if a chain and its time-reversal are both UES then $\mathcal{T}_{\mathrm{mix}}$ is finite. (So our results imply the equivalence of $\{\mathcal{T}_{\mathrm{mix}}, \mathcal{T}_{\mathrm{set}}, \mathcal{T}_{\mathrm{G}}\}$ for reversible chains, which was part of Theorem C). In fact,

$$\mathcal{T}_{\mathrm{mix}} \leq 2(\mathcal{T}_{\mathrm{set}} + \mathcal{T}_{\mathrm{reset}}).$$

Because $A = \{x : h(x, \pi) \leq 2\mathcal{T}_{\mathrm{reset}}\}$ has $\pi(A) \geq 1/2$, and so $\sup_x E_x H_A \leq 2\mathcal{T}_{\mathrm{set}}$. Thus for any initial distribution, we run the chain until hitting A, then until an optimal stopping time attaining π, and this stopping time has mean $\leq \sup_x E_x H_A \ + 2\mathcal{T}_{\mathrm{reset}}$.

Theorems 1 and 3 reflect the spirit of Theorems B and C in dealing with inequalities. Theorem A is in the spirit of standard results on maximal coupling and minimal strong stationary times (see section 6.3) giving "optimal constructions" or "min-max characterizations". Lovász and Winkler [22] proved another remarkable identity in the same spirit.

Theorem D *For a finite-state irreducible chain,* $\mathcal{T}^*_{\mathrm{reset}} = \mathcal{T}_{\mathrm{forget}}$.

Assuming this extends to our general-space setting, one could use Theorem D in place of Lemma 7 in establishing Corollary 8.

5.1 Proofs

Proof of Lemma 6. In the definition of $\mathcal{T}_{\mathrm{minorize}}$, the probability measure μ must satisfy $\mu \ll \pi$. Setting $V = d\mu/d\pi$ and using the symmetry relation (32) leads to the stated expression for $\mathcal{T}^*_{\mathrm{minorize}}$. Next, we may rewrite the definition of $\bar{d}(t)$ as

$$\bar{d}(t) = 2 \text{ ess } \sup_x \int (k_t(x, y) - 1)^- \ \pi(dy)$$

and then the expression for $\mathcal{T}^*_{\mathrm{uniform}}(c)$ follows by symmetry. Similarly, writing $\mathcal{T}_{\mathrm{G}} = 2 \text{ ess } \sup_x \int g^-(x, y) \ \pi(dy)$ gives by symmetry

$$\mathcal{T}^*_{\mathrm{G}} = 2 \text{ ess } \sup_y \int g^-(x, y) \ \pi(dx) \tag{33}$$

and the stated expression for $\mathcal{T}^*_{\mathrm{G}}$ follows because $\pi G = 0$.

Proof of Lemma 7.

$$\begin{aligned} \tfrac{1}{2}\mathcal{T}^*_{\mathrm{G}} &= \text{ ess } \sup_y \int g^-(x, y) \ \pi(dx) \text{ by } (33) \\ &\leq \int \text{ ess } \sup_y (-g(x, y)) \ \pi(dx) \\ &= \int h(x, \pi) \ \pi(dx) \text{ by } (13) \\ &= \mathcal{T}_{\mathrm{reset}}. \end{aligned}$$

For the second inequality, consider $\{m, \delta, V\}$ as in the definition of $\mathcal{T}^*_{\mathrm{minorize}}$, so that the chain Y with kernel $Q = K_m$ satisfies

$$Q(x, \cdot) \geq \delta V(x) \pi(\cdot) \ \forall x,$$

Let $Y(0)$ have distribution π and define a randomized stopping time $T \geq 1$ for Y by

$$P(T = t + 1 | T > t, Y(t) = x, Y(t + 1) = y) = \delta V(x) \ \frac{1}{\frac{dQ(x,\cdot)}{d\pi}(y)}.$$

One can verify inductively that

$$\begin{aligned} P(Y(t) \in \cdot | T > t) &= \pi(\cdot) \\ P(Y(t+1) \in \cdot | T = t+1, Y(t) = x) &= \pi(\cdot) \\ P(T = t+1 | T > t) &= \delta. \end{aligned}$$

So $ET = 1/\delta$ and $Y(T)$ has distribution π, independent of $Y(0)$. Since Y is the uniformly-sampled X-chain, this construction gives a stopping time S for X such that $X(S)$ has distribution π, independent of $X(0)$. So $\mathcal{T}_{\text{reset}} \leq ES = \frac{m}{2\delta}$ and thus $\mathcal{T}_{\text{reset}} \leq \frac{1}{2}\mathcal{T}^*_{\text{minorize}}$.

6 Discussion

6.1 Transforming a chain

Let $X(t)$ be a chain with forget time $\mathcal{T}_{\text{forget}}$. Corollary 9 gives three constructions of chains of the form $\widetilde{X}(t) = X(N(t))$ for different definitions of $N(t)$, and relates their forget times $\widetilde{\mathcal{T}}_{\text{forget}}$ to $\mathcal{T}_{\text{forget}}$.

Corollary 9 *(a) [jump chain]. Define*

$$N(t+1) = \min\{s > N(t) : X(s) \neq X(N(t))\}.$$

Then $\widetilde{\mathcal{T}}_{\text{forget}} \leq \mathcal{T}_{\text{forget}}$.

(b) [slowed-down chain]. Take $N(t+1) = N(t)$ *or* $N(t)+1$, *with*

$$P(N(t+1) = N(t) + 1 | X(N(t)) = x) = a(x).$$

Then $\widetilde{\mathcal{T}}_{\text{forget}} \leq a^{-1}\mathcal{T}_{\text{forget}}$, *where* $a = \inf_x a(x)$.

(c) [chain watched only on A]. Take

$$N(t+1) = \min\{s > N(t) : X(s) \in A\}.$$

Then $\widetilde{\mathcal{T}}_{\text{forget}} \leq \mathcal{T}_{\text{forget}}$.

The proofs are immediate, by considering the minimizing σ in the definition of $\mathcal{T}_{\text{forget}}$, and using $\tilde{\sigma} = \sigma$ (cases (a,b)) or $\tilde{\sigma} = P_\sigma(X(H_A) \in \cdot)$ (case (c)) as target distributions in the definition of $\widetilde{\mathcal{T}}_{\text{forget}}$. Intuitively, any reasonable definition of "mixing time" should satisfy similar inequalities. But note that with the traditional definition using total variation at fixed times (14), inequality (a) fails (the jump chain may be periodic) and the other inequalities do not seem simple to establish.

6.2 Technical remarks

(a) We have used "uniform smoothing" rather than "geometric smoothing" throughout, though there is no essential difference. Our statement of Theorem B skips some further, similar-style, assertions in [23] Theorem 16.0.2. Our statement of the drift condition (iv) is superficially different from theirs, but is clearly equivalent. Our statement was chosen to highlight the quantitative equality $\mathcal{T}_{\text{petite}} = \mathcal{T}_{\text{drift}}$, which is a consequence of the following observations. Given a petite set C, the function $V(x) = E_x H_C$ satisfies the inequality in Theorem B (iv) with $\beta = 1$ and $b = \sup_x E_x H_C$. Conversely, if V satisfies the inequality in Theorem B (iv) then $E_x H_C \leq V(x)/\beta$ by the obvious supermartingale argument.

(b) In the deterministic chain $X(t) = t$ modulo n, the parameters in Theorems 1 and 3 are $\Theta(n)$. This example shows that in Theorem 1 we cannot replace $\mathcal{T}_{\text{uniform}}(c)$ by $\mathcal{T}_{\text{continuize}}$ or by any smoothing essentially weaker than uniform.

(c) We glossed over two related technical points. For a periodic chain the limit (1) defining G may not exist; and in the setting of Theorem 1 we do not know *a priori* that G exists. What is important about G is that it satisfies

$$(I - P)G = I - \Pi$$

where I is the identity kernel and $\Pi(x, \cdot) = \pi(\cdot)$. In the period-$d$ setting where P^d is uniformly ergodic on each cyclic component, we can modify (1) by taking averages over $\{t, t+1, \ldots, t+d-1\}$, and then the $t \to \infty$ average of these limits exists. Using the general-space decomposition of a periodic chain into cyclic components ([23] sec. 5.4.1; [29] sec. 6.3), one can show that a UES chain is of this form, and so G exists.

(d) In our proofs we assumed that infima in the definitions of parameters $\mathcal{T}$ are attained. Pedantically, we should have considered attaining $\mathcal{T} + \varepsilon$, and then let $\varepsilon \to 0$.

(e) The parameters $\mathcal{T}_{\text{G}}$ and $\mathcal{T}_{\text{mix}}$ give probabilistic interpretations of certain L^1 and L^∞ norms of G, so it is natural to ask whether the analogous L^2 norm $\sup_x \int_{y:g(x,y)<0} g^2(x, y)\pi(dy)$ has a probabilistic interpretation as a mixing time.

(f) Our proofs used the same general set of techniques as in [1], though at the level of detailed proofs the overlap with [1, 3, 20] is quite small.

(g) A more complicated example of a UES chain whose time-reversal is not UES can be found in [29] exercise 8.3.11.

6.3 Conceptual remarks

(a) The asymptotic geometric rate of convergence of a chain is controlled by its *spectral gap*. That parameter is rather different from our mixing time parameters. See [2] for a extensive discussion in the reversible setting.

(b) Fix an initial distribution μ and define $s(t) = \min\{c : \mathcal{L}_\mu X(t) \geq (1 - c)\pi\}$. A *strong stationary time* is a stopping time T such that $X(T)$ has distribution π and is independent of T (e.g. the minorization construction (15) gave a strong stationary time). Such a T must satisfy the inequalities $P(T > t) \geq s(t)\ \forall t$; and it is easy to construct an optimal strong stationary time T satisfying $P(T > t) = s(t)\ \forall t$. See Diaconis - Fill [9] for developments of such theory. This construction, and the conceptually similar notion [14, 18] of *maximal coupling*, are in the same spirit as Theorem A.

(c) One of the themes of [23] is a sequence of theorems, in the general format of Theorem B, which treat successively stronger notions of convergence (ergodicity, geometric ergodicity, V-uniform ergodicity, uniform ergodicity), and relate each to drift and "return time to petite sets" conditions. Their presentation thus emphasizes "general" results such as the existence of minorizing measures as a consequence of irreducibility, and the "split chain" construction. But results at that level of generality are inherently non-quantitative. We are deliberately approaching these results from the opposite direction in order to get quantitative results. Whether analogs of Theorem 1 hold for these more general notions of convergence is an interesting question.

(d) Informally, our parameters are defined to "scale as time". This is easier to formalize in continuous time: if $X(t)$ has parameter value $\mathcal{T}$ then $X^*(t) = X(ct)$ should have parameter value $c^{-1}\mathcal{T}$. For two such parameters in continuous time, the existence of some universal inequality $\mathcal{T}_2 \leq \psi(\mathcal{T}_1)$ clearly implies a linear inequality $\mathcal{T}_2 \leq K_{1,2}\mathcal{T}_1$. Thus the existence of linear inequalities in Theorems 1 and 3 is not surprising.

(e) A quite different setting where mixing times might be studied is "randomly-perturbed chaos". Consider a topological space $\mathcal{X}$ and a continuous function $f : \mathcal{X} \to \mathcal{X}$ for which π is invariant, and suppose we define kernels $P^{(n)}(x, \cdot)$ such that $P^{(n)}(x, \cdot) \to \delta_{f(x)}$ weakly and whose stationary distributions $\pi^{(n)} \to \pi$ in total variation. Theorem 1 implies there is a well-defined order of magnitude $\Theta(t(n))$ for the parameters therein, providing an indirect formalization of the time until the underlying deterministic process $(f^i(x); i \geq 0)$ becomes chaotic. Getting explicit results in this setting seems challenging, even in simple-looking examples like the following. Fix $0 \leq a < 1$ and consider the random walk on the reals modulo 1 whose step-distribution is Normal(a, σ^2) modulo 1. If $a = 0$, clearly the mixing time parameters in Theorems 1 and 3 are $\Theta(\sigma^{-2})$ as $\sigma \to 0$. For general a, the behavior of our parameters as $\sigma \to 0$ is related to equidistribution of $\{ia \bmod 1 : \ i = 1, 2, \ldots\}$. Heuristically it appears that for typical a the mixing times are $\Theta(\sigma^{-2/3})$, but this appears non-elementary to prove.

Acknowledgements. We thank J. Baxter and D. Revuz and R. Syski for helpful comments on the literature.

References

[1] D.J. Aldous. Some inequalities for reversible Markov chains. *J. London Math. Soc. (2)*, 25:564–576, 1982.

[2] D.J. Aldous and J.A. Fill. Reversible Markov chains and random walks on graphs. Book in preparation, 1997.

[3] J.R. Baxter and R.V. Chacon. Stopping times for recurrent Markov processes. *Illinois J. Math.*, 20:467–475, 1976.

[4] D. Bayer and P. Diaconis. Trailing the dovetail shuffle to its lair. *Ann. Appl. Probab.*, 2:294–313, 1992.

[5] F.R.K. Chung and R.L. Graham. Stratified random walks on the n-cube. Technical report, Univ. Pennsylvania, 1996.

[6] C. Dellacherie and P.-A. Meyer. *Probabilités et Potentiel: Théorie Discrète du Potentiel.* Hermann, Paris, 1983.

[7] P. Diaconis. *Group Representations in Probability and Statistics.* Institute of Mathematical Statistics, Hayward CA, 1988.

[8] P. Diaconis. The cut-off phenomenon in finite Markov chains. *Proc. Nat. Acad. Sci. USA*, 93:1659–1664, 1996.

[9] P. Diaconis and J.A. Fill. Strong stationary times via a new form of duality. *Ann. Probab.*, 18:1483–1522, 1990.

[10] P. Diaconis and L. Saloff-Coste. Nash inequalities for finite Markov chains. *J. Theoretical Probab.*, 9:459–510, 1996.

[11] P. Diaconis and L. Saloff-Coste. Walks on generating sets of groups. Technical report, Mathematics Dept., Harvard University, 1996.

[12] H. Dinges. Stopping sequences. In *Séminaire de Probabilités VIII*, pages 27–36. Springer-Verlag, 1974. Lecture Notes in Math. 381.

[13] A. Frieze, R. Kannan, and N. Polson. Sampling from log-concave distributions. *Ann. Appl. Probab.*, 4:812–837, 1994.

[14] S. Goldstein. Maximal coupling. *Z. Wahrsch. Verw. Gebiete*, 46:193–204, 1979.

[15] M. Jerrum and A. Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18:1149–1178, 1989.

[16] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O(n^5)$ volume algorithm. *Random Struct. Alg.*, 8:xxx, 1997.

[17] J.G. Kemeny and J.L. Snell. *Finite Markov Chains.* Van Nostrand, 1960.

[18] T. Lindvall. *Lectures on the Coupling Method.* Wiley, 1992.

[19] L. Lovász and M. Simonovits. Random walks in a convex body and an improved volume algorithm. *Random Struct. Alg.*, 4:359–412, 1993.

[20] L. Lovász and P. Winkler. Efficient stopping rules for Markov chains. In *Proc. 27th ACM Symp. Theory of Computing*, pages 76–82, 1995.

[21] L. Lovász and P. Winkler. Fast mixing in a Markov chain. In preparation, 1997.

[22] L. Lovász and P. Winkler. Reversal of Markov chains and the forget time. In preparation, 1997.

[23] S.P. Meyn and R.L. Tweedie. *Markov Chains and Stochastic Stability.* Springer-Verlag, 1993.

[24] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.

[25] E. Nummelin. *General Irreducible Markov Chains and Non-Negative Operators.* Cambridge University Press, 1984.

[26] S. Orey. *Limit Theorems for Markov Chain Transition Probabilities.* Van Nostrand, 1971.

[27] J.W. Pitman. Occupation measures for Markov chains. *Adv. in Appl. Probab.*, 9:69–86, 1977.

[28] D. Revuz. Remarks on the filling scheme for recurrent Markov chains. *Duke Math. J.*, 45:681–689, 1978.

[29] D. Revuz. *Markov Chains.* North-Holland, second edition, 1984.

[30] H. Rost. The stopping distributions of a Markov process. *Inventiones Math.*, 14:1–16, 1971.

[31] A. J. Sinclair. *Algorithms for Random Generation and Counting.* Birkhauser, 1993.

[32] R. Syski. *Passage Times for Markov Chains.* IOS Press, Amsterdam, 1992.