

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a preprint version which may differ from the publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/135168>

Please be advised that this information was generated on 2019-10-23 and may be subject to change.

TIGHT INEQUALITIES AMONG SET HITTING TIMES IN MARKOV CHAINS

SIMON GRIFFITHS, ROSS J. KANG, ROBERTO IMBUZEIRO OLIVEIRA,
AND VIRESH PATEL

ABSTRACT. Given an irreducible discrete-time Markov chain on a finite state space, we consider the largest expected hitting time $T(\alpha)$ of a set of stationary measure at least α for $\alpha \in (0, 1)$. We obtain tight inequalities among the values of $T(\alpha)$ for different choices of α . One consequence is that $T(\alpha) \leq T(1/2)/\alpha$ for all $\alpha < 1/2$. As a corollary we have that, if the chain is lazy in a certain sense as well as reversible, then $T(1/2)$ is equivalent to the chain's mixing time, answering a question of Peres. We furthermore demonstrate that the inequalities we establish give an almost everywhere pointwise limiting characterisation of possible hitting time functions $T(\alpha)$ over the domain $\alpha \in (0, 1/2]$.

1. INTRODUCTION

Hitting times are a classical topic in the theory of finite Markov chains, with connections to mixing times, cover times and electrical network representations [5, 6]. In this paper, we consider a natural family of extremal problems for maximum expected hitting times. In contrast to most earlier work on hitting times that considered the maximum expected hitting times of individual states, we focus on hitting *sets* of states of at least a given stationary measure. Informally, we are interested in the following basic question: how much more difficult is it to hit a smaller set than a larger one? (We note that other, quite different extremal problems about hitting times have been considered, e.g. [3].)

Following the notation of Levin, Peres and Wilmer [5], we let a sequence of random variables $X = (X_t)_{t=0}^{\infty}$ denote an irreducible Markov chain with finite state space Ω , transition matrix P , and stationary distribution π . We denote by μ_0 some initial distribution of the chain and by \mathbb{P}_{μ_0} the corresponding law. In the case that $\mu_0 = x$ almost surely, for some $x \in \Omega$, we write \mathbb{P}_x for the corresponding law.

2010 *Mathematics Subject Classification*. Primary 60J10.

Key words and phrases. Markov chains, hitting times.

SG is supported by CNPq Proc. 500016/2010-2.

This work was begun while RJK was at Durham University, supported by EPSRC grant EP/G066604/1. He is currently supported by a NWO Veni grant.

RIO is supported by a Bolsa de Produtividade em Pesquisa and a Universal grant from CNPq, Brazil.

This work was begun while VP was at Durham University, supported by EPSRC grant EP/G066604/1. He is currently supported by EPSRC grant EP/J008087/1.

Given a subset $A \subseteq \Omega$, the *hitting time* of A is the random variable τ_A defined as follows:

$$\tau_A \equiv \min\{t : X_t \in A\}.$$

We shall take particular interest in the maximum expected hitting times of sets of at least a given size. For $\alpha \in (0, 1)$ we define $T(\alpha) = T^P(\alpha)$ as follows:

$$T(\alpha) \equiv \max\{\mathbb{E}_x[\tau_A] : x \in \Omega, A \subseteq \Omega, \pi(A) \geq \alpha\}.$$

In other words, $T(\alpha) = T^P(\alpha)$ is the maximum, over all starting states $X_0 = x \in \Omega$ and all sets $A \subseteq \Omega$ of stationary measure at least α , of the expected hitting time of A from x .

1.1. The extremal ratio problem. Note the obvious fact that, given $0 < \alpha < \beta < 1$, $T(\alpha)$ is lower bounded by $T(\beta)$ always. Informally in other words, it is more difficult to hit smaller subsets of the state space. A natural problem then is to determine how much more difficult this is, i.e. how large the ratio between $T(\alpha)$ and $T(\beta)$ can become. We dub this the *extremal ratio problem*.

Problem 1.1. *Given $0 < \alpha < \beta < 1$, what is the largest possible value of $T(\alpha)/T(\beta)$ over all irreducible finite Markov chains (on at least two states)?*

A first result on this problem was noted by the third author [7, Corollary 1.7].

Theorem 1.2. *Fix $0 < \alpha < \beta < 1/2$. There exists a constant $C_\beta > 0$ such that the following holds. For any irreducible finite Markov chain,*

$$T(\alpha) \leq C_\beta \cdot \frac{T(\beta)}{\alpha}.$$

This can be shown via Cèsaro mixing time, specifically as a consequence of an equivalence between $T(\beta)$ for $\beta \in (0, 1/2)$ and Cèsaro mixing time for any irreducible chain. This equivalence, which was recently proved independently by the third author [7] and by Peres and Sousi [9], we will discuss in more detail in Subsection 1.3.

In this paper, we improve upon the above result significantly, without recourse to any results on mixing time. Our first main result implies that the optimal constant in Theorem 1.2 is $C_\beta = 1$ and that moreover we can include the case $\beta = 1/2$.

Theorem 1.3. *Fix $0 < \alpha < \beta \leq 1/2$. For any irreducible finite Markov chain,*

$$(\star) \quad T(\alpha) \leq T(\beta) + \left(\frac{1}{\alpha} - 1\right) \cdot T(1 - \beta) \leq \frac{T(\beta)}{\alpha}.$$

This bound on $T(\alpha)$ is tight: for any $0 < \alpha < \beta \leq 1/2$, there exists an irreducible finite Markov chain for which the three terms in (\star) are all equal. Furthermore, $\beta = 1/2$ represents a boundary case for Theorem 1.3: for each $\beta > 1/2$, there is a class of irreducible finite Markov chains such that $T(\alpha)/T(\beta)$ is arbitrarily large. Thus we have completely settled the extremal ratio problem.

As an application of Theorem 1.3, we show in Subsection 1.3 how mixing time is equivalent to $T(1/2)$ for any irreducible chain, under the added restriction that the chain is lazy in a certain sense as well as reversible; this resolves a problem posed by Peres [4].

Our strategy for proving Theorem 1.3 relies on a simple, but useful proposition, which can be deduced from the ergodic properties of irreducible finite Markov chains. We require the following definitions. Given two sets $A, B \subseteq \Omega$, we define

$$d^+(A, B) \equiv \max_{x \in A} \mathbb{E}_x [\tau_B] \quad \text{and} \quad d^-(A, B) \equiv \min_{x \in A} \mathbb{E}_x [\tau_B].$$

Proposition 1.4. *Given an irreducible Markov chain with finite state space Ω and stationary distribution π , let $A, C \subseteq \Omega$. Then*

$$\pi(A) \leq \frac{d^+(A, C)}{d^+(A, C) + d^-(C, A)}.$$

Both Theorem 1.3 and Proposition 1.4 are proved in Section 2. The examples mentioned after the statement of Theorem 1.3 are presented in Section 3. We remark that we have replaced our original proof of Proposition 1.4 by a shorter and more elegant argument of Peres and Sousi [10].

1.2. The shape problem. In consideration of Theorem 1.3, it is natural to wonder what form the ratio $T(\alpha)/T(\beta)$ may possibly take. The second problem we treat is what we call the *shape problem*.

Problem 1.5. *What is the minimal set of constraints on the possible “shape” of the function $T(\alpha)$ over the domain $\alpha \in (0, 1/2]$ over irreducible finite Markov chains (on at least two states)?*

We show that, in the appropriate limit, the constraints imposed by (\star) in Theorem 1.3 are the only non-trivial constraints on $T(\alpha)$ over the domain $\alpha \in (0, 1/2]$. (The trivial constraint is that T must be a decreasing function.)

We now make this statement rigorous. Let \mathcal{F} denote the set of decreasing functions $f : (0, 1/2] \rightarrow \mathbb{R}$ given by $f(\alpha) = T(\alpha)/T(1/2)$ for some irreducible finite Markov chain (on at least two states). We also consider limits of such functions. Let $\overline{\mathcal{F}}$ denote the set of decreasing functions $f : (0, 1/2] \rightarrow \mathbb{R}$ each of which may be obtained as the almost everywhere (a.e.) pointwise limit of functions in \mathcal{F} . Our second main result is as follows.

Theorem 1.6. *Let $f : (0, 1/2] \rightarrow \mathbb{R}$ be a decreasing function. Then $f \in \overline{\mathcal{F}}$ if and only if $f(1/2) = 1$ and*

$$f(\alpha) \leq \frac{1}{\alpha} \quad \text{for all } \alpha \in (0, 1/2).$$

We prove this by way of a class of chains we call L -shaped Markov chains, for which the hitting time functions $T(\alpha)$ can be straightforwardly determined. We show Theorem 1.6 in Section 3.

As it turns out, the constraints given by (\star) for $0 < \alpha < \beta \leq 1/2$ are not the only non-trivial constraints on $T(\alpha)$ over the larger domain $\alpha \in (0, 1)$. We demonstrate this in Section 4. The shape problem over that larger domain remains an interesting open problem.

1.3. The connection to mixing times. To put our results into wider context, we now describe the relationship between Theorem 1.3 and mixing times. Recall that the (standard) *mixing time* of a chain with state space Ω , transition matrix P , and stationary distribution π is defined as

$$t_{\text{mix}}^P \equiv \min \left\{ t \in \mathbb{N} : \forall x \in \Omega, \forall A \subset \Omega, |P^t(x, A) - \pi(A)| \leq \frac{1}{4} \right\}.$$

This parameter has various connections to the analysis of MCMC algorithms, to phase transitions in statistical mechanics, and to other pure and applied problems [5]. Aldous [1] showed that it is also related to other parameters of the chain, including the following hitting time parameter:

$$t_{\text{prod}}^P \equiv \max\{\pi(A)\mathbb{E}_x[\tau_A] : x \in \Omega, \emptyset \neq A \subset \Omega\}.$$

Theorem 1.7. *There exists a universal constant $C > 0$ such that the following holds. Consider a reversible, irreducible finite Markov chain with transition matrix P that is lazy in the sense that $P_{xx} \geq 1/2$ for all x in the state space. Then*

$$\frac{t_{\text{mix}}^P}{C} \leq t_{\text{prod}}^P \leq C t_{\text{mix}}^P.$$

We remark that Aldous proved Theorem 1.7 in continuous time, but there are standard methods to transfer his result to discrete time (cf. [5, Theorem 20.3]).

Aldous's theorem is typically summed up by saying that t_{mix}^P and t_{prod}^P are “equivalent up to universal constants”, or simply “equivalent”. A similar equivalence was proved for all irreducible finite Markov chains (not necessarily lazy or reversible), with t_{mix}^P replaced by *Cèsaro mixing time* [2]:

$$t_{\text{Ces}}^P \equiv \min \left\{ t \in \mathbb{N} : \forall x \in \Omega, \forall A \subset \Omega, \left| \frac{1}{t} \sum_{s=0}^{t-1} P^s(x, A) - \pi(A) \right| \leq \frac{1}{4} \right\}.$$

A drawback of Theorem 1.7 and its Cèsaro mixing version is that it might seem that the mixing time depends on the hitting times of arbitrarily small sets. On the contrary, it transpires that the maximum hitting times of only sets that are large enough is also equivalent to t_{mix}^P and t_{Ces}^P (in the analogous senses). The following was proved independently by Peres and Sousi [9] and by the third author [7].

Theorem 1.8. *For each $\alpha \in (0, 1/2)$, there exists a constant $c(\alpha) > 0$ such that the following holds. Consider a reversible, irreducible finite Markov chain with transition matrix P that is lazy in the sense that $P_{xx} \geq 1/2$ for all x in the state space. Then*

$$\frac{t_{\text{mix}}^P}{c(\alpha)} \leq T^P(\alpha) \leq c(\alpha) t_{\text{mix}}^P.$$

Moreover, for any irreducible finite Markov chain (not necessarily reversible or lazy),

$$\frac{t_{\text{Ces}}^P}{c(\alpha)} \leq T^P(\alpha) \leq c(\alpha) t_{\text{Ces}}^P.$$

Note that, together with the Cèsaro mixing time form of Theorem 1.7, Theorem 1.2 now follows.

There is no analogue of Theorem 1.8 if one allows $\alpha > 1/2$: a simple counterexample is given by a random walk on a graph consisting of two large cliques connected by a single edge [8]. Until now, it was not known whether $T^P(1/2)$ is also equivalent to t_{mix}^P and t_{Ces}^P . We prove here that this is the case, answering a question of Peres [4].

Theorem 1.9. *There exists a universal constant $c > 0$ such that the following holds. Consider a reversible, irreducible finite Markov chain with transition matrix P that is lazy in the sense that $P_{xx} \geq 1/2$ for all x in the state space. Then*

$$\frac{t_{\text{mix}}^P}{c} \leq T^P(1/2) \leq c t_{\text{mix}}^P.$$




FIGURE 1. An illustration of the situation in Theorem 1.3.

Moreover, for any irreducible finite Markov chain (not necessarily reversible or lazy),

$$\frac{t_{\text{Ces}}^P}{c} \leq T^P(1/2) \leq c t_{\text{Ces}}^P.$$

Proof. By Theorem 1.7 and its Cèsaro mixing time version, it suffices to show that t_{prod}^P is equivalent to $T^P(1/2)$. But this is simple: on the one hand,

$$\frac{T^P(1/2)}{2} \leq \max\{\pi(A)\mathbb{E}_x[\tau_A] : x \in \Omega, A \subset \Omega, \pi(A) \geq 1/2\} \leq t_{\text{prod}}^P,$$

whereas, on the other hand, Theorem 1.3 implies that

$$\pi(A)\mathbb{E}_x[\tau_A] \leq \pi(A)T^P(\pi(A)) \leq T^P(1/2)$$

if $\pi(A) \leq 1/2$, and the fact that $T^P(\cdot)$ is monotone decreasing implies the above inequality also holds if $\pi(A) > 1/2$. \square

1.4. Organization. The remainder of the article is organised as follows. In Section 2, we prove Theorem 1.3. In Section 3, we show Theorem 1.3 is tight by presenting some two- and three-state Markov chains. We also prove Theorem 1.6 in Section 3. Finally, in Section 4 we consider the behaviour of $T(\alpha)$ over the larger domain $\alpha \in (0, 1)$ and make some concluding remarks.

2. PROOFS FOR THEOREM 1.3

We begin by showing that Theorem 1.3 is an easy consequence of Proposition 1.4.

Proof of Theorem 1.3. Consider an irreducible Markov chain with finite state space Ω and stationary distribution π . Fix a state $x \in \Omega$ and a set $A \subseteq \Omega$ with $\pi(A) \geq \alpha$. We prove that

$$\mathbb{E}_x[\tau_A] \leq T(\beta) + \left(\frac{1}{\alpha} - 1\right) \cdot T(1 - \beta).$$

Since x and A are arbitrary, this will suffice to prove the theorem.

Define the set $C = C_A^\beta$ as follows:

$$C \equiv \left\{ y \in \Omega : \mathbb{E}_y(\tau_A) > \left(\frac{1}{\alpha} - 1\right) \cdot T(1 - \beta) \right\}.$$

We claim that $\pi(C) < 1 - \beta$. Indeed, if, on the contrary, $\pi(C)$ were at least $1 - \beta$, then it would follow that $d^+(A, C) \leq T(1 - \beta)$ while $d^-(C, A) > (\alpha^{-1} - 1)T(1 - \beta)$. This would imply, by Proposition 1.4, that $\pi(A) < \alpha$, a contradiction. Thus, letting $B \equiv \Omega \setminus C$, we have established that $\pi(B) > \beta$. Our route from x to A is now clear — proceed from x to B and then on from B to A . See Figure 1. That is, using the Markovian property of the chain, the expected hitting time of A from x may be bounded by

$$\mathbb{E}_x[\tau_A] \leq \mathbb{E}_x[\tau_B] + d^+(B, A).$$

Combining the bound $\mathbb{E}_x [\tau_B] \leq T(\beta)$ (since $\pi(B) \geq \beta$) with the bound $d^+(B, A) \leq (\alpha^{-1} - 1) \cdot T(1 - \beta)$ (since B is the complement of C), we obtain

$$\mathbb{E}_x [\tau_A] \leq T(\beta) + \left(\frac{1}{\alpha} - 1\right) \cdot T(1 - \beta),$$

as required. \square

All that remains is to prove Proposition 1.4. As remarked in the introduction, we have replaced our original proof by a shorter and more elegant argument suggested by Peres and Sousi [10]. Our original proof, which may be obtained at <http://arxiv.org/abs/1209.0039v1>, relied on the ergodic theorem for irreducible Markov chains combined with a martingale concentration inequality.

Proof of Proposition 1.4 [10]. Denote the Markov chain by X . Our approach is to define a distribution μ on A and a distribution ν on C such that

$$(2.1) \quad \pi(A)\mathbb{E}_\nu [\tau_A] \leq (1 - \pi(A))\mathbb{E}_\mu [\tau_C].$$

Doing so will complete a proof of the proposition. Indeed, re-arranging inequality (2.1), we obtain

$$\pi(A) \leq \frac{\mathbb{E}_\mu [\tau_C]}{\mathbb{E}_\mu [\tau_C] + \mathbb{E}_\nu [\tau_A]} \leq \frac{d^+(A, C)}{d^+(A, C) + d^-(C, A)},$$

as required.

We now define the distributions μ and ν to satisfy inequality (2.1). Consider an auxiliary Markov chain on A defined by the following transitions: for each $x, y \in A$, let Q_{xy} be the probability that, started from x , the first state of A hit by X after time τ_C is y (i.e. that y is the first state of A hit after the original chain has reached C from x). Let μ denote a stationary distribution of this new chain, and let ν be the hitting distribution on C when the original chain is started from μ , i.e. $\nu(y) = \mathbb{P}_\mu (X_{\tau_C} = y)$ for each $y \in C$.

It remains to prove that (2.1) holds for this choice of μ and ν . First observe that, started from the distribution μ , the expected time the chain X spends in A before it reaches C and returns to A is given by $\mathbb{E}_\mu [\tau] \pi(A)$, where τ denotes the number of steps in such a cycle (from A to C then back to A). This observation is not difficult to verify, but we have included a proof below in Lemma A.1 of the appendix. Next, since all visits to A occur before the chain reaches C , we have that $\mathbb{E}_\mu [\tau] \pi(A) \leq \mathbb{E}_\mu [\tau_C]$. Finally, inequality (2.1) follows since $\mathbb{E}_\mu [\tau] = \mathbb{E}_\mu [\tau_C] + \mathbb{E}_\nu [\tau_A]$. \square

3. EXAMPLES AND A PROOF OF THEOREM 1.6

This section is devoted to exhibiting classes of Markov chains which demonstrate that Theorem 1.3 is tight, in a few different senses.

We first show that equality in (\star) is attained. For each $0 < \alpha < \beta \leq 1/2$ we exhibit an irreducible three-state chain with $T(\alpha) = T(\beta)/\alpha$ and hence $T(\alpha) = T(\beta) + (\alpha^{-1} - 1)T(\beta) \geq T(\beta) + (\alpha^{-1} - 1)T(1 - \beta)$, as required. Consider the three-state chain with transition matrix

$$\begin{pmatrix} 0 & 1 & 0 \\ \frac{\varepsilon}{(1-\alpha-\varepsilon)} & 1 - \frac{\alpha+\varepsilon}{(1-\alpha-\varepsilon)} & \frac{\alpha}{(1-\alpha-\varepsilon)} \\ 0 & 1 & 0 \end{pmatrix},$$

where $0 < \varepsilon < \beta - \alpha$. We note immediately that $(\varepsilon, 1 - \alpha - \varepsilon, \alpha)$ is the stationary distribution of the chain. It can be easily checked that $T(\beta) = 1$ and $T(\alpha) = 1/\alpha$.

We next show that the condition $\beta \leq 1/2$ in Theorem 1.3 is necessary by writing down an irreducible finite chain with $T(\beta) = 0$ and $T(\alpha)$ arbitrarily large when $\beta > 1/2$. Supposing $\beta > 1/2$, let N be an arbitrarily large number and let γ be such that $\max\{\alpha, 1/2\} < \gamma < \beta$. Consider the two-state Markov chain with transition matrix

$$\begin{pmatrix} 1 - \frac{1}{\gamma N} & \frac{1}{\gamma N} \\ \frac{1}{(1-\gamma)N} & 1 - \frac{1}{(1-\gamma)N} \end{pmatrix}.$$

The stationary distribution of the chain is $(\gamma, 1 - \gamma)$. It is an exercise to verify that $T(\beta) = 0$ and $T(\alpha) \geq (1 - \gamma)N$, as desired.

We now turn to the proof of Theorem 1.6. We must prove that each decreasing function $f : (0, 1/2] \rightarrow \mathbb{R}$ satisfying

$$f(\alpha) \leq \frac{1}{\alpha} \quad \text{for all } \alpha \in (0, 1/2)$$

may be obtained as the a.e. pointwise limit of a sequence of functions f_1, f_2, \dots in \mathcal{F} (i.e. functions f_i such that $f_i(\alpha) = T^{P_i}(\alpha)/T^{P_i}(1/2)$ for some irreducible finite Markov chain with transition matrix P_i). We first prove this for a certain class of step functions. Then we consider general functions as limits of these step functions in order to obtain the theorem.

The class of decreasing step functions $f : (0, 1/2] \rightarrow \mathbb{R}$ we consider are those that may be written in the form

$$f(\alpha) = 1 + \sum_{i=1}^k \lambda_i \cdot \mathbf{1}_{\alpha \leq \alpha_i},$$

where the λ_i and α_i are positive reals satisfying

$$(3.1) \quad \sum_{j=1}^i \lambda_j \leq \alpha_i^{-1} - 1 \quad \text{for each } i \in \{1, \dots, k\},$$

and $0 < \alpha_k < \dots < \alpha_1 < 1/2$. We call such a step function *hittable*. We note that if f is a hittable step function then $f(1/2) = 1$ and $f(\alpha) \leq 1/\alpha$ for all $\alpha \in (0, 1/2)$.


Given a hittable step function $f(\alpha) = 1 + \sum_{i=1}^k \lambda_i \cdot \mathbf{1}_{\alpha \leq \alpha_i}$, we define the ε -error set for f to be the set

$$\text{Err}_f(\varepsilon) \equiv \bigcup_{i=0}^k [\alpha_i, \alpha_i + \varepsilon],$$

where we interpret $\alpha_0 = 0$.

Lemma 3.1. *Let $f : (0, 1/2] \rightarrow \mathbb{R}$ be a hittable step function and $\varepsilon > 0$. Then there exists an irreducible finite Markov chain such that $f(\alpha) = T(\alpha)/T(1/2)$ for all $\alpha \in (0, 1/2] \setminus \text{Err}_f(\varepsilon)$.*

The examples of Markov chains we shall use in the proof of the lemma are all of the same type. An *L-shaped Markov chain* is a chain whose state space may be labelled $\Omega = \{v_{-1}, v_0, v_1, \dots, v_k\}$ in such a way that the transition matrix of the chain has non-zero entries only at $P_{i(i-1)}, P_{ii}, P_{(i-1)i}, P_{i0}$ for $i \in \{0, 1, \dots, k\}$. Note that v_0 is the only state that may be reached directly from a non-adjacent

FIGURE 2. A depiction of an L -shaped Markov chain.

state. Thus, with the exception of jumps to v_0 , all transitions are to a neighbour in the sequence $v_{-1}, v_0, v_1, \dots, v_k$. See Figure 2. In proving Lemma 3.1, we need only consider L -shaped chains. Indeed, it is because the hitting times of such Markov chains are relatively easy to determine that they are suitable for our purposes. The following lemma, though somewhat specialised, is exactly what we shall require in our proof of Lemma 3.1.

Lemma 3.2. *Suppose we are given an L -shaped Markov chain on state space $\Omega = \{v_{-1}, v_0, v_1, \dots, v_k\}$ with the property that $\mathbb{E}_{v_j} [\tau_{v_0}]$ is maximised at $j = -1$. If $i \in \{0, \dots, k\}$ and $\alpha \in (0, 1)$ satisfy*

$$(3.2) \quad \pi(\{v_{i+1}, \dots, v_k\}) + \pi(v_{-1}) < \alpha \leq \pi(\{v_i, \dots, v_k\}),$$

then

$$T(\alpha) = \mathbb{E}_{v_{-1}} [\tau_{\{v_i, \dots, v_k\}}] = \mathbb{E}_{v_{-1}} [\tau_{v_i}].$$

Proof. The second equality is obvious since, starting from v_{-1} , the chain first arrives in the set $\{v_i, \dots, v_k\}$ at v_i . It is also immediate that $T(\alpha) \geq \mathbb{E}_{v_{-1}} [\tau_{\{v_i, \dots, v_k\}}]$, by the definition of $T(\alpha)$ and the assumption that $\pi(\{v_i, \dots, v_k\}) \geq \alpha$. Thus all that remains is to prove, for any state v_j and set A with $\pi(A) \geq \alpha$, that $\mathbb{E}_{v_j} [\tau_A] \leq \mathbb{E}_{v_{-1}} [\tau_{\{v_i, \dots, v_k\}}]$.

Fix $j \in \{-1, \dots, k\}$ and a set A with $\pi(A) \geq \alpha$. Let i' be the minimal non-negative integer for which $v_{i'} \in A$. The condition on α implies that $i' \leq i$. Now (using the property that $\mathbb{E}_{v_j} [\tau_{v_0}]$ is maximised at $j = -1$, and the fact that $i' \leq i$) we have that

$$\mathbb{E}_{v_j} [\tau_A] \leq \mathbb{E}_{v_j} [\tau_{v_0}] + \mathbb{E}_{v_0} [\tau_{v_{i'}}] \leq \mathbb{E}_{v_{-1}} [\tau_{v_0}] + \mathbb{E}_{v_0} [\tau_{v_i}].$$

Since any path from v_{-1} to v_i necessarily passes through v_0 , the final expression is equal to $\mathbb{E}_{v_{-1}} [\tau_{\{v_i, \dots, v_k\}}]$, completing the proof. \square

The intuition of the above lemma (at least for our intended application) is that if v_{-1} has a very small measure (ε say) then for almost all values of α (except on a set of measure at most $k\varepsilon$) we know how to express $T(\alpha)$ directly as a hitting time. This is central to our proof of Lemma 3.1.

Proof of Lemma 3.1. We shall prove the following assertion: for every hittable step function $f(\alpha) = 1 + \sum_{i=1}^k \lambda_i \cdot \mathbf{1}_{\alpha \leq \alpha_i}$, every $0 < \varepsilon < 1/2 - \alpha_1$ and every sufficiently large natural number N , there exists an L -shaped Markov chain with transition matrix P , state space $\Omega = \{v_{-1}, v_0, v_1, \dots, v_k\}$ and stationary measure π satisfying

- (i) $\pi(v_{-1}) = \varepsilon$, $\pi(v_0) = 1 - \alpha_1 - \varepsilon$, and $\pi(\{v_i, \dots, v_k\}) = \alpha_i$ for each $i \in \{1, \dots, k\}$,
- (ii) $\mathbb{E}_{v_i} [\tau_{v_0}] \leq N$ for each $i \in \{-1, 0, 1, \dots, k\}$ with equality if $i = -1$, and

(iii) $\mathbb{E}_{v_{i-1}}[\tau_{v_i}] = \lambda_i N$ for each $i \in \{1, \dots, k\}$.

From this assertion Lemma 3.1 easily follows. Indeed, since $\pi(v_0) = 1 - \alpha_1 - \varepsilon > 1/2$ we have that $T(1/2)$ is precisely the maximum expected hitting time of v_0 , and it follows immediately from (ii) that $T(1/2) = N$. Given $\alpha \in (0, 1/2] \setminus \text{Err}_f(\varepsilon)$, we shall determine $T(\alpha)$ using Lemma 3.2 and condition (iii). In order to apply Lemma 3.2, first notice that condition (ii) ensures that $\mathbb{E}_{v_j}[\tau_{v_0}]$ is maximised at $j = -1$. Let $i \in \{1, \dots, k\}$ be smallest such that $\alpha \leq \alpha_i$. Using (i) and the fact that $\alpha \in (0, 1/2] \setminus \text{Err}_f(\varepsilon)$, it is straightforward to verify that (3.2) holds in the statement of Lemma 3.2. Thus, applying Lemma 3.2 and using condition (iii), we have

$$T(\alpha) = \mathbb{E}_{v_{-1}}[\tau_{v_i}] = \mathbb{E}_{v_{-1}}[\tau_{v_0}] + \sum_{j=1}^i \mathbb{E}_{v_{j-1}}[\tau_{v_j}] = \left(1 + \sum_{j=1}^i \lambda_j\right) N = f(\alpha)T(1/2),$$

as required.

We now prove the above assertion by stating explicitly the entries of the transition matrix P . First, we set

$$\begin{aligned} P_{-10} &= \frac{1}{N}, & P_{0-1} &= \frac{\varepsilon}{(1 - \alpha_1 - \varepsilon)N}, & P_{-1-1} &= 1 - P_{-10} \\ P_{01} &= \frac{1 - \alpha_1}{(1 - \alpha_1 - \varepsilon)\lambda_1 N}, & P_{10} &= \frac{1 - \alpha_1 - \lambda_1 \alpha_2}{(\alpha_1 - \alpha_2)\lambda_1 N} \quad \text{and} \quad P_{00} = 1 - P_{0-1} - P_{01}. \end{aligned}$$

Next, for each $i \in \{2, \dots, k\}$, we set

$$\begin{aligned} P_{(i-1)i} &= \frac{1 - \alpha_i(1 + \sum_{j=1}^{i-1} \lambda_j)}{(\alpha_{i-1} - \alpha_i)\lambda_i N}, & P_{i(i-1)} &= \frac{1 - \alpha_i(1 + \sum_{j=1}^i \lambda_j)}{(\alpha_i - \alpha_{i+1})\lambda_i N}, \\ P_{i0} &= \frac{1}{N} \quad \text{and} \quad P_{ii} = 1 - P_{i0} - P_{i(i-1)} - P_{i(i+1)}. \end{aligned}$$

Last, we set $P_{11} = 1 - P_{10} - P_{12}$. It is routine to verify that each entry in the transition matrix P of our Markov chain is in $[0, 1]$ using (3.1), $0 < \varepsilon < 1/2 - \alpha_1$, $0 < \alpha_k < \dots < \alpha_2 < \alpha_1$, and a large enough choice of N .

Some straightforward calculations confirm that the resulting stationary distribution π satisfies condition (i) above. Condition (ii) follows easily from checking that $P_{i0} \geq 1/N$ (so that $\mathbb{E}_{v_i}[\tau_{v_0}] \leq N$) for all i and that $\mathbb{E}_{v_{-1}}[\tau_{v_0}] = N$. To verify condition (iii) for each $i \in \{1, \dots, k\}$, we compute the expected hitting time from v_{i-1} to v_i by considering the chain started at v_{i-1} and conditioning on the first step. We use induction on i . For the base case ($i = 1$), we have that

$$\begin{aligned} \mathbb{E}_{v_0}[\tau_{v_1}] &= 1 + P_{00}\mathbb{E}_{v_0}[\tau_{v_1}] + P_{0-1}\mathbb{E}_{v_{-1}}[\tau_{v_1}] \\ &= 1 + P_{00}\mathbb{E}_{v_0}[\tau_{v_1}] + P_{0-1}(N + \mathbb{E}_{v_0}[\tau_{v_1}]), \end{aligned}$$

which implies (after substitution and rearrangement) that $\mathbb{E}_{v_0}[\tau_{v_1}] = \lambda_1 N$. Next,

$$\begin{aligned} \mathbb{E}_{v_1}[\tau_{v_2}] &= 1 + P_{11}\mathbb{E}_{v_1}[\tau_{v_2}] + P_{10}\mathbb{E}_{v_0}[\tau_{v_2}] \\ &= 1 + P_{11}\mathbb{E}_{v_1}[\tau_{v_2}] + P_{10}(\lambda_1 N + \mathbb{E}_{v_1}[\tau_{v_2}]), \end{aligned}$$

which implies that $\mathbb{E}_{v_1}[\tau_{v_2}] = \lambda_2 N$. Finally, for $i \in \{3, \dots, k\}$, we have

$$\begin{aligned} \mathbb{E}_{v_{i-1}}[\tau_{v_i}] &= 1 + P_{(i-1)(i-2)} \mathbb{E}_{v_{i-2}}[\tau_{v_i}] + P_{(i-1)(i-1)} \mathbb{E}_{v_{i-1}}[\tau_{v_i}] + P_{(i-1)0} \mathbb{E}_{v_0}[\tau_{v_i}] \\ &= 1 + P_{(i-1)(i-2)} (\lambda_{i-1} N + \mathbb{E}_{v_{i-1}}[\tau_{v_i}]) + P_{(i-1)(i-1)} \mathbb{E}_{v_{i-1}}[\tau_{v_i}] \\ &\quad + P_{(i-1)0} \left(\sum_{j=1}^{i-1} \lambda_j N + \mathbb{E}_{v_{i-1}}[\tau_{v_i}] \right), \end{aligned}$$

where the second equality uses the inductive assumption that $\mathbb{E}_{v_{j-1}}[\tau_{v_j}] = \lambda_j N$ for $j \in \{1, \dots, i-1\}$. This implies that $\mathbb{E}_{v_{i-1}}[\tau_{v_i}] = \lambda_i N$, as desired. \square

It is now straightforward to deduce Theorem 1.6.

Proof of Theorem 1.6. The only if part is an immediate consequence of Theorem 1.3. Now, fix a decreasing function $f : (0, 1/2] \rightarrow \mathbb{R}$ with $f(1/2) = 1$ that satisfies $f(\alpha) \leq \alpha^{-1}$ for all $\alpha \in (0, 1/2)$. Denote by $D = D(f) \subseteq (0, 1/2]$ the set of discontinuity points of f . Since f is decreasing, the set D is countable by Froda's theorem¹. For each positive integer n , define the function $f_n : (0, 1/2] \rightarrow \mathbb{R}$ by

$$f_n(x) = f(\lceil 2^n x \rceil 2^{-n}).$$

One easily notes that $f_n(x) \rightarrow f(x)$ for all $x \in (0, 1/2] \setminus D$.

We observe that each f_n is a hittable step function, because it can be written

$$1 + \sum_{i=1}^{2^{n-1}-1} \lambda_i \mathbf{1}_{\alpha \leq \alpha_i},$$

where $\alpha_i = 1/2 - i2^{-n}$, and $\lambda_i = f(\alpha_i) - f(\alpha_{i-1})$. Condition (3.1) is easily seen to hold since

$$1 + \sum_{j=1}^i \lambda_j = 1 + f(\alpha_i) - f(\alpha_0) = f(\alpha_i) \leq \alpha_i^{-1}.$$

To prove the theorem we must find a sequence of functions $g_n \in \mathcal{F}$ such that $g_n(x) \rightarrow f(x)$ except on a set of measure zero. By Lemma 3.1 there exists for each n a function $g_n \in \mathcal{F}$ such that $g_n(x) = f_n(x)$ for all $x \in (0, 1/2] \setminus \text{Err}_{f_n}(2^{-2n})$, where

$$\text{Err}_{f_n}(2^{-2n}) = \bigcup_{i=0}^{2^{n-1}-1} \left[\frac{i}{2^n}, \frac{i}{2^n} + \frac{1}{2^{2n}} \right].$$

We now prove that $g_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$ for each $x \in (0, 1/2] \setminus (D \cup D')$, where D' denotes the set of points that lie in infinitely many intervals of $\text{Err}_{f_n}(2^{-2n})$. Since $D \cup D'$ has measure zero, this will complete the proof of the theorem.

To this end, fix $x \in (0, 1/2] \setminus (D \cup D')$. Since $x \notin D$, we have that $f_n(x) \rightarrow f(x)$ as $n \rightarrow \infty$. Furthermore, since $x \notin D'$, there exists n_0 such that

$$x \notin \bigcup_{n \geq n_0} \bigcup_{i=0}^{2^{n-1}-1} \left[\frac{i}{2^n}, \frac{i}{2^n} + \frac{1}{2^{2n}} \right],$$

and so $g_n(x) = f_n(x)$ for all $n \geq n_0$. Thus $\lim_{n \rightarrow \infty} g_n(x) = \lim_{n \rightarrow \infty} f_n(x) = f(x)$, completing the proof of the theorem. \square

¹See http://en.wikipedia.org/wiki/Froda's_theorem.

4. ONE FURTHER RESULT AND CONCLUDING REMARKS

For $0 < \alpha < \beta \leq 1/2$ we proved the tight inequality $T(\alpha) \leq T(\beta) + (\alpha^{-1} - 1)T(1 - \beta)$ relating hitting times of large enough sets in irreducible finite Markov chains. Furthermore, we demonstrated that this is the only non-trivial restriction on $T(\alpha)$ as a function over $\alpha \in (0, 1/2]$, in the sense made rigorous in Theorem 1.6.

The most obvious remaining question then is whether there are other non-trivial inequalities relating the values of $T(\alpha)$ for all $\alpha \in (0, 1)$. In one further result, we demonstrate that $T : (0, 1) \rightarrow \mathbb{R}$ is further constrained. However, determining the set of all inequalities that hold among the values of $T(\alpha)$ for all $\alpha \in (0, 1)$ and thereby giving a characterisation in the spirit of Theorem 1.6 of the possible behaviour of $T : (0, 1) \rightarrow \mathbb{R}$ remains an interesting open problem.

To demonstrate that $T : (0, 1) \rightarrow \mathbb{R}$ is further constrained it suffices to give a single example of such an additional restriction, which is as follows.

Proposition 4.1. *Given an irreducible finite Markov chain, suppose $T(0.01) = 99.9T(0.02)$. Then $T(0.99) \geq 0.1T(0.02)$.*

We note that this restriction is indeed outside of the class of restrictions imposed by Theorem 1.3. Writing T for $T(0.02)$, first one can check using Lemma 3.1 that there exist Markov chains satisfying the equality $T(0.01) = 99.9T$. Furthermore, assuming this equality, the application of Theorem 1.3 gives that $T(0.01) \leq T + 99T(0.98)$. Although this inequality demands that $T(0.98)$ be very close to T — specifically, $T(0.98) \in [(98.9/99)T, T]$ — there is no restriction on $T(0.99)$. Thus Proposition 4.1 does indeed represent an additional restriction. We require the following lemma.

Lemma 4.2. *Given an irreducible Markov chain with finite state space Ω , let $A, B, C \subseteq \Omega$ and T be a real number such that*

$$d^+(\Omega, B) \leq T, \quad d^+(\Omega, A \cup C) \leq T, \quad d^+(\Omega, A) \leq 99.9T \quad \text{and} \quad d^-(B, A) \geq 98.9T.$$

Then $d^+(B, C) < 14T$.

Proof. Let $y \in B$. Consider running the chain for $10T$ steps and denote by p_y the probability $\mathbb{P}_y(\tau_A \leq 10T)$. The assumptions on the hitting time of A imply that

$$98.9T \leq \mathbb{E}_y[\tau_A] \leq 10T + (1 - p_y)99.9T.$$

Thus $p_y < 0.111 < 1/8$. On the other hand, $\mathbb{P}_y(\tau_{A \cup C} \leq 10T) \geq 9/10$ by Markov's inequality, and so $\mathbb{P}_y(\tau_C \leq 10T) \geq 9/10 - 1/8 > 3/4$.

We may now bound $d^+(B, C)$ as follows. Note that, in the event that the chain does not hit C after $10T$ steps, the expected remaining time to hit C may be bounded by T (an upper bound on expected time to return to B) plus $d^+(B, C)$ (an upper bound on the expected time to hit C from an element of B). Thus

$$d^+(B, C) \leq 10T + \frac{1}{4}(T + d^+(B, C)).$$

It follows that $d^+(B, C) \leq 41T/3 < 14T$, as required. \square

We now prove Proposition 4.1.

Proof of Proposition 4.1. Let us write T for $T(0.02)$. Since $T(0.01) = 99.9T$ there exists a set $A \subseteq \Omega$ with $\pi(A) \geq 0.01$ and a state $x \in \Omega$ such that $\mathbb{E}_x[\tau_A] = 99.9T$.




FIGURE 3. An illustration of the situation in Proposition 4.1.

Define sets

$$B' \equiv \{y \in \Omega : \mathbb{E}_y [\tau_A] \leq 99T\} \quad \text{and} \quad B \equiv \{y \in \Omega : \mathbb{E}_y [\tau_A] \in [98.9T, 99T]\}.$$

Arguing as in the proof of Theorem 1.3, one obtains that $\pi(B') \geq 0.98$ — specifically, if this were not the case, then one would have $d^+(A, \Omega \setminus B') \leq T$ and $d^-(\Omega \setminus B', A) > 99T$, which contradicts the bound of $\pi(A) \geq 0.01$ using Proposition 1.4. We now claim that $\pi(B) \geq 0.96$. Indeed, if on the contrary $\pi(B' \setminus B)$ were greater than 0.02, then one would obtain $\mathbb{E}_x [\tau_A] < \mathbb{E}_x [\tau_{B' \setminus B}] + 98.9T \leq 99.9T$, a contradiction.

Now, define

$$C \equiv \{y \in \Omega : \mathbb{E}_y [\tau_A] \geq 99.8T\}.$$

See Figure 3. We claim that $\pi(C) \leq 0.01$. Indeed, if $\pi(C)$ were greater than 0.01, then the set $A \cup C$ would have stationary measure at least 0.02, so that $d^+(\Omega, A \cup C) \leq T$. And we would then obtain from Lemma 4.2 that $d^+(B, C) < 14T$. On the other hand, $d^-(C, B) \geq 0.8T$ (otherwise, $d^-(C, A) \leq d^-(C, B) + d^+(B, A) < 0.8T + 99T$, which contradicts the definition of C). And so, by Proposition 1.4, $\pi(B) < 14T/14.8T = 70/74 < 0.96$, a contradiction. Thus we have $\pi(\Omega \setminus C) \geq 0.99$ and the inequality $99.9T = \mathbb{E}_x [\tau_A] \leq \mathbb{E}_x [\tau_{\Omega \setminus C}] + 99.8T$ implies that $\mathbb{E}_x [\tau_{\Omega \setminus C}] \geq 0.1T$. Therefore $T(0.99) \geq 0.1T$, as required. \square

Acknowledgements. We are grateful to Yuval Peres and Perla Sousi for their kind permission to include their proof of Proposition 1.4. We also thank the anonymous referee for helpful comments and corrections.

REFERENCES

1. David J. Aldous, *Some inequalities for reversible Markov chains*, J. London Math. Soc. (2) **25** (1982), no. 3, 564–576. MR 657512 (83f:60098)
2. David J. Aldous, László Lovász, and Peter Winkler, *Mixing times for uniformly ergodic Markov chains*, Stochastic Process. Appl. **71** (1997), no. 2, 165–185. MR 1484158 (98i:60062)
3. Graham Brightwell and Peter Winkler, *Maximum hitting time for random walks on graphs*, Random Struct. Algorithms **1** (1990), no. 3, 263–276.
4. Navin Goyal, *Problems from the AIM Workshop on Algorithmic Convex Geometry*, Accessed 11 February 2012: <http://www.aimath.org/WWN/convexgeometry/convexgeometry.pdf>, 2007.
5. David A. Levin, Yuval Peres, and Elizabeth L. Wilmer, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2009, With a chapter by James G. Propp and David B. Wilson. MR 2466937 (2010c:60209)
6. László Lovász, *Random walks on graphs: a survey*, Combinatorics, Paul Erdős is Eighty, Vol. 2, Bolyai Soc. Math. Stud., vol. 2, János Bolyai Math. Soc., Budapest, 1996, pp. 353–397. MR 1395866
7. Roberto Imbuzeiro Oliveira, *Mixing and hitting times for finite markov chains*, Electron. J. Probab. **17** (2012), no. 70, 1–12.

8. Yuval Peres, Personal communication, 2012.
9. Yuval Peres and Perla Sousi, *Mixing times are hitting times of large sets*, <http://arxiv.org/abs/1108.0133>, 2011.
10. ———, Personal communication, 2012.

APPENDIX A. A LEMMA FOR THE PROOF OF PROPOSITION 1.4

For completeness, we include here a proof of an assertion used in the proof of Proposition 1.4. This is similar to the proof of Proposition 1.14 of [5].

Lemma A.1. *Suppose we are given an irreducible Markov chain X with finite state space Ω and stationary distribution π . Let μ be a distribution on Ω and τ be an a.s. positive stopping time such that $\mathbb{P}_\mu(X_\tau = x) = \mu(x)$ for any $x \in \Omega$. Then for any $S \subset \Omega$ the expected time X spends in S up to time τ starting from μ equals $\pi(S) \mathbb{E}_\mu[\tau]$.*

Proof. For each $x \in \Omega$, define $\tilde{\pi}(x)$ as the expected time X spends at x up to time τ when started from μ , i.e.

$$\tilde{\pi}(x) \equiv \mathbb{E}_\mu \left[\sum_{t=0}^{\infty} \mathbf{1}_{\{X_t=x, t < \tau\}} \right].$$

We shall prove that $(\tilde{\pi} P)(x) = \tilde{\pi}(x)$ for all $x \in \Omega$, which implies that $\tilde{\pi}$ is a multiple of the (unique) stationary distribution π . The lemma then follows from $\sum_{x \in \Omega} \tilde{\pi}(x) = \mathbb{E}_\mu[\tau]$, so that $\tilde{\pi}(x) = \pi(x) \mathbb{E}_\mu[\tau]$ for all $x \in \Omega$.

Let us now fix $x \in \Omega$ and compute $(\tilde{\pi} P)(x)$:

$$(\tilde{\pi} P)(x) = \sum_{y \in \Omega} \mathbb{E}_\mu \left[\sum_{t=0}^{\infty} \mathbf{1}_{\{X_t=y, t < \tau\}} \right] P_{yx} = \sum_{t=0}^{\infty} \sum_{y \in \Omega} \mathbb{P}_\mu(X_t = y, t < \tau) P_{yx},$$

since all terms are non-negative. By the Markov property, each term of the double sum equals $\mathbb{P}_\mu(X_t = y, X_{t+1} = x, t < \tau)$, and resolving the inner sum gives

$$(\tilde{\pi} P)(x) = \sum_{t=0}^{\infty} \mathbb{P}_\mu(X_{t+1} = x, t < \tau).$$

Now we split each term in the last summation into two parts as follows:

$$\mathbb{P}_\mu(X_{t+1} = x, t < \tau) = \mathbb{P}_\mu(X_{t+1} = x, t + 1 < \tau) + \mathbb{P}_\mu(X_{t+1} = x, t + 1 = \tau).$$

Summing the first part over t gives

$$\begin{aligned} \sum_{t=0}^{\infty} \mathbb{P}_\mu(X_{t+1} = x, t + 1 < \tau) &= \tilde{\pi}(x) - \mathbb{P}_\mu(X_0 = x, \tau > 0) \\ &= \tilde{\pi}(x) - \mathbb{P}_\mu(X_0 = x) = \tilde{\pi}(x) - \mu(x), \end{aligned}$$

where the second equality uses the fact that $\tau > 0$ a.s. Summing the second part,

$$\sum_{t=0}^{\infty} \mathbb{P}_\mu(X_{t+1} = x, t + 1 = \tau) = \mathbb{P}_\mu(X_\tau = x) = \mu(x),$$

where we have used $\tau > 0$ a.s. for the first equality and the assumption on μ for the second. It follows that $(\tilde{\pi} P)(x) = \tilde{\pi}(x) - \mu(x) + \mu(x) = \tilde{\pi}(x)$, as desired. \square

INSTITUTO NACIONAL DE MATEMÁTICA PURA E APLICADA (IMPA), RIO DE JANEIRO, BRAZIL
E-mail address: `sgriff@impa.br`

CENTRUM WISKUNDE & INFORMATICA, AMSTERDAM, NETHERLANDS
E-mail address: `ross.kang@gmail.com`

INSTITUTO NACIONAL DE MATEMÁTICA PURA E APLICADA (IMPA), RIO DE JANEIRO, BRAZIL
E-mail address: `rimfo@impa.br`

UNIVERSITY OF BIRMINGHAM, BIRMINGHAM, UNITED KINGDOM
E-mail address: `viresh.s.patel@gmail.com`