Monotonicity of Multiple-Cover Times under Uniform Speed-Up

Mark Holmes* Yuval Peres†

15 November 2026

Abstract

Let P be an irreducible transition matrix on a finite state space V, and for $0 < q \leq 1$ let $P_q = (1-q)I + qP$ be the lazy chain that stays put with probability $1-q$ at each step (and otherwise makes a transition according to P). Fix positive integers L_v, $v \in V$, and consider the first time at which every state v has been visited at least L_v times by the P_q-chain, where each time unit spent at a state, including holds, counts as a visit. Holding at a state is a cheap way to accumulate visits there, so it is not clear how this time depends on q. We prove that its expectation is strictly decreasing in q whenever $|V| \geq 2$. The same is true in continuous time: if all jump rates of a finite irreducible chain are multiplied by c, then the expected time until the occupation time of each state v reaches a given level a_v is strictly decreasing in c. Both results follow from an inequality for the original chain, which we prove using the exchangeability of excursions. Examples with two states show that the time for the lazy chain need not be stochastically monotone in q, and that the expectation can increase when the speed-up is not uniform in space or in time.

1 Introduction

Let V be a finite set and let P be an irreducible transition matrix on V. For $0 < q \leq 1$ consider the lazy version of P,

$$P_q = (1-q)I + qP.$$

A chain with transition matrix P_q takes a step according to P with probability q and stays where it is with probability $1-q$, independently at each time. Let $Y = (Y_t)_{t \geq 0}$ be such a chain and let

$$\ell_t(v) = \#\{0 \leq s \leq t : \ Y_s = v\}$$

be the number of visits to v by time t. Every time unit spent at v is counted here, so the starting state receives a visit at time 0, and a hold at v is another visit to v. Given a vector $\boldsymbol{L} = (L_v)_{v \in V}$ of positive integers, let

$$\tau_{\boldsymbol{L}}(q) = \min\{t \geq 0 : \ \ell_t(v) \geq L_v \ \text{ for all } v \in V\}$$

be the first time by which each state v has been visited at least L_v times. If $L_v = L$ for all v this is the L-cover time of Y, and $L = 1$ gives the usual cover time. We write $\mathbb{P}_x$ and $\mathbb{E}_x$ for probabilities and expectations when the chain starts at x, and $\mathbb{N} = \{1, 2, \dots\}$.

How does $\tau_{\boldsymbol{L}}(q)$ depend on q? For the cover time the answer is easy. A hold never leads to a new state, so the lazy chain is the P-chain run on a slower clock, and

$$\mathbb{E}_x[\tau_{\mathbf{1}}(q)] = \frac{1}{q}\,\mathbb{E}_x[\tau_{\mathbf{1}}(1)],$$

*School of Mathematics and Statistics, The University of Melbourne, Australia.
†Beijing Institute of Mathematical Sciences and Applications, Beijing, China.

where $\mathbf{1}$ is the vector with all entries equal to 1. When some L_v is larger than 1 the question is more delicate, because a hold at v is itself a visit to v. If v still needs visits, then holding there is the cheapest way to get one, while a step of P may take the chain on a long excursion before it returns to v. Decreasing q therefore slows the chain down but makes each visit of the underlying P-chain worth more, and it is not obvious which effect dominates. Our first result is that the slower chain always takes longer on average.

Theorem 1.1. *Let P be an irreducible transition matrix on a finite set V with $|V| \geq 2$, and let $\boldsymbol{L} \in \mathbb{N}^V$. Then for every $x \in V$ the function*

$$q \mapsto \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)]$$

is strictly decreasing on $(0, 1]$. The same holds for $\max_{x \in V} \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)]$.

For the graph consisting of a single edge, Staite [11] found an explicit formula for the expected L-cover time of the lazy walk and proved that it is monotone in q, for every L.

There is a continuous-time version of Theorem 1.1, in which the number of visits is replaced by the occupation time, and the amount of time required at each state can be an arbitrary positive number. Let Q be the generator of an irreducible continuous-time Markov chain on V. For $c > 0$ let $Z^{(c)}$ be the chain with generator cQ, in which all the jump rates are multiplied by c, and let

$$\ell_t^{(c)}(v) = \int_0^t \mathbb{1}_{\{Z_s^{(c)} = v\}}\, ds$$

be the time that it spends at v up to time t. For $\boldsymbol{a} = (a_v)_{v \in V} \in (0, \infty)^V$ define

$$\theta_{\boldsymbol{a}}(c) = \inf\{t \geq 0 :\ \ell_t^{(c)}(v) \geq a_v \ \text{ for all } v \in V\}.$$

Theorem 1.2. *Let Q be the generator of an irreducible Markov chain on a finite set V with $|V| \geq 2$, and let $\boldsymbol{a} \in (0, \infty)^V$. Then for every $x \in V$ the function $c \mapsto \mathbb{E}_x[\theta_{\boldsymbol{a}}(c)]$ is strictly decreasing on $(0, \infty)$. The same holds for $\max_{x \in V} \mathbb{E}_x[\theta_{\boldsymbol{a}}(c)]$.*

If V consists of a single state x, then $\tau_{\boldsymbol{L}}(q) = L_x - 1$ and $\theta_{\boldsymbol{a}}(c) = a_x$ do not depend on q and c. Note that reversibility is not assumed in either theorem.

The theorems cannot be strengthened in two natural directions. First, they are statements about expectations only. The time $\tau_{\boldsymbol{L}}(q)$ is in general not stochastically decreasing in q, so Theorem 1.1 cannot be proved by coupling the chains for two values of q in such a way that the faster one never finishes after the slower one; see Section 5.1. Second, it matters that the chain is sped up by the same factor at all states and at all times. In Sections 5.2–5.4 we give examples with two states in which the expected time goes up when the holding probability is decreased at some of the states, or at some of the times, and also when all the jump rates of a continuous-time chain are increased, but not by the same factor.

Outline of the proof. Write the lazy chain as $Y_t = X_{N_t}$, where X is a chain with transition matrix P and N_t is the number of P-steps taken by time t. Each visit of X to v lasts a geometric number of time units, and it follows that the number K_v of visits that X must pay to v before Y has visited v at least L_v times satisfies $K_v - 1 \sim \mathrm{Bin}(L_v - 1, q)$. Moreover, the variables K_v are independent of each other and of X. This leads to the identity

$$q\, \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] = \mathbb{E}[f_x(\boldsymbol{K})],$$

where $f_x(\boldsymbol{k})$ denotes the expected time for X to visit every state v at least k_v times (Proposition 2.1). The law of $\boldsymbol{K}$ depends on q in a simple way, and differentiating in q we find that

$$\frac{d}{dq}\, \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] = -\frac{1}{q^2}\, \mathbb{E}\Big[f_x(\boldsymbol{K}) - \sum_{v \in V} (K_v - 1)\, \Delta_v f_x(\boldsymbol{K})\Big],$$

where $\Delta_v f_x(\boldsymbol{k})$ is the expected time that is saved when the requirement k_v is lowered by one (the terms with $K_v = 1$ vanish). Theorem 1.1 thus reduces to the inequality

$$\sum_{v \in V} (k_v - 1)\, \Delta_v f_x(\boldsymbol{k}) \;\le\; f_x(\boldsymbol{k}),$$

which is Lemma 3.1. The reason for it is the following. Lowering the requirement at v makes a difference only if v is the state whose requirement is fulfilled last, and then it saves at most the last of the $k_v - 1$ excursions of X from v. These excursions are exchangeable, so the last one takes on average a fraction $1/(k_v - 1)$ of the time that all of them take together. In continuous time the same argument works, with Poisson variables in place of binomial ones.

Related work. Cover times of finite Markov chains have been studied for a long time; see Matthews [10] and [9, Chapter 11] for background. Winkler and Zuckerman [12] considered the time needed for a random walk on a graph to visit every vertex many times. They introduced the *blanket time*, the first time at which every vertex has been visited at least a constant fraction of the number of times predicted by the stationary distribution, and conjectured that its expectation is bounded by a constant multiple of the cover time. This was proved by Ding, Lee and Peres [5]. Chan, Ding and Li [3] introduced the k-cover time, which is the expected time to visit every state at least k times, and used it to bound the number of samples needed to learn or test a Markov chain from a single trajectory. Ding [4, Chapter 3] bounds the additional time needed to cover a graph a second time after the first cover.

Theorems 1.1 and 1.2 say that a uniform speed-up always helps. Other natural modifications of a chain can hurt: Barlow, Ding, Nachmias and Peres [2] showed that adding an edge to a graph can increase the cover time, though by a factor of at most 4. Cover times of walks whose transition rule varies with time were studied in [1, 7]; see also [6], where the holding probabilities of a random walk on $\mathbb{Z}$ are chosen by a controller.

Organization. Section 2 contains the representation of the lazy chain in terms of holding times. The key inequality is proved in Section 3, and the two theorems in Section 4. The examples are in Section 5.

2 Holding times

In this section and the next two, $X = (X_n)_{n \ge 0}$ is a Markov chain on V with irreducible transition matrix P, started at x, and we assume that $|V| \ge 2$. For $v \in V$ and $j \ge 1$ let $S_v(j)$ be the time of the j-th visit of X to v, where a visit at time 0 is counted. For $\boldsymbol{k} = (k_v)_{v \in V} \in \mathbb{N}^V$ define

$$\sigma(\boldsymbol{k}) = \max_{v \in V} S_v(k_v) \qquad \text{and} \qquad f_x(\boldsymbol{k}) = \mathbb{E}_x[\sigma(\boldsymbol{k})]. \tag{2.1}$$

Thus $\sigma(\boldsymbol{k})$ is the first time by which X has visited every state v at least k_v times, and $\sigma(\boldsymbol{L}) = \tau_{\boldsymbol{L}}(1)$ in the notation of the introduction. To see that f_x is finite, let $T_v^+ = \min\{n \ge 1 : X_n = v\}$ and $\kappa = \max_{u,v \in V} \mathbb{E}_u[T_v^+]$, which is finite because V is finite and P is irreducible. The strong Markov property gives $\mathbb{E}_x[S_v(j)] \le j\kappa$, and since $\sigma(\boldsymbol{k}) \le \sum_v S_v(k_v)$ we obtain

$$f_x(\boldsymbol{k}) \le \kappa \sum_{v \in V} k_v. \tag{2.2}$$

The next proposition relates the expectations in Theorems 1.1 and 1.2 to the function f_x. For the continuous-time statement we let $\gamma = \max_{v \in V} |Q(v,v)|$, which is positive, and take $P = I + Q/\gamma$. This is an irreducible transition matrix on V.

Proposition 2.1. (i) *Let $0 < q \le 1$ and $\boldsymbol{L} \in \mathbb{N}^V$. Let $\boldsymbol{K} = (K_v)_{v \in V}$ be a random vector with independent coordinates such that $K_v - 1$ has the binomial distribution with parameters $L_v - 1$ and q. Then*

$$q\, \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] = \mathbb{E}[f_x(\boldsymbol{K})]. \tag{2.3}$$

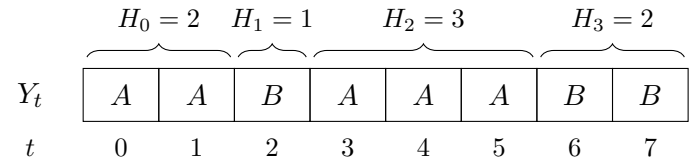


Figure 1: The lazy chain Y when X begins with A, B, A, B and the first four holding times are $2, 1, 3, 2$. If $L_A = 4$ and $L_B = 2$, then $K_A = K_B = 2$ and $\tau_{\boldsymbol{L}}(q) = 6$. At time 6 the chain X has made three steps, and indeed $\sigma(\boldsymbol{K}) = \max\{S_A(2), S_B(2)\} = 3$.

(ii) *Let $c > 0$ and $\boldsymbol{a} \in (0,\infty)^V$, and let $P = I + Q/\gamma$ as above. Let $\boldsymbol{K} = (K_v)_{v\in V}$ be a random vector with independent coordinates such that $K_v - 1$ has the Poisson distribution with mean $c\gamma a_v$. Then*

$$c\gamma\, \mathbb{E}_x[\theta_{\boldsymbol{a}}(c)] = \mathbb{E}[f_x(\boldsymbol{K})]. \tag{2.4}$$

By (2.2), the right-hand sides of (2.3) and (2.4) are finite. In the proof we will construct $\boldsymbol{K}$ on the same probability space as X, independently of X, in such a way that the chain X has made exactly $\sigma(\boldsymbol{K})$ steps when the lazy chain completes its task. Figure 1 illustrates the construction.

Proof of Proposition 2.1. (i) Let $\xi_1, \xi_2, \dots$ be i.i.d. Bernoulli(q) random variables that are independent of X, and let $N_t = \xi_1 + \cdots + \xi_t$ and $Y_t = X_{N_t}$ for $t \geq 0$. Then Y is a Markov chain with transition matrix P_q started at x, and we define $\tau = \tau_{\boldsymbol{L}}(q)$ in terms of this chain. Since X visits every state infinitely often, $\tau < \infty$ almost surely.

For $n \geq 0$ let $H_n = \#\{t \geq 0 : N_t = n\}$ be the number of time units that Y spends at X_n before the next step of X. The holding times $H_0, H_1, \dots$ are i.i.d. with $\mathbb{P}(H_n = h) = (1-q)^{h-1}q$ for $h \geq 1$, and they are independent of X. During the j-th visit of X to v, the chain Y makes $H_{S_v(j)}$ visits to v. Hence the number of visits to v that X has to make before Y has visited v at least L_v times is

$$K_v = \min\bigl\{k \geq 1 :\ H_{S_v(1)} + \cdots + H_{S_v(k)} \geq L_v\bigr\},$$

and the L_v-th visit of Y to v occurs at a time t with $N_t = S_v(K_v)$. As N_t is nondecreasing in t, taking the maximum over v gives

$$N_\tau = \max_{v\in V} S_v(K_v) = \sigma(\boldsymbol{K}). \tag{2.5}$$

We now check that $\boldsymbol{K} = (K_v)_{v\in V}$ has the law described in the statement and is independent of X. The indices $S_v(j)$, for $v \in V$ and $j \geq 1$, are distinct and are determined by X. Therefore, conditionally on X, the variables $H_{S_v(j)}$ are i.i.d. with the same geometric law as H_0. It follows that the whole array $(H_{S_v(j)})_{v\in V,\, j\geq 1}$ is independent of X and has i.i.d. entries, and in particular the variables K_v are independent of each other and of X. To find the law of K_v, think of $H_{S_v(1)}, H_{S_v(2)}, \dots$ as the gaps between successive points of a renewal process on $\{1, 2, \dots\}$. Since the gaps are geometric, each integer is a renewal point with probability q, independently of the others, and $K_v - 1$ is the number of renewal points in $\{1, \dots, L_v - 1\}$. So $K_v - 1$ is binomial with parameters $L_v - 1$ and q.

It remains to compute $\mathbb{E}_x[N_\tau]$. Let $\mathcal{F}_t$ be the σ-field generated by X and $\xi_1, \dots, \xi_t$. Then τ is a stopping time with respect to $(\mathcal{F}_t)$, so the event $\{\tau \geq t\}$ is in $\mathcal{F}_{t-1}$ and is independent of ξ_t. Consequently

$$\mathbb{E}_x[N_\tau] = \sum_{t\geq 1} \mathbb{E}_x\bigl[\xi_t \mathbf{1}_{\{\tau\geq t\}}\bigr] = q \sum_{t\geq 1} \mathbb{P}_x(\tau \geq t) = q\, \mathbb{E}_x[\tau].$$

On the other hand, (2.5) and the independence of $\boldsymbol{K}$ and X give $\mathbb{E}_x[N_\tau] = \mathbb{E}_x[\sigma(\boldsymbol{K})] = \mathbb{E}[f_x(\boldsymbol{K})]$. This proves (2.3).

(ii) Let $(N_t)_{t\ge 0}$ be a Poisson process of rate $c\gamma$ which is independent of X, and let $Z_t = X_{N_t}$. Then Z is a continuous-time Markov chain with generator $c\gamma(P-I) = cQ$, started at x. This is the usual uniformization [8], and we define $\theta = \theta_{\boldsymbol{a}}(c)$ in terms of Z. For $n \ge 0$ let H_n be the length of the time interval on which $N_t = n$. The variables $H_0, H_1, \dots$ are i.i.d. exponential with rate $c\gamma$ and are independent of X. Let

$$K_v = \min\bigl\{k \ge 1 :\ H_{S_v(1)} + \cdots + H_{S_v(k)} \ge a_v\bigr\}.$$

As in part (i), the variables K_v are independent of each other and of X, and $K_v - 1$ is the number of points in $[0, a_v]$ of a Poisson process of rate $c\gamma$, so it is Poisson with mean $c\gamma a_v$. Almost surely none of the sums $H_{S_v(1)} + \cdots + H_{S_v(k)}$ is equal to a_v. Hence the occupation time of v reaches a_v in the interior of the time interval on which $N_t = S_v(K_v)$, and as before we get $N_\theta = \sigma(\boldsymbol{K})$ almost surely.

Finally, $N_t - c\gamma t$ is a martingale with respect to the filtration generated by X and $(N_s)_{s\le t}$, and θ is a stopping time for this filtration. Optional stopping at the bounded time $\theta \wedge t$ gives $\mathbb{E}_x[N_{\theta\wedge t}] = c\gamma\,\mathbb{E}_x[\theta \wedge t]$, and letting $t \to \infty$ we obtain $\mathbb{E}_x[N_\theta] = c\gamma\,\mathbb{E}_x[\theta]$ by monotone convergence. Since $\mathbb{E}_x[N_\theta] = \mathbb{E}[f_x(\boldsymbol{K})]$, this is (2.4). □

3 The key inequality

We keep the notation of Section 2; in particular X is a chain with transition matrix P started at x, and σ and f_x are defined in (2.1). Let $\mathbf{e}_v$ denote the v-th unit vector, and for $\boldsymbol{k} \in \mathbb{N}^V$ and $v \in V$ let

$$\Delta_v f_x(\boldsymbol{k}) = \begin{cases} f_x(\boldsymbol{k}) - f_x(\boldsymbol{k} - \mathbf{e}_v) & \text{if } k_v \ge 2, \\ 0 & \text{if } k_v = 1. \end{cases}$$

When $k_v \ge 2$, this is the expected time that is saved if the requirement at v is lowered by one.

Lemma 3.1. *For every $\boldsymbol{k} \in \mathbb{N}^V$,*

$$\sum_{v\in V} (k_v - 1)\,\Delta_v f_x(\boldsymbol{k}) \ \le\ f_x(\boldsymbol{k}). \tag{3.1}$$

The idea of the proof is simple. Lowering the requirement at v by one makes a difference only if v is the state whose requirement is fulfilled last, and in that case it saves at most the last of the $k_v - 1$ excursions from v that X has completed by then. These excursions are exchangeable, so on average the last one accounts for a fraction $1/(k_v - 1)$ of their total length, and this total is at most $\sigma(\boldsymbol{k})$. Hence the term for v in (3.1) is at most the contribution to $f_x(\boldsymbol{k})$ from the event that v is last.

In the proof, for a random variable ζ and an event G we write $\mathbb{E}_x[\zeta; G]$ for $\mathbb{E}_x[\zeta \mathbb{1}_G]$.

Proof of Lemma 3.1. Write $\sigma = \sigma(\boldsymbol{k})$, and let $W = X_\sigma$. The times $S_u(k_u)$, $u \in V$, are almost surely finite, and they are distinct because $X_{S_u(k_u)} = u$. So W is the unique state at which the maximum in (2.1) is attained, that is, the state whose requirement is fulfilled last.

Fix $v \in V$ with $k_v \ge 2$ and let $r = k_v - 1$. For $i \ge 1$ let $R_i = S_v(i+1) - S_v(i)$ be the length of the i-th excursion of X from v. We first compare σ with

$$\sigma(\boldsymbol{k} - \mathbf{e}_v) = \max\Bigl\{S_v(r),\ \max_{u\ne v} S_u(k_u)\Bigr\}.$$

If $W \ne v$, then the two times are equal. If $W = v$, then $\sigma = S_v(r+1)$ and $\sigma(\boldsymbol{k} - \mathbf{e}_v) \ge S_v(r)$, so the difference is at most R_r. In both cases

$$\sigma(\boldsymbol{k}) - \sigma(\boldsymbol{k} - \mathbf{e}_v) \ \le\ R_r \mathbb{1}_{\{W = v\}}. \tag{3.2}$$

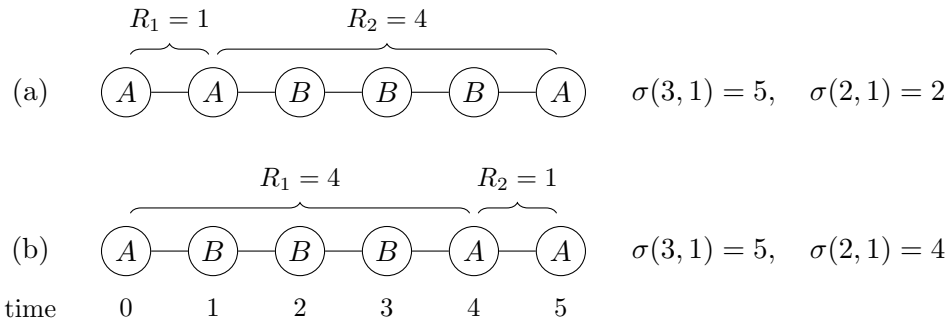


Figure 2: Two paths that differ only in the order of the first two excursions from A. Let $(k_A, k_B) = (3, 1)$. In both paths A is the last state to meet its requirement, at time 5. Lowering k_A to 2 saves three steps in (a) and one step in (b). For the path (a) we have $(k_A - 1)\,[\sigma(3,1) - \sigma(2,1)] = 6 > 5 = \sigma(3,1)$.

Next we use the exchangeability of the excursions. For $u \neq v$ let $C_i(u)$ be the number of visits to u during the i-th excursion from v, and let b_u be the number of visits to u before time $S_v(1)$. By the strong Markov property, the excursions are i.i.d. and are independent of the path of X up to time $S_v(1)$. Now $W = v$ if and only if every $u \neq v$ has been visited at least k_u times before time $S_v(k_v) = S_v(r+1)$, that is,

$$\{W = v\} = \Big\{ b_u + \sum_{i=1}^{r} C_i(u) \geq k_u \ \text{ for all } u \neq v \Big\}.$$

This event is not affected by a permutation of the first r excursions. So $\mathbb{E}_x[R_i;\, W = v]$ is the same for all $i \in \{1, \ldots, r\}$. Since $R_1 + \cdots + R_r = S_v(k_v) - S_v(1)$, and $S_v(k_v) = \sigma$ on the event $\{W = v\}$, it follows that

$$r\,\mathbb{E}_x[R_r;\, W = v] = \mathbb{E}_x\big[S_v(k_v) - S_v(1);\, W = v\big] \ \leq \ \mathbb{E}_x[\sigma;\, W = v].$$

Taking expectations in (3.2) and multiplying by $r = k_v - 1$, we obtain

$$(k_v - 1)\,\Delta_v f_x(\boldsymbol{k}) \ \leq \ \mathbb{E}_x[\sigma;\, W = v].$$

Finally we sum over all v with $k_v \geq 2$. The events $\{W = v\}$ are disjoint, so the right-hand sides add up to at most $\mathbb{E}_x[\sigma] = f_x(\boldsymbol{k})$, and (3.1) follows. □

The inequality (3.1) is a statement about expectations. It does not hold path by path, that is, with $\sigma(\boldsymbol{k})$ in place of $f_x(\boldsymbol{k})$ and $\sigma(\boldsymbol{k}) - \sigma(\boldsymbol{k} - \mathbf{e}_v)$ in place of $\Delta_v f_x(\boldsymbol{k})$; see Figure 2.

It will be convenient to write

$$g_x(\boldsymbol{k}) = f_x(\boldsymbol{k}) - \sum_{v \in V} (k_v - 1)\,\Delta_v f_x(\boldsymbol{k}), \tag{3.3}$$

so that $g_x \geq 0$ by Lemma 3.1. Note that $g_x(\mathbf{1}) = f_x(\mathbf{1})$ is the expected cover time of X, which is at least $|V| - 1 \geq 1$.

4 Proofs of the theorems

By Proposition 2.1, the expectations in Theorems 1.1 and 1.2 are of the form $\mathbb{E}[f_x(\boldsymbol{K})]$ divided by the parameter, where the coordinates of $\boldsymbol{K} - \mathbf{1}$ are independent binomial or Poisson variables. The next lemma shows how such an expectation changes with the parameter.

Lemma 4.1. *Let $\boldsymbol{K}(s) = (K_v(s))_{v\in V}$ be a random vector with independent coordinates, and suppose that one of the following holds.*

(a) *$K_v(s) - 1$ is binomial with parameters m_v and s, where $0 < s < 1$ and the m_v are nonnegative integers.*

(b) *$K_v(s) - 1$ is Poisson with mean $s\lambda_v$, where $s > 0$ and $\lambda_v > 0$.*

Let $F(s) = \mathbb{E}[f_x(\boldsymbol{K}(s))]$. Then

$$sF'(s) = \mathbb{E}\Big[\sum_{v\in V}\big(K_v(s)-1\big)\,\Delta_v f_x(\boldsymbol{K}(s))\Big], \tag{4.1}$$

and consequently

$$F(s) - sF'(s) = \mathbb{E}[g_x(\boldsymbol{K}(s))] \;\geq\; 0. \tag{4.2}$$

Proof. Let $p_v(j;s) = \mathbb{P}(K_v(s) = j)$ for $j \geq 1$. In both cases a direct computation gives

$$s\,\frac{\partial}{\partial s}\,p_v(j;s) = (j-1)\,p_v(j;s) - j\,p_v(j+1;s), \qquad j \geq 1. \tag{4.3}$$

Since $F(s) = \sum_{\boldsymbol{k}\in\mathbb{N}^V} f_x(\boldsymbol{k}) \prod_{u\in V} p_u(k_u;s)$, we can differentiate the product one factor at a time and use (4.3). This gives

$$sF'(s) = \sum_{v\in V}\sum_{\boldsymbol{k}\in\mathbb{N}^V} f_x(\boldsymbol{k})\big[(k_v-1)\,p_v(k_v;s) - k_v\,p_v(k_v+1;s)\big]\prod_{u\neq v} p_u(k_u;s).$$

In the terms that contain $p_v(k_v+1;s)$ we now replace $\boldsymbol{k}$ by $\boldsymbol{k}-\mathbf{e}_v$, and find that

$$sF'(s) = \sum_{v\in V}\sum_{\boldsymbol{k}:\,k_v\geq 2} (k_v-1)\big[f_x(\boldsymbol{k}) - f_x(\boldsymbol{k}-\mathbf{e}_v)\big]\prod_{u\in V} p_u(k_u;s),$$

which is (4.1). In case (a) all the sums are finite. In case (b) the differentiation under the sum and the rearrangement are justified by the bound (2.2) and the rapid decay of the Poisson weights, which is locally uniform in s. Finally, (4.2) follows from (4.1), the definition (3.3) and Lemma 3.1. □

Proof of Theorem 1.1. For $0 < q < 1$ let $\boldsymbol{K}(q)$ be as in Proposition 2.1(i), so that Lemma 4.1 applies in case (a) with $s = q$ and $m_v = L_v - 1$. By (2.3) we have $\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] = F(q)/q$, and therefore (4.2) gives

$$\frac{d}{dq}\,\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] = \frac{qF'(q) - F(q)}{q^2} = -\frac{\mathbb{E}[g_x(\boldsymbol{K}(q))]}{q^2} \;\leq\; 0, \qquad 0 < q < 1. \tag{4.4}$$

To see that the derivative is strictly negative, recall that $g_x \geq 0$ and that $g_x(\mathbf{1}) = f_x(\mathbf{1}) \geq 1$. Since $\mathbb{P}(\boldsymbol{K}(q) = \mathbf{1}) = (1-q)^m$ with $m = \sum_{v\in V}(L_v - 1)$, it follows from (4.4) that

$$\frac{d}{dq}\,\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)] \;\leq\; -\frac{(1-q)^m f_x(\mathbf{1})}{q^2} \;<\; 0, \qquad 0 < q < 1.$$

So $q \mapsto \mathbb{E}_x[\tau_{\boldsymbol{L}}(q)]$ is strictly decreasing on $(0,1)$. It is also continuous on $(0,1]$, since by (2.3) the function $q\,\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)]$ is a polynomial in q. Hence it is strictly decreasing on $(0,1]$.

For the last assertion, fix $0 < q' < q \leq 1$ and let x^* be a state that maximizes $\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)]$ over $x \in V$. Then

$$\max_{x\in V}\mathbb{E}_x[\tau_{\boldsymbol{L}}(q')] \;\geq\; \mathbb{E}_{x^*}[\tau_{\boldsymbol{L}}(q')] \;>\; \mathbb{E}_{x^*}[\tau_{\boldsymbol{L}}(q)] = \max_{x\in V}\mathbb{E}_x[\tau_{\boldsymbol{L}}(q)].$$ □

Proof of Theorem 1.2. For $c > 0$ let $\boldsymbol{K}(c)$ be as in Proposition 2.1(ii). Now Lemma 4.1 applies in case (b) with $s = c$ and $\lambda_v = \gamma a_v$, and (2.4) says that $\mathbb{E}_x[\theta_{\boldsymbol{a}}(c)] = F(c)/(c\gamma)$. As in the previous proof,

$$\frac{d}{dc}\,\mathbb{E}_x[\theta_{\boldsymbol{a}}(c)] = -\frac{\mathbb{E}[g_x(\boldsymbol{K}(c))]}{c^2\gamma} \;\le\; -\frac{f_x(\mathbf{1})}{c^2\gamma}\exp\Bigl(-c\gamma\sum_{v\in V} a_v\Bigr) \;<\; 0$$

for every $c > 0$, where we used that $\mathbb{P}(\boldsymbol{K}(c) = \mathbf{1}) = \exp(-c\gamma\sum_v a_v)$. The statement about the maximum over x follows as before. □

5 Examples

In this section we show that $\tau_{\boldsymbol{L}}(q)$ is not stochastically monotone in q, and that Theorems 1.1 and 1.2 fail for speed-ups that are not uniform. Two states are enough for all of this, so we let $V = \{A, B\}$, except at the end of Section 5.1.

5.1 No stochastic monotonicity

Let P be the transition matrix of the chain that alternates between A and B. Start at A and take $L_A = 3$ and $L_B = 1$. For $q = 1$ the chain follows the path $ABABA$, and $\tau_{\boldsymbol{L}}(1) = 4$. For $q < 1$ the lazy chain follows the path $AAAB$ with probability $(1-q)^2 q$, and then $\tau_{\boldsymbol{L}}(q) = 3$. So

$$\mathbb{P}_A\bigl(\tau_{\boldsymbol{L}}(q) \le 3\bigr) > 0 = \mathbb{P}_A\bigl(\tau_{\boldsymbol{L}}(1) \le 3\bigr),$$

and $\tau_{\boldsymbol{L}}(q)$ does not stochastically dominate $\tau_{\boldsymbol{L}}(1)$. In particular, there is no coupling of the two chains under which $\tau_{\boldsymbol{L}}(q) \ge \tau_{\boldsymbol{L}}(1)$ almost surely. The expectations are of course ordered. Here $f_A(\boldsymbol{k}) = 2(k_A - 1)$ if $k_A \ge 2$ and $k_B = 1$, and $f_A(\mathbf{1}) = 1$, so Proposition 2.1 gives

$$\mathbb{E}_A[\tau_{\boldsymbol{L}}(q)] = \frac{(1-q)^2 + 4q(1-q) + 4q^2}{q} = q + 2 + \frac{1}{q}.$$

The same thing happens when all the requirements are equal. For instance, consider simple random walk on a path with three vertices, started at an endpoint, and let $L_v = 2$ for all v. The walk is at the middle vertex at every odd time, so it needs the even times $0, 2, 4, 6$ in order to visit each endpoint twice, and its 2-cover time is at least 6. The lazy walk, on the other hand, can finish at time 5 by holding once at each vertex.

5.2 Holding probabilities that depend on the state

Consider the chain on $\{A, B\}$ that jumps from A to B with probability α and from B to A with probability β, and otherwise stays where it is. Let τ be the first time by which both states have been visited at least twice. We claim that

$$\mathbb{E}_A[\tau] = 1 - \alpha + \frac{1+\beta}{\alpha} + \frac{\alpha}{\beta}. \tag{5.1}$$

To see this, consider the first step. With probability $1 - \alpha$ the chain holds at A. It then takes $1/\alpha$ steps on average to reach B, and one step later the chain either has made its second visit to B, or is back at A (with probability β) and needs another $1/\alpha$ steps on average to return to B. In this case the expected remaining time after the first step is $1 + (1+\beta)/\alpha$. With probability α the chain jumps to B at the first step. One step later one of the two states has been visited twice, and the expected additional time needed for the other one is $1/\beta$ if the chain held at B and $1/\alpha$ if it jumped back to A. This gives an expected remaining time of $1 + (1-\beta)/\beta + \beta/\alpha$. Hence

$$\mathbb{E}_A[\tau] = 1 + (1-\alpha)\Bigl(1 + \frac{1+\beta}{\alpha}\Bigr) + \alpha\Bigl(\frac{1}{\beta} + \frac{\beta}{\alpha}\Bigr),$$

which simplifies to (5.1). By symmetry, $\mathbb{E}_B[\tau]$ is obtained by exchanging α and β.

Now take $(\alpha, \beta) = (1/4, 1/2)$, so that $\mathbb{E}_A[\tau] = 29/4$ and $\mathbb{E}_B[\tau] = 5$. If β is raised to $3/4$, these become $97/12$ and $59/12$. So the expected time from A, which is also the larger of the two, goes up. This is not very surprising, since leaving B more quickly makes it harder to collect visits there. However, the same can happen when the chain is sped up at both states: for $(\alpha, \beta) = (3/10, 1)$ we get $\mathbb{E}_A[\tau] = 23/3 > 29/4$ and $\mathbb{E}_B[\tau] = 139/30$, so the larger of the two expectations increases again.

5.3 Holding probabilities that depend on time

The holding probability must also be constant in time. Consider again the chain on $\{A, B\}$ started at A, with $L_A = L_B = 2$, but suppose now that the step from time $t-1$ to time t is a jump to the other state with probability q_t and a hold otherwise, where

$$q_1 = \frac{1}{4}, \qquad q_2 = \frac{3}{4}, \qquad q_3 = d, \qquad q_t = \frac{1}{2} \quad (t \geq 4).$$

Clearly $\tau \geq 3$. The possible paths up to time 2, their probabilities and the corresponding conditional expectations of $\tau - 3$ are as follows.

$Y_0Y_1Y_2$	AAA	AAB	ABB	ABA
probability	$3/16$	$9/16$	$1/16$	$3/16$
$\mathbb{E}[\tau - 3 \mid Y_0Y_1Y_2]$	$4 - 2d$	$2d$	$2 - 2d$	$2 - 2d$

For example, after AAB only one more visit to B is needed. With probability $1 - d$ the chain holds at the third step and $\tau = 3$. Otherwise it jumps to A, and then it needs two more steps on average to return to B. The other entries are found in the same way. Therefore

$$\mathbb{E}_A[\tau] = 3 + \frac{3(4 - 2d) + 9 \cdot 2d + (2 - 2d) + 3(2 - 2d)}{16} = \frac{17 + d}{4},$$

which is increasing in d. By symmetry the expected time is the same when the chain starts at B.

5.4 Jump rates increased by different factors

Finally we give a continuous-time example in which all the jump rates are increased, but not by the same factor. Consider the chain on $\{A, B\}$ that jumps from A to B at rate r_A and from B to A at rate r_B. Start at A, and let θ be the first time at which the chain has spent at least one unit of time in each state. Denote by J_A and J_B the numbers of jumps out of A and out of B before time θ. The number of jumps out of A by time t, minus r_A times the time spent at A by time t, is a martingale, and by optional stopping the expected time spent at A before θ is $\mathbb{E}_A[J_A]/r_A$. The same holds for B, and so

$$\mathbb{E}_A[\theta] = \frac{\mathbb{E}_A[J_A]}{r_A} + \frac{\mathbb{E}_A[J_B]}{r_B}. \tag{5.2}$$

Let M_A be the number of sojourns at A that end before the time spent at A reaches 1, and define M_B in the same way. Then M_A and M_B are independent Poisson variables with means r_A and r_B. The chain alternates between sojourns at A and at B, beginning with A, and θ falls in the $(M_A + 1)$-st sojourn at A or in the $(M_B + 1)$-st sojourn at B, whichever comes later. It follows that

$$J_A = \max\{M_A,\, M_B + 1\} \qquad \text{and} \qquad J_B = \max\{M_A,\, M_B\}.$$

If $r_A = r_B = 1$, then (5.2) gives $\mathbb{E}_A[\theta] \leq \mathbb{E}[M_A + M_B + 1] + \mathbb{E}[M_A + M_B] = 5$. If $(r_A, r_B) = (2, 10)$, it gives

$$\mathbb{E}_A[\theta] \;\geq\; \frac{\mathbb{E}[M_B + 1]}{2} + \frac{\mathbb{E}[M_B]}{10} = \frac{13}{2}.$$

(Numerically, the two expectations are about 3.70 and 6.50.) So $\mathbb{E}_A[\theta]$ increases when the rates are changed from $(1,1)$ to $(2,10)$, although both of them go up. The common factor c in Theorem 1.2 cannot be replaced by factors that depend on the state.

Acknowledgments

The research of Y. Peres was supported by the National Natural Science Foundation of China, RFIS grant No. W2531011.

The proofs of the theorems are due to the authors. Chatgpt 5.6 was used to explore prior work on multiple cover times; Chatgpt 6 and Claude opus 5.5 were used to generate examples and to polish the exposition.

References

[1] C. Avin, M. Koucký and Z. Lotker, Cover time and mixing time of random walks on dynamic graphs, *Random Structures Algorithms* **52** (2018), 576–596.

[2] M. T. Barlow, J. Ding, A. Nachmias and Y. Peres, The evolution of the cover time, *Combin. Probab. Comput.* **20** (2011), 331–345.

[3] S. O. Chan, Q. Ding and S. H. Li, Learning and testing irreducible Markov chains via the k-cover time, in *Proceedings of the 32nd International Conference on Algorithmic Learning Theory*, Proc. Mach. Learn. Res. **132** (2021), 458–480.

[4] Q. (Devon) Ding, *Aspects of local time processes with applications*, PhD thesis, University of California, Berkeley, 2026. Technical Report UCB/EECS-2026-169.

[5] J. Ding, J. R. Lee and Y. Peres, Cover times, blanket times, and majorizing measures, *Ann. of Math. (2)* **175** (2012), 1409–1471.

[6] O. Gurel-Gurevich, Y. Peres and O. Zeitouni, Localization for controlled random walks and martingales, *Electron. Commun. Probab.* **19** (2014), no. 24, 8 pp.

[7] J. Haslegrave, T. Sauerwald and J. Sylvester, Time dependent biased random walks, *ACM Trans. Algorithms* **18** (2022), no. 2, Art. 12, 30 pp.

[8] A. Jensen, Markoff chains as an aid in the study of Markoff processes, *Skand. Aktuarietidskr.* **36** (1953), 87–91.

[9] D. A. Levin and Y. Peres, *Markov chains and mixing times*, 2nd ed., with contributions by E. L. Wilmer, American Mathematical Society, Providence, RI, 2017.

[10] P. Matthews, Covering problems for Markov chains, *Ann. Probab.* **16** (1988), 1215–1228.

[11] D. S. (Seamus) Staite, *On the L-cover time of lazy random walks on graphs*, MSc thesis, University of Melbourne (2026).

[12] P. Winkler and D. Zuckerman, Multiple cover time, *Random Structures Algorithms* **9** (1996), 403–411.