Comparison Principle for Cover Times of Random Walks on Dynamical Percolation

Ishaan Bhadoo

Abstract

We answer Question 1.12 from [HS20]. The key is to use Lehec's inequality for general chains and the Ding–Lee–Peres theorem for simple random walks. More precisely, the auxiliary walk of Hermon–Sousi is generally non-reversible, so one cannot apply Ding–Lee–Peres directly to it. Instead, Lehec's one-sided non-reversible cover-time upper bound applies to the auxiliary walk, while Ding–Lee–Peres is used only for the reversible simple random walk.

Contents

1 Introduction

Percolation has been a popular model of study ever since it was introduced in 1957 by Broadbent and Hammersley [BH57]. This paper focuses on a dynamical adaptation of percolation where the percolation environment evolves as a two-state Markov chain. Each edge refreshes at rate μ and becomes open with probability p and closed with probability $1-p$. We study a model of random walk introduced by Peres, Steif, and Stauffer [PSS15] on this evolving environment, where the walker chooses one of its neighbors at rate 1 and jumps only if that edge is open.

A key step towards understanding the dynamical walk is to relate it to the well-understood simple random walk (SRW). This is when the walker chooses one of its neighbors at rate 1 and moves there. In this spirit, Hermon and Sousi [HS20] established several comparison principles between important random walk quantities, such as the hitting times of the dynamical walker and the SRW. They showed, for example,

$$t_{\text{hit}}^{(\mu,p)} \le \frac{C(\mu)}{p} t_{\text{hit}}^{\text{SRW}},$$

where μ is the refresh rate for the evolution of the configuration, and p is the probability that the edge is open after refreshment.

While [HS20] resolved many key aspects of hitting time behavior, they left open the question of whether a comparison principle could also be established for cover times [HS20, Question 1.12]. This paper addresses that open problem.

The following theorem gives an affirmative answer and thereby solves Question 1.12 of [HS20].

Theorem 1.1 (Cover-time comparison for RWDP)**.** *For every $\mu > 0$ there is a finite constant $C(\mu)$ such that for every finite connected simple graph G and every $p \in (0, 1]$,*

$$t_{\text{cov}}^{(\mu,p)}(G) \le \frac{C(\mu)}{p}\, t_{\text{cov}}^{\text{SRW}}(G). \tag{1}$$

Overview of the proof. We embed the auxiliary chain of Hermon–Sousi in the trajectory of the dynamical walker. Its additive symmetrization is compared with simple random walk through a Dirichlet-form estimate. Since the auxiliary chain need not be reversible, we apply Lehec's non-reversible cover-time upper bound to it and use the Ding–Lee–Peres theorem only for simple random walk. Finally, the regeneration-time construction transfers the resulting cover-time estimate from the auxiliary chain to the full dynamical walk.

Acknowledgements. The author thanks Perla Sousi for suggesting this topic of study and for many helpful discussions. This paper developed from the author's Part III essay[1] (for certain special classes of graphs, such a comparison was established in the essay.)

Statement on A.I. Use. The key reference [Leh14] was pointed out by GPT 5.5 after prompting it to find a cover time comparison result for non-reversible chains. Codex was used in the writing process. The author takes responsibility for the proofs in this note.

2 The model and random-walk quantities

2.1 Dynamical percolation

In this paper, we consider a dynamical version of percolation on a graph $G = (V, E)$. This model was first introduced in [HPS97] and is a model for an evolving percolation process. Given $p \in (0, 1]$ and a refresh rate $\mu > 0$, we construct a time-evolving random environment as follows. At time $t = 0$, the configuration is initialized according to some law on $\{0, 1\}^E$.

Each edge $e \in E(G)$ is equipped with an independent Poisson clock of rate μ. When the clock of edge e rings, the edge refreshes: it becomes open with probability p and closed with probability $1 - p$, independently of its previous state and independently of the rest of the configuration. We denote the configuration at time t by η_t, where $\eta_t(e) = 1$ if edge e is open at time t and $\eta_t(e) = 0$ otherwise.

The process $(\eta_t)_{t\ge 0}$ forms a continuous-time Markov process on $\{0, 1\}^E$. It is easy to check that π_p (the percolation measure on edges) is a stationary distribution for this process. In particular, if the initial distribution is π_p, then $\eta_t \sim \pi_p$ for all $t \ge 0$.

2.2 Random walk on dynamical percolation

We now define the model that we discuss for the remainder of the paper: random walk on dynamical percolation, introduced in [PSS15]. Let $(X_t)_{t\ge 0}$ denote the position of the walker at

[1] Available at https://ishaan44.github.io/assets/pdf/Part_III_Essay.pdf

time $t \geq 0$. The joint process $(X_t, \eta_t)_{t\geq 0}$ evolves as follows. The environment $(\eta_t)_{t\geq 0}$ evolves independently according to the dynamical percolation process described above.

Meanwhile, the walker (independently) evolves as a continuous-time random walk of rate 1. At each time when the walker's clock rings, the walker chooses one of its neighbors uniformly at random, jumps to that neighbour only if that edge is open at that moment (i.e., open in η_t). Otherwise, it remains at its current vertex.

This describes a Markov process (X_t, η_t) on $V \times \{0,1\}^E$. We emphasize that while $(X_t, \eta_t)_{t\geq 0}$ is Markovian on its own, the process (X_t) alone is not Markovian due to the dependence on the evolving environment.

From now on we restrict ourselves to the case of a finite connected simple graph $G = (V, E)$ with at least two vertices; the one-vertex case is trivial. It is easy to check that the product measure $\pi_{\text{full,p}} = \pi \times \pi_p$ is a stationary distribution for the full process (X_t, η_t), where π denotes the measure on V given by

$$\pi(x) = \frac{\deg(x)}{2|E|} \qquad \text{for each } x \in V.$$

Here, $\deg(x)$ denotes the degree of the vertex x in the graph G.

2.3 Hitting times and cover times

In this section, we introduce several important quantities associated with random walks on graphs, both in the static setting (simple random walk, or SRW) and the dynamic setting (random walk on dynamical percolation). Understanding these quantities will be crucial for the comparisons between the SRW and random walk on dynamical percolation that we develop later. We start by defining these concepts for a finite state Markov chain and then give specific notations for the case of SRW and random walk on dynamical percolation.

Definition 2.1 (Hitting times and cover times)**.** *Consider an ergodic Markov chain $(Z_t)_{t\geq 0}$ with state space Ω, transition matrix P and stationary distribution π. Then for each set $A \subseteq \Omega$ define $T_A := \inf\{t : Z_t \in A\}$; when A is a singleton $\{y\}$ we simply write T_y. The hitting time and cover time of the Markov chain are defined as*

$$t_{\text{hit}} := \max_{x,y} \mathbb{E}_x(T_y) \qquad \textit{and} \qquad t_{\text{cov}} := \max_{x} \mathbb{E}_x(\tau_{\text{cov}}),$$

where $\tau_{\text{cov}} = \inf\{t : \text{for every } x \in \Omega \text{ there exists } s \leq t \text{ such that } Z_s = x\}$ is the time taken to visit all points.

We refer to the Markov chain $(X_t, \eta_t)_{t\geq 0}$ as the full process. Let $\mathbb{P}^{(\mu,p)}_{x,\eta}$ and $\mathbb{E}^{(\mu,p)}_{x,\eta}$ denote the law and the expectation of the full process started from (x, η). We want to compare a simple random walker on G to the random walker on dynamical percolation. We consider the continuous time version of the simple random walker with jump rate 1 to avoid ergodic issues. Let us start by defining the following hitting times:

$$t^{\text{SRW}}_{\text{hit}} := \max_{x,y\in V} \mathbb{E}_x(T_y) \qquad \text{and} \qquad t^{(\mu,p)}_{\text{hit}} := \max_{x,y\in V,\, \eta\in\{0,1\}^E} \mathbb{E}^{(\mu,p)}_{x,\eta}(T_y).$$

Similarly, we define the following cover times,

$$t^{\text{SRW}}_{\text{cov}}(G) := \max_{x\in V} \mathbb{E}_x(\tau^{\text{SRW}}_{\text{cov}}) \qquad \text{and} \qquad t^{(\mu,p)}_{\text{cov}}(G) := \max_{x\in V,\, \eta\in\{0,1\}^E} \mathbb{E}^{(\mu,p)}_{x,\eta}(\tau^{(\mu,p)}_{\text{cov}}),$$

where $\tau^{(\mu,p)}_{\text{cov}}$ is the first time the random walk on dynamical percolation visits every vertex.

Remark. These are not the cover and hitting times in the usual sense of considering (X_t, η_t) a Markov chain. We define them in the above manner since we only care about the coordinate of the random walker.

3 The auxiliary chain and comparison inputs

We now introduce the auxiliary chain of [HS20]. It is embedded in the trajectory of the dynamical walker and will serve as a bridge between the full process and simple random walk.

Fix an ordering $E = \{e_1, \ldots, e_n\}$. For every edge e_i, create infinitely many formal copies $e_{i,1}, e_{i,2}, \ldots$; these copies are not themselves edges of G. Start with $X_0 = x$, $\eta_0 \sim \pi_p$, and define a time-evolving multiset of *infected edges* R_t. Initially $R_0 = \varnothing$. Whenever the clock of the walker rings, add to R_t the edge that the walker examines. If that edge is already infected, add instead its lowest-numbered copy that is not already in R_t.

When $R_{t-} = A$, give every edge in $E \setminus A$ its usual clock of rate μ, and generate one clock of rate $|A|\mu$ for the infected multiset. When an uninfected edge rings, refresh it to be open with probability p and closed with probability $1 - p$. When the infected-set clock rings, choose an element of A uniformly and remove it. If the chosen element is a genuine edge of E, refresh its state as well. With this construction, (X_t, η_t) has the transition rates of the full process.

Definition 3.1 (Auxiliary chain)**.** *Start with $X_0 = x$, $\eta_0 \sim \pi_p$, and $R_0 = \varnothing$. Set $\tau_0 = 0$. If τ_i has been defined, let $\tau_i + S_i$ be the first time after τ_i at which R_t becomes nonempty, and set*

$$\tau_{i+1} = \inf\{t \geq \tau_i + S_i : R_t = \varnothing\}.$$

The auxiliary chain is the discrete-time chain

$$Y_i = X_{\tau_i}, \qquad i \geq 0.$$

The process $|R_t|$ is a birth-and-death chain with transition rates

$$q(k, k+1) = 1, \qquad q(k, k-1) = \mu k.$$

This immediately gives the following properties; see Lemmas 3.5, 3.6, and 4.4 of [HS20].

Lemma 3.2 (Properties of the auxiliary chain)**.** *The chain Y is irreducible and has invariant distribution*

$$\pi(x) = \frac{\deg(x)}{2|E|}.$$

The increments $\sigma_i := \tau_i - \tau_{i-1}$ are i.i.d., have exponential tails, and satisfy

$$m_\mu := \mathbb{E}[\sigma_i] = e^{1/\mu}. \tag{2}$$

In particular,

$$\mathbb{E}[\sigma_i \mid \mathcal{F}_{\tau_{i-1}}] = m_\mu.$$

If P_{aux} is the transition matrix of Y and $P = P_{\text{SRW}}$ is the discrete-time SRW kernel, then for every edge $\{x, y\} \in E$,

$$P_{\text{aux}}(x, y) \geq \frac{p\mu}{1+\mu} P(x, y). \tag{3}$$

Consequently, if P^*_{aux} is the time reversal with respect to π and

$$S_{\text{aux}} := \frac{P_{\text{aux}} + P^*_{\text{aux}}}{2},$$

then

$$\mathcal{E}_{S_{\text{aux}}}(f, f) \geq \frac{p\mu}{1+\mu} \mathcal{E}_P(f, f) \qquad \textit{for all } f : V \to \mathbb{R}. \tag{4}$$

Proof. Only the transition comparison requires explanation. Starting from x, consider the event that the first attempted jump is to y, that $\{x,y\}$ is open, and that this edge refreshes before the walker's next attempt. On this event the next regeneration occurs with the walker at y. Its probability is

$$P(x,y)p\frac{\mu}{1+\mu},$$

which proves (3). Since P_{aux} and P have the same invariant distribution and P is reversible, the same bound holds for P^*_{aux}; averaging gives (4). □

We shall also use the infected-set construction to pass from an arbitrary initial environment to a stationary one. Start the infected set with every edge infected, and let θ be the first time it becomes empty. Since $|R_t|$ has upward rate 1 and downward rate $\mu|R_t|$,

$$\sup_{x,\eta}\mathbb{E}_{x,\eta}\theta \le c_\mu \log(|E|+1). \tag{5}$$

At time θ, every edge is distributed according to π_p, and the environment is independent of X_θ. Indeed, the present state of each edge comes from the last refresh at which that genuine edge left the infected set. Any later examination would infect it again, so at time θ none of these refreshed states has subsequently been examined by the walker.

For a finite irreducible Markov chain Q on V, define the commute-time metric

$$d_Q(x,y) := \mathbb{E}^Q_x T_y + \mathbb{E}^Q_y T_x.$$

Lemma 3.3. *The function d_Q is a metric on V.*

Proof. Symmetry is built into the definition, and positivity follows from irreducibility. The strong Markov property gives

$$\mathbb{E}_x T_z \le \mathbb{E}_x T_y + \mathbb{E}_y T_z, \qquad \mathbb{E}_z T_x \le \mathbb{E}_z T_y + \mathbb{E}_y T_x.$$

Adding these inequalities proves the triangle inequality. □

Definition 3.4. *Let $\gamma_2(V,\rho)$ denote Talagrand's admissible-partition functional:*

$$\gamma_2(V,\rho) = \inf_{(\mathcal{A}_k)}\sup_{x\in V}\sum_{k\ge 0} 2^{k/2}\operatorname{diam}_\rho(\mathcal{A}_k(x)).$$

Here $\mathcal{A}_0=\{V\}$, each $\mathcal{A}_{k+1}$ refines $\mathcal{A}_k$, and $|\mathcal{A}_k|\le 2^{2^k}$ for $k\ge 1$.

We now isolate the two cover-time results used in the proof. Once these inputs are stated, the comparison is short.

The important point is that the auxiliary chain need not be reversible. The following comparison with its additive symmetrization is therefore essential; it follows from the commute-time variational principle of Doyle and Steiner [DS11].

Lemma 3.5 (Symmetrization comparison)**.** *Let P be an irreducible finite Markov chain with stationary distribution π, and let*

$$S = \frac{P+P^*}{2}$$

be its additive symmetrization. Then, for all x,y,

$$d_P(x,y) \le d_S(x,y).$$

Lehec extended the cover-time upper bound of Ding–Lee–Peres from reversible chains to general chains. This is sufficient for our purposes and is the key input:

Lemma 3.6 ([Leh14], Upper bound of cover time)**.** *There is a universal constant $C_L < \infty$ such that, for every irreducible positive recurrent Markov chain P on a finite state space V,*

$$t_{\mathrm{cov}}(P) \le C_L \gamma_2(V, \sqrt{d_P})^2.$$

No reversibility assumption is required.

Next we state the following celebrated result of Ding-Lee-Peres:

Lemma 3.7 (Ding–Lee–Peres)**.** *There is a universal constant $C_D < \infty$ such that, for every finite reversible irreducible Markov chain R on V,*

$$\gamma_2(V, \sqrt{d_R})^2 \le C_D t_{\mathrm{cov}}(R).$$

In fact, Ding–Lee–Peres prove the corresponding two-sided comparison for reversible chains, but we only require this direction.

The next lemma transfers a comparison between metrics—in our case, commute-time metrics—to a comparison of γ_2.

Lemma 3.8 (Metric monotonicity of γ_2)**.** *If ρ_1 and ρ_2 are metrics on V and $\rho_1(x, y) \le a\rho_2(x, y)$ for all $x, y \in V$, then*

$$\gamma_2(V, \rho_1) \le a\gamma_2(V, \rho_2).$$

Consequently, if $d_1(x, y) \le Ld_2(x, y)$ for all x, y, then

$$\gamma_2(V, \sqrt{d_1})^2 \le L\gamma_2(V, \sqrt{d_2})^2.$$

Proof. For every $A \subseteq V$,

$$\mathrm{diam}_{\rho_1}(A) \le a \, \mathrm{diam}_{\rho_2}(A).$$

Thus every admissible partition sequence has γ_2 cost for ρ_1 at most a times its cost for ρ_2. Taking the infimum over admissible partition sequences proves the first claim. The second claim follows by applying the first with $\rho_i = \sqrt{d_i}$ and $a = \sqrt{L}$. □

Next up we have a comparison theorem between two chains one of which is reversible and the other one is not. This is exactly how we compare the auxiliary chain with the simple random walk.

Proposition 3.9 (A non-reversible cover-time comparison)**.** *Let P be an irreducible finite Markov chain on V, and let R be a reversible irreducible Markov chain on the same state space. Suppose that*

$$d_P(x, y) \le Ld_R(x, y) \qquad \forall x, y \in V. \tag{6}$$

Then there is a universal constant $C < \infty$ such that

$$t_{\mathrm{cov}}(P) \le CL \, t_{\mathrm{cov}}(R).$$

Proof. Since P is finite and irreducible, it is positive recurrent. By Lemma 3.6,

$$t_{\mathrm{cov}}(P) \le C_L \gamma_2(V, \sqrt{d_P})^2.$$

By Lemma 3.8 and (6),

$$\gamma_2(V, \sqrt{d_P})^2 \le L\gamma_2(V, \sqrt{d_R})^2.$$

Since R is reversible, Lemma 3.7 gives

$$\gamma_2(V, \sqrt{d_R})^2 \le C_D t_{\mathrm{cov}}(R).$$

Combining the three inequalities proves the result. □

4 Proof of the cover-time comparison

In this section we prove Theorem 1.1. The proof goes in the same way as in [HS20] using the above estimates.

The appearance of $1/p$ can also be explained. It is the factor by which the walker is slowed down because at any point, it only sees a p-fraction of the edges that the simple random walk sees.

Proof of Theorem 1.1. We first prove the cover-time comparison for the auxiliary chain. By Lemma 3.2,

$$\mathcal{E}_{S_{\text{aux}}}(f,f) \geq \alpha \mathcal{E}_P(f,f), \qquad \alpha = \frac{p\mu}{1+\mu}.$$

Since both S_{aux} and P are reversible with respect to π, the Dirichlet principle gives

$$d_{S_{\text{aux}}}(x,y) \leq \alpha^{-1} d_{\text{SRW}}(x,y) \qquad \forall x,y \in V.$$

By the comparison, Lemma 3.5,

$$d_{\text{aux}}(x,y) \leq d_{S_{\text{aux}}}(x,y).$$

Therefore

$$d_{\text{aux}}(x,y) \leq \frac{1+\mu}{p\mu}\, d_{\text{SRW}}(x,y) \qquad \forall x,y \in V. \tag{7}$$

Applying Proposition 3.9 with $P = P_{\text{aux}}$, $R = P$, and $L = (1+\mu)/(p\mu)$ gives

$$t_{\text{cov}}(Y) \leq C \frac{1+\mu}{p\mu}\, t_{\text{cov}}^{\text{SRW}}(G), \tag{8}$$

where C is universal.

We now transfer the auxiliary estimate to the full walk. First start the environment from π_p, independently of $X_0 = x$. Let

$$N_{\text{cov}}^{\text{aux}} := \inf\{k \geq 0 : \{Y_0, \ldots, Y_k\} = V\}.$$

Because $Y_k = X_{\tau_k}$, the full walk has certainly covered all vertices by the real time $\tau_{N_{\text{cov}}^{\text{aux}}}$. Hence, pathwise,

$$\tau_{\text{cov}}^{(\mu,p)} \leq \tau_{N_{\text{cov}}^{\text{aux}}}. \tag{9}$$

Let $\mathcal{F}_i := \mathcal{F}_{\tau_i}$. Since the event $\{N_{\text{cov}}^{\text{aux}} \geq i\}$ is determined by $Y_0, \ldots, Y_{i-1}$, it is $\mathcal{F}_{i-1}$-measurable. Lemma 3.2 gives

$$\mathbb{E}[\sigma_i \mid \mathcal{F}_{i-1}] = m_\mu.$$

Thus

$$\begin{aligned}
\mathbb{E}_x[\tau_{N_{\text{cov}}^{\text{aux}}}] &= \mathbb{E}_x\left[\sum_{i\geq 1} \mathbf{1}_{\{N_{\text{cov}}^{\text{aux}} \geq i\}} \sigma_i\right] \\
&= \sum_{i\geq 1} \mathbb{E}_x\left[\mathbf{1}_{\{N_{\text{cov}}^{\text{aux}} \geq i\}} \mathbb{E}[\sigma_i \mid \mathcal{F}_{i-1}]\right] \\
&= m_\mu \sum_{i\geq 1} \mathbb{P}_x(N_{\text{cov}}^{\text{aux}} \geq i) = m_\mu \mathbb{E}_x[N_{\text{cov}}^{\text{aux}}].
\end{aligned}$$

Together with (9) and (8), this yields

$$\max_x \mathbb{E}_{x,\pi_p}^{(\mu,p)}[\tau_{\text{cov}}^{(\mu,p)}] \leq C m_\mu \frac{1+\mu}{p\mu}\, t_{\text{cov}}^{\text{SRW}}(G). \tag{10}$$

It remains to pass from stationary environment to worst-case environment. The stopping time θ constructed above satisfies, for all initial (x, η),

$$\mathbb{E}_{x,\eta}\theta \leq c_\mu \log(|E| + 1),$$

while at time θ the environment is distributed as π_p and is independent of the current walk coordinate. Thus the strong Markov property and (10) give

$$t_{\mathrm{cov}}^{(\mu,p)}(G) \leq c_\mu \log(|E| + 1) + C m_\mu \frac{1+\mu}{p\mu}\, t_{\mathrm{cov}}^{\mathrm{SRW}}(G).$$

For connected simple graphs with $|V| \geq 2$, $|E| \leq |V|^2$ and $t_{\mathrm{cov}}^{\mathrm{SRW}}(G) \geq |V| - 1$, so the logarithmic term is absorbed into the right-hand side of (1); the one-vertex graph is trivial. This proves the theorem. □

References

[BH57] S. R. Broadbent and J. M. Hammersley. "Percolation processes: I. Crystals and mazes". In: *Mathematical Proceedings of the Cambridge Philosophical Society* 53.3 (1957), pp. 629–641.

[DLP12] J. Ding, J. R. Lee, and Y. Peres. "Cover times, blanket times, and majorizing measures". In: *Annals of Mathematics* 175.3 (2012), pp. 1409–1471.

[DS11] P. G. Doyle and J. Steiner. "Commuting time geometry of ergodic Markov chains". In: *arXiv preprint arXiv:1107.2612* (2011). arXiv: 1107.2612 [math.PR].

[HPS97] O. Häggström, Y. Peres, and J. E. Steif. "Dynamical percolation". In: *The Annals of Probability* 25.3 (1997), pp. 1269–1295.

[HS20] J. Hermon and P. Sousi. "A comparison principle for random walk on dynamical percolation". In: *The Annals of Probability* 48.6 (2020), pp. 2952–2987.

[Leh14] J. Lehec. "Cover times and generic chaining". In: *Journal of Applied Probability* 51.1 (2014), pp. 247–261.

[PSS15] Y. Peres, A. Stauffer, and J. E. Steif. "Random walks on dynamical percolation: mixing times, mean squared displacement and hitting times". In: *Probability Theory and Related Fields* 162.3–4 (2015), pp. 487–530.