Quenched exit times for random walk on dynamical percolation

Yuval Peres[*]   Perla Sousi[†]   Jeffrey E. Steif[‡]

November 2, 2018

Abstract

We consider random walk on dynamical percolation on the discrete torus $\mathbb{Z}_n^d$. In previous work, mixing times of this process for $p < p_c(\mathbb{Z}^d)$ were obtained in the annealed setting where one averages over the dynamical percolation environment. Here we study exit times in the *quenched* setting, where we condition on a typical dynamical percolation environment. We obtain an upper bound for all p which for $p < p_c$ matches the known upper bound.

Keywords and phrases. Dynamical percolation, random walk, hitting times, mixing times.
MSC 2010 *subject classifications.* Primary 60K35, 60K37

1 Introduction

In this paper, we study quenched mixing results for random walk on dynamical percolation on the torus $\mathbb{Z}_n^d$ of side length n with parameters p and $\mu \leq 1/2$. Let each edge evolve independently where an edge in state 0 (absent, closed) switches to state 1 (present, open) at rate $p\mu$ and an edge in state 1 switches to state 0 at rate $(1-p)\mu$. Let $(\eta_t)_{t \geq 0}$ denote the resulting Markov process on $\{0,1\}^{E(\mathbb{Z}_n^d)}$ whose stationary distribution is product measure with density p, denoted by π_p; this model is called *dynamical percolation*. We next perform a random walk on the evolving graph $(\eta_t)_{t \geq 0}$ by having the random walker at rate 1 choose a neighbour (in the original graph) uniformly at random and move there if (and only if) the connecting edge is open at that time. Letting $(X_t)_{t \geq 0}$ denote the position of the walker at time t, we have, when initial configurations are given, that

$$(M_t)_{t \geq 0} := ((X_t, \eta_t))_{t \geq 0}$$

is a Markov process while $(X_t)_{t \geq 0}$ of course is not. It is easy to see that the unique stationary distribution for (M_t) is $\pi \times \pi_p$, where π is uniform on $\mathbb{Z}_n^d$, and that this measure is reversible.

In [5], a number of annealed results were obtained for this model where one has d and p fixed while μ and n are considered the important parameters with respect to which we want to study the model. Also in [1] the authors study the mixing time of the non-backtracking random walk on certain dynamical configuration models.

We summarise here the relevant results obtained in [5] concerning mixing time.

[*]Microsoft Research, Redmond WA, U.S.A. Email: peres@microsoft.com
[†]University of Cambridge, Cambridge, UK. Email: p.sousi@statslab.cam.ac.uk
[‡]Chalmers University of Technology and Gothenburg University, Gothenburg, Sweden. Email: steif@chalmers.se

Even though mixing times are traditionally defined only for Markov chains, one can easily adapt the definition to cases like X above as follows. For $\varepsilon \in (0,1)$ and η_0 a configuration of edges we write
$$t_{\mathrm{mix}}(\varepsilon, \eta_0) = \min \left\{ t \geq 0 : \max_x \|\mathbb{P}_{x,\eta_0}(X_t = \cdot) - \pi\|_{\mathrm{TV}} \leq \varepsilon \right\}.$$

We write $\mathbb{P}_{x,\pi_p}()$, when the environment process starts from stationarity and the random walk starts from x. We then write
$$t_{\mathrm{mix}}(\varepsilon, \pi_p) = \min \left\{ t \geq 0 : \left\|\mathbb{P}_{0,\pi_p}(X_t = \cdot) - \pi\right\|_{\mathrm{TV}} \leq \varepsilon \right\}.$$

As usual we let $p_c(\mathbb{Z}^d)$ be the critical value for bond percolation on $\mathbb{Z}^d$ and $\theta_d(p)$ be the probability that the origin is in an infinite cluster of $\mathbb{Z}^d$ at percolation parameter p.

The following describes the subcritical picture very well.

Theorem 1.1 ([5]). *For any $d \geq 1$, $\varepsilon > 0$ and $p \in (0, p_c(\mathbb{Z}^d))$, there exists $C = C(d, \varepsilon, p) \in (0, \infty)$ and $n_0 = n_0(d, \varepsilon, p) \in \mathbb{N}$ such that, for all $n \geq n_0$ and for all $\mu \leq 1/2$, we have*

$$\frac{n^2}{C\mu} \leq t_{\mathrm{mix}}(\varepsilon, \pi_p) \leq \sup_{\eta_0} t_{\mathrm{mix}}(\varepsilon, \eta_0) \leq \frac{Cn^2}{\mu}.$$

The following yields lower bounds throughout the whole parameter space of p.

Theorem 1.2 ([5]). *(i) Given $d \geq 1$ and $\varepsilon > 0$, there exist $C_1 = C_1(d, \varepsilon) > 0$ and $n_0 = n_0(d, \varepsilon)$ such that, for all p, for all $n \geq n_0$ and for all $\mu \leq 1/2$, we have*

$$t_{\mathrm{mix}}(\varepsilon, \pi_p) \geq C_1 n^2.$$

(ii) Given $d \geq 1$, p and $\varepsilon < 1 - \theta_d(p)$, there exists $C_2 = C_2(d, p, \varepsilon) > 0$ and $n_0 = n_0(d, p, \varepsilon)$ such that, for all $n \geq n_0$ and for all $\mu \leq 1/2$, we have

$$t_{\mathrm{mix}}(\varepsilon, \pi_p) \geq \frac{C_2}{\mu}. \tag{1.1}$$

In particular, for $\varepsilon < 1 - \theta_d(p)$, we get a lower bound for $t_{\mathrm{mix}}(\varepsilon, \pi_p)$ of order $n^2 + \frac{1}{\mu}$.

Remarks (i). In the usual theory of Markov chains, a lower bound on the ε-mixing time for a fixed ε (small) would yield a lower bound of a similar order (depending on ε) on the mixing time when $\varepsilon = 1/4$; this is however not the case here which is not a contradiction since $(X_t)_{t \geq 0}$ is not a Markov chain.
(ii). We believe (as stated in [5]) that in the supercritical regime, the mixing time is much faster than in the subcritical regime and has order at most $\frac{1}{\mu} + n^2$. Despite this, the methods in [5] did not even yield the much larger (subcritical) upper bound of $\frac{n^2}{\mu}$ in the supercritical regime. One of the corollaries of one of our main results is to obtain a related upper bound uniform in p (away from 0 and 1).

The above results concerned *annealed* mixing times, meaning that the marginal distribution of $(X_t)_{t \geq 0}$ is studied. Here we study mixing and exit times of the conditional distribution of $(X_t)_{t \geq 0}$ given (typical) $(\eta_t)_{t \geq 0}$; in other words, we study the *quenched* mixing and exit time behaviour of $(X_t)_{t \geq 0}$.

For certain results, an annealed version immediately yields a quenched version. This is true (due to Fubini's Theorem) for almost sure results such as recurrence and transience. Note, on the

2

other hand, that annealed upper bounds on mixing times do not automatically pass to quenched upper bounds on mixing times. One way to see this is to observe that if a convex combination of probability measures is close in total variation to some probability measure ν, it still of course may be the case that all the probability measures appearing in the convex combination are far in total variation from ν. An example of this, in the context of a Markov chain in a randomly evolving environment, is the following. Let $(M_n)_{n \geq 1}$ be an i.i.d. sequence of 2×2 matrices where each matrix is either

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

or

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

each with probability $1/2$. Let $(X_k)_{k \geq 0}$ be the process on $\{0, 1\}$ which at time n jumps according to the matrix M_{n+1}. It is clear that the annealed mixing time is 1 since, independent of the starting distribution for X_0, the distribution of X_1 is uniform. However the quenched mixing time is always infinite, since if we condition on any "environment" $(M_n)_{n \geq 1}$, the resulting time inhomogeneous Markov chain is such that for every k, X_k is deterministic. On the other hand, an appropriately defined quenched mixing upper bound easily yields an annealed mixing upper bound. A version of this is given by Proposition 1.5.

Next we write $\mathbb{P}_{x,\eta}(\cdot)$ for the probability measure of the walk, when the environment process is conditioned to be $\eta = (\eta_t)_{t \geq 0}$ and the walk starts from x. We write $\mathcal{P}$ for the distribution of the environment which is dynamical percolation on the torus, a measure on càdlàg paths $[0, \infty) \mapsto \{0, 1\}^{E(\mathbb{Z}_n^d)}$. We write $\mathcal{P}_{\eta_0}$ to denote the measure $\mathcal{P}$ when the starting environment is η_0. Abusing notation we write $\mathbb{P}_{x,\eta_0}(\cdot)$ to mean the law of the full system when the walk starts from x and the initial configuration of the environment is η_0. To distinguish it from the quenched law, we always write η_0 in the subscript as opposed to η.

Now we discuss hitting time bounds in both the quenched and annealed settings. The bounds we obtain are valid for all values of the percolation parameter p.

Let $A \subseteq \mathbb{Z}_n^d$. We denote by τ_A the first hitting time of A by X, i.e.

$$\tau_A = \inf\{t \geq 0 : X_t \in A\}.$$

Theorem 1.3. *For all $d \geq 1$ and $\delta > 0$, there exists $C = C(d, \delta) < \infty$ and $c = c(d, \delta) < 1$, so that for all $p \in [\delta, 1]$, for all n and for all $\mu \leq 1/2$, random walk in dynamical percolation on $\mathbb{Z}_n^d$ with parameters p and μ satisfies that for all $A \subseteq \mathbb{Z}_n^d$ with $|A| \geq n^d/2$ and for all k*

$$\max_{\eta_0} \mathcal{P}_{\eta_0}\left(\eta = (\eta_t)_{t \geq 0} : \max_x \mathbb{E}_{x,\eta}[\tau_A] \geq k \cdot \frac{Cn^2 \log n}{\mu}\right) \leq c^k \ \text{and}$$

$$\max_{x,\eta_0} \mathbb{E}_{x,\eta_0}[\tau_A] \leq \frac{Cn^2}{\mu}.$$

Remark 1.4. We note that the second statement for $p < p_c$ follows from Theorem 1.1.

For $\varepsilon \in (0, 1)$, $x \in \mathbb{Z}_n^d$ and a fixed environment $\eta = (\eta_t)_{t \geq 0}$ we write

$$t_{\text{mix}}(\varepsilon, x, \eta) = \min\left\{t \geq 0 : \|\mathbb{P}_{x,\eta}(X_t = \cdot) - \pi\|_{\text{TV}} \leq \varepsilon\right\}.$$

We also write

$$t_{\text{mix}}(\varepsilon, \eta) = \max_x t_{\text{mix}}(\varepsilon, x, \eta)$$

3

which we refer to as the quenched ε-mixing time. We remark that $t_{\mathrm{mix}}(\varepsilon, \eta)$ could be infinite. Using the obvious definitions, the standard inequality $t_{\mathrm{mix}}(\varepsilon) \leq \log_2(1/\varepsilon) t_{\mathrm{mix}}(1/4)$ does not necessarily hold for time-inhomogeneous Markov chains and therefore also not for quenched mixing times. Therefore, in such situations, to describe the rate of convergence to stationarity, it is more natural to give bounds on $t_{\mathrm{mix}}(\varepsilon, \eta)$ for all ε rather than just considering $\varepsilon = 1/4$.

Proposition 1.5. *For all $d \geq 1$ and $\delta > 0$, there exists $C = C(d, \delta) < \infty$ so that for all $p \in [\delta, 1]$, for all n, for all $\mu \leq 1/2$ and for all ε, random walk in dynamical percolation on $\mathbb{Z}_n^d$ with parameters p and μ satisfies for all $x \in \mathbb{Z}_n^d$*

$$\max_{\eta_0} \mathcal{P}_{\eta_0}\left(\eta = (\eta_t)_{t \geq 0} : t_{\mathrm{mix}}(\varepsilon, \eta, x) \geq \frac{Cn^2 \log(1/\varepsilon)}{\mu^4}\right) \leq \varepsilon. \tag{1.2}$$

We next discuss quenched lower bounds on the mixing time. It is now important whether p belongs to the sub or supercritical regime for percolation. In [5], it was proved that $\frac{n^2}{\mu}$ is the correct order of the (annealed) mixing time in the subcritical regime and, as already stated, it was conjectured there that the mixing time in the supercritical regime is much faster.

Proposition 1.6. *For all $d \geq 1$, $p \in (0, p_c(\mathbb{Z}^d))$, $\varepsilon > 0$ and M, there exists $\beta = \beta(d, p, \varepsilon, M) > 0$ and $n_0 = n_0(d, p, \varepsilon, M)$ so that if $(\eta_t)_{t \geq 0}$ is dynamical percolation started in stationarity, then for all $n \geq n_0$ we have*

$$\mathcal{P}_{\pi_p}\left(\eta = (\eta_t)_{t \geq 0} : t_{\mathrm{mix}}(1 - \varepsilon, 0, \eta) \leq \frac{\beta n^2}{\mu}\right) \leq \frac{1}{M}. \tag{1.3}$$

Proof. Fix $d \geq 1$, $p \in (0, p_c(\mathbb{Z}^d))$, ε and M. Let $\sigma := \min\{\varepsilon^2, \frac{1}{M^2}\}$. By Theorem 1.1 there exists $\beta = \beta(d, p, \sigma)$ so that for all large n and for all $\mu \leq 1/2$,

$$\left\|\mathbb{P}_{0, \pi_p}\left(X_{\frac{\beta n^2}{\mu}} = \cdot\right) - \pi\right\|_{\mathrm{TV}} \geq 1 - \sigma.$$

Since

$$\mathbb{P}_{0, \pi_p}\left(X_{\frac{\beta n^2}{\mu}} = \cdot\right) = \int \mathbb{P}_{0, \eta}\left(X_{\frac{\beta n^2}{\mu}} = \cdot\right) d\mathcal{P}_{\eta_0}((\eta_t)_{t \geq 0}) \, d\pi_p(\eta_0),$$

convexity of the total variation norm in the sense that

$$\left\|\int \mu_\alpha d\rho(\alpha) - \pi\right\|_{\mathrm{TV}} \leq \int \|\mu_\alpha - \pi\|_{\mathrm{TV}} \, d\rho(\alpha)$$

yields that

$$\int \left\|\mathbb{P}_{0, \eta}\left(X_{\frac{\beta n^2}{\mu}} = \cdot\right) - \pi\right\|_{\mathrm{TV}} d\mathcal{P}_{\eta_0}((\eta_t)_{t \geq 0})) \, d\pi_p(\eta_0) \geq 1 - \sigma, \tag{1.4}$$

where $\eta = (\eta_t)_{t \geq 0}$. This now implies that

$$\mathcal{P}\left(\eta = (\eta_t)_{t \geq 0} : \left\|\mathbb{P}_{0, \eta}\left(X_{\frac{\beta n^2}{\mu}} = \cdot\right) - \pi\right\|_{\mathrm{TV}} \leq 1 - \sigma^{\frac{1}{2}}\right) \leq \sigma^{\frac{1}{2}}$$

Since $\sigma := \min\{\varepsilon^2, \frac{1}{M^2}\}$, this gives the result. $\square$

The following is a quenched lower bound in the context of Theorem 1.2. The first statement is proved as the previous result. The second one follows from the fact that with high probability the environment will be such that there will exist a vertex isolated throughout the interval $[0, \beta/\mu]$.

4

Proposition 1.7. *Given $d \geq 1$, $\varepsilon \in (0,1)$, $p < 1$ and M, there exist $\beta > 0$ and $n_0 > 0$ such that, for all $n \geq n_0$ and for all $\mu \leq 1/2$, if $(\eta_t)_{t \geq 0}$ is dynamical percolation started in stationarity, then*

$$\mathcal{P}_{\pi_p}\left(\eta = (\eta_t)_{t \geq 0} : t_{\mathrm{mix}}(\varepsilon, 0, \eta) \leq \beta n^2\right) \leq \frac{1}{M} \quad and$$

$$\mathcal{P}_{\pi_p}\left(\eta = (\eta_t)_{t \geq 0} : t_{\mathrm{mix}}(\varepsilon, 0, \eta) \leq \frac{\beta}{\mu}\right) \leq \frac{1}{M}.$$

Mixing times in the supercritical case will be studied in [4]. Some of the results in this paper will be used there.

In Section 2, we state a general result, Theorem 2.1, giving an upper bound on the quenched mixing time for a Markov process in a random environment. In Section 3, Theorem 2.1 is applied to obtain Theorem 1.3 and Proposition 1.5. Finally Theorem 2.1 is proved in Section 4.

2   General setup

Given a general finite state Markov chain $p(\cdot, \cdot)$ with state space Ω and a stationary distribution π, we let

$$Q_p(A, B) = Q_{p,\pi}(A, B) := \sum_{x \in A, y \in B} \pi(x) p(x, y) \tag{2.1}$$

for $A, B \subseteq \Omega$. Also, for $S \subseteq \Omega$, we let

$$\varphi_p(S) = \varphi_{p,\pi}(S) := \frac{Q_p(S, S^c)}{\pi(S)}.$$

Observe that

$$\varphi_p(S) = \mathbb{P}(X_1 \notin S \mid X_0 \in S)$$

where $(X_k)_{k \in \mathbb{Z}}$ is the stationary Markov chain associated to $p(x, y)$ and π. We call $\varphi_p(S)$ the expansion of S (relative to the Markov chain $p(\cdot, \cdot)$ and stationary distribution π). Note that $p(\cdot, \cdot)$ may have more than one stationary distribution and so we need to make π explicit.

Finally we recall standard notation. Let μ and ν be two probability measures on the same space Ω. We write

$$\chi^2(\mu, \nu) = \sum_{x \in \Omega} \nu(y) \left(\frac{\mu(y)}{\nu(y)} - 1\right)^2 = \sum_y \frac{\mu(y)^2}{\nu(y)} - 1.$$

By Cauchy-Schwarz we have

$$2 \left\| \mu - \nu \right\|_{\mathrm{TV}} \leq \chi(\mu, \nu).$$

We will now consider the following general set up of a finite state Markov chain in a Markovian evolving environment.

Let $\mathcal{E}$ be a state space for a discrete time homogeneous Markov chain η with transition kernel R. Moreover, for every $\zeta \in \mathcal{E}$ let $(p_\zeta(x, y))_{x,y \in \mathcal{S}}$ be a transition matrix on a finite state space $\mathcal{S}$. Assume π is a probability distribution on $\mathcal{S}$ which is stationary for p_ζ for all $\zeta \in \mathcal{E}$ and has full support.

We now define an annealed discrete time Markov chain (X, η) on $\mathcal{S} \times \mathcal{E}$ evolving as follows: when in state (x, ζ), it jumps to the state (x', ζ') by first choosing ζ' at random according to $R(\zeta, \cdot)$ and then choosing x' at random according to $p_{\zeta'}(x, \cdot)$. In symbols if Q denotes the annealed transition kernel we have

$$Q((x, \zeta), (x', d\zeta')) = R(\zeta, d\zeta') p_{\zeta'}(x, x').$$

5

Given a realisation of the environment process, $\eta = (\eta_i)_{i \geq 0}$, the coordinate X becomes a time inhomogeneous Markov chain with transition matrix p_{η_i} at time $i - 1$.

Observe that since $\pi p_\zeta = \pi$ for all $\zeta \in \mathcal{E}$, it follows easily that $\max_x \|\mathbb{P}_{x,\eta}(X_k = \cdot) - \pi\|_{\mathrm{TV}}$ is decreasing in k for any fixed environment η. Next, as defined in the introduction we let

$$t_{\mathrm{mix}}(\varepsilon, \eta) := \inf \left\{ k : \max_x \|\mathbb{P}_{x,\eta}(X_k = \cdot) - \pi\|_{\mathrm{TV}} \leq \varepsilon \right\} = \max_{x \in \mathcal{S}} t_{\mathrm{mix}}(\varepsilon, x, \eta)$$

be the ε-mixing time in the environment η.

The following general theorem yields quenched upper bounds on the mixing time in our general set up of a Markov chain in a Markovian evolving environment. We let $\pi_\star := \min_x \pi(x)$ below. For $\zeta \in \mathcal{E}$ and $S \subseteq \mathcal{S}$ we let

$$\varphi(\zeta, S) := \mathbb{E}_\zeta \left[\varphi_{p_{\eta_1}}(S) \right]. \tag{2.2}$$

Notice that in the previous expression we average over the new environment η_1, i.e. we run the environment process for one step starting from ζ and use the transition matrix that it yields. For $r \in [\pi_\star, \frac{1}{2}]$, let

$$\varphi(r) := \inf\{\varphi(\zeta, S) : \zeta \in \mathcal{E}, \pi(S) \leq r\}$$

and $\varphi(r) := \varphi(\frac{1}{2})$ for $r \geq \frac{1}{2}$. Clearly $\varphi(r)$ is weakly decreasing in r. It is crucial for our applications that in the above definitions, the minimisation over S occurs outside of the expectation in (2.2) rather than inside. If the minimum was taken on the inside, then $\varphi(r)$ would be much smaller and the following result therefore would be much weaker.

Theorem 2.1. *Consider a finite state Markov chain X in a Markovian evolving environment satisfying $p_\zeta(x, x) \geq \gamma$ for all ζ and all x with $\gamma \in (0, 1/2]$. For all $\varepsilon > 0$ and $x \in \mathcal{S}$ if*

$$n \geq 1 + \frac{2(1 - \gamma)^2}{\gamma^2} \int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u \varphi^2(u)}$$

then for all $\zeta \in \mathcal{E}$,

$$\mathcal{P}_\zeta \left(\eta = (\eta_t)_{t \geq 0} : \chi(\mathbb{P}_{x,\eta}(X_n = \cdot), \pi) \geq \varepsilon^{1/4} \right) \leq \varepsilon^{1/4}.$$

Remark 2.2. We note that the above theorem remains true in the following variant of the Markov chain described above. Suppose that at every step, the chain X remains in place with probability $1/2$ and with probability $1/2$ it jumps according to the transition kernel given by the environment at this time. When X stays in place (because of laziness), then the environment at the next step also stays in place, otherwise it moves according to its transition kernel. So the transition matrix of the environment depends on the extra randomness coming from whether X made an actual jump or not. For the changes needed in the proof see Remark 4.2. This version will be used when proving Theorem 1.3.

3 Returning to random walk on dynamical percolation

In this section we prove Proposition 1.5 using the general result Theorem 2.1 stated in the previous section. We also prove Theorem 1.3. In order to apply Theorem 2.1 we need to obtain a bound on φ valid for all values of p. We also need to discretise our process.

Before starting the proofs we introduce some notation and prove some preliminary results.

Let $\mathcal{E} := D_{[0,1]}(\{0,1\}^{E(\mathbb{Z}_n^d)})$ be the space of right continuous paths with left limits from $[0,1]$ into $\{0,1\}^{E(\mathbb{Z}_n^d)}$. Letting $\bar{\eta}_k := \eta_{[k-1,k]}$, it is easy to see that $((X_k, \bar{\eta}_k))_{k\geq 0}$ is a Markov chain in a Markovian evolving environment. It is clear that $(\bar{\eta}_k)_{k\geq 0}$ is a Markov chain with state space $\mathcal{E}$ and that for all $\zeta \in \mathcal{E}$, the corresponding Markov chain p_ζ simply corresponds to doing random walk on $\mathbb{Z}_n^d$ for time 1 during which time the bond configuration evolution is fixed to be ζ.

Lemma 3.1. *For all $\delta > 0$, there exists $\sigma = \sigma(\delta) > 0$ so that for all $d \geq 1$, for all n, for all $\mu \leq 1/2$, for all $p \in [\delta, 1]$, for all $A \subseteq E(\mathbb{Z}_n^d)$ and for all η_0,*

$$\mathcal{P}_{\eta_0}(|a \in A : \eta_t(a) = 1 \text{ for all } t \in [1/2, 1]| \geq |A|\sigma\mu) \geq \sigma\mu.$$

Proof. The left hand size is minimised when $\eta_0 \equiv 0$. In this case, the left hand side equals $\mathbb{P}(\text{Bin}(|A|, (1 - e^{-\mu/2})pe^{-\mu(1-p)/2}) \geq |A|\sigma\mu)$ where $\text{Bin}(m, q)$ denotes a Binomial random variable with parameters m and q. Since $\mu \leq 1/2$ and $p \geq \delta$, it is easy to see that there exists a σ satisfying the requirements. $\square$

For $S \subseteq \mathbb{Z}_n^d$ we now let $\partial_E(S)$ denote the *edge boundary* of S which is the set of edges from S to S^c. We simply write $\varphi_{\eta_{[0,1]}}$ to denote $\varphi_{p_{\eta_{[0,1]}}, \pi}$, i.e. we run the environment process for time 1 starting from η_0 and use $p_{\eta_{[0,1]}}$ for the transition probability of the random walk.

Lemma 3.2. *For all d, there exists $c_d > 0$ so that for all n, for all $\mu \leq 1/2$, for all p, for all β, for all nonempty sets $S \subseteq \mathbb{Z}_n^d$ with $\pi(S) \leq 1/2$, for all $\eta_{[0,1]}$ satisfying*

$$|e \in \partial_E(S) : \eta_t(e) = 1 \text{ for all } t \in [1/2, 1]| \geq |\partial_E(S)|\beta,$$

we have

$$\varphi_{\eta_{[0,1]}}(S) \geq \frac{c_d\beta}{n(\pi(S))^{1/d}}.$$

Proof. Let $S_{\text{good}} := \{s \in S : \exists e \text{ from } s \text{ to } S^c \text{ which is open during } [1/2, 1]\}$ and let $S_{\text{bad}} := S \backslash S_{\text{good}}$. (Note that S_{good} is a subset of the *internal vertex boundary* of S.) Note that

$$|S_{\text{good}}| \geq \frac{1}{2d} \times |e \in \partial_E(S) : \eta_t(e) = 1 \text{ for all } t \in [1/2, 1]|$$

and hence $|S_{\text{good}}| \geq \frac{|\partial_E(S)|\beta}{2d}$. Since $\pi(S) \leq 1/2$, by the standard isoperimetric inequality on $\mathbb{Z}_n^d$, we have that $|\partial_E(S)| \geq c'_d |S|^{(d-1)/d}$ for some universal constant c'_d only depending on d. It follows that

$$|S_{\text{good}}| \geq \frac{c'_d}{2d} |S|^{(d-1)/d}\beta. \tag{3.1}$$

Consider now

$$\varphi_{\eta_{[0,1]}}(S) = \mathbb{P}_{\pi, \eta_{[0,1]}}(X_1 \notin S \mid X_0 \in S).$$

The subscript $\eta_{[0,1]}$ means that the environment is fixed to be this realisation. The conditioning $X_0 \in S$ gives probability $1/|S|$ to each point in S. Since the uniform distribution is stationary for all realisations of the environment by the definition of the random walk, one can infer that

$$\max_{y \in \mathbb{Z}_n^d} \mathbb{P}_{\pi, \eta_{[0,1]}}\left(X_{\frac{1}{2}} = y \mid X_0 \in S\right) \leq \frac{1}{|S|}. \tag{3.2}$$

Now

$$\begin{aligned}
\mathbb{P}_{\pi, \eta_{[0,1]}}(X_1 \notin S \mid X_0 \in S) \geq &\mathbb{P}_{\pi, \eta_{[0,1]}}\left(X_{\frac{1}{2}} \in S_{\text{good}} \cup S^c \mid X_0 \in S\right) \\
&\times \mathbb{P}_{\pi, \eta_{[0,1]}}\left(X_1 \notin S \mid X_0 \in S, X_{\frac{1}{2}} \in S_{\text{good}} \cup S^c\right).
\end{aligned} \tag{3.3}$$

7

By (3.2), the first factor on the right hand side is at least $1 - \frac{|S_{\mathrm{bad}}|}{|S|}$. This equals $\frac{|S_{\mathrm{good}}|}{|S|}$, which, by (3.1), is at least $\frac{c'_d}{2d}|S|^{-1/d}\beta = \frac{c'_d}{2d}(\pi(S))^{-1/d}n^{-1}\beta$. For the second term, if $X_{\frac{1}{2}} \in S_{\mathrm{good}}$, we fix an arbitrary edge e from $X_{\frac{1}{2}}$ to S^c which is open during $[1/2, 1]$. The probability that the random walk attempts only one jump during $[\frac{1}{2}, 1]$ and the attempted jump is along this edge is at least a constant $\gamma = \gamma(d) > 0$ only depending upon d. On the other hand, if $X_{\frac{1}{2}} \in S^c$, there is a fixed probability the walk does not move during $[1/2, 1]$, which we can also take to be $\gamma(d)$.

This gives that the left hand side of (3.3) is at least $\frac{\gamma(d)c'_d}{2d}(\pi(S))^{-1/d}n^{-1}\beta$. Letting $c(d) := \frac{\gamma(d)c'_d}{2d}$ yields the claim. $\qquad\square$

Proof of Proposition 1.5. We will apply Theorem 2.1 with $\mathcal{E}$ being the space of right continuous paths with left limits on $[0, 1]$. Observe that for all x

$$t_{\mathrm{mix}}(\varepsilon, x, \eta) \leq t_{\mathrm{mix}}(\varepsilon, x, (\overline{\eta}_k)_{k \geq 0})).$$

We now show that for all $d \geq 1$ and $\delta > 0$, there exists $C_1 = C_1(d, \delta) > 0$ so that for all $p \in [\delta, 1]$, for all $\mu \leq 1/2$, for all n, for all $\eta_0 \in \{0, 1\}^{E(\mathbb{Z}_n^d)}$ and for all S with $\pi(S) \leq 1/2$,

$$\mathcal{P}_{\eta_0}\left(\varphi_{\overline{\eta}_1}(S) \geq \frac{C_1\mu}{n(\pi(S))^{1/d}}\right) \geq C_1\mu. \tag{3.4}$$

Fix d and δ. Choose $\sigma(\delta)$ from Lemma 3.1 and c_d from Lemma 3.2. Fix S with $\pi(S) \leq 1/2$. Combining Lemmas 3.1 and 3.2 with A in Lemma 3.1 taken to be $\partial_E(S)$ and β in 3.2 taken to be $\sigma\mu$, the two lemmas imply that

$$\mathcal{P}_{\eta_0}\left(\varphi_{\overline{\eta}_1}(S) \geq \frac{c_d\sigma\mu}{n(\pi(S))^{1/d}}\right) \geq \sigma\mu,$$

establishing (3.4). From (3.4) we now get that for all sets S with $\pi(S) \leq 1/2$ and for all η_0

$$\varphi(\eta_0, S) \geq \frac{C_1^2\mu^2}{n(\pi(S))^{1/d}},$$

and hence

$$\int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u\varphi^2(u)} = \int_{4/n^d}^{1/2} \frac{du}{u\varphi^2(u)} + \frac{1}{\varphi^2(1/2)} \int_{1/2}^{4/\varepsilon} \frac{du}{u} \leq C_2 \left(\frac{n}{\mu^2}\right)^2 \log\left(\frac{1}{\varepsilon}\right),$$

where C_2 is a positive constant. Since for any environment η and any $x \in \mathbb{Z}_n^d$ we have

$$\mathbb{P}_\eta(X_1 = x \mid X_0 = x) \geq \frac{1}{e},$$

applying Theorem 2.1 completes the proof. $\qquad\square$

We turn to prove Theorem 1.3. We now let $\widetilde{\mathcal{E}} := D_{[0,1/\mu]}(\{0, 1\}^{E(\mathbb{Z}_n^d)})$ be the space of right continuous paths with left limits from $[0, 1/\mu]$ into $\{0, 1\}^{E(\mathbb{Z}_n^d)}$. Let $\widetilde{\eta}_k := \eta_{[\frac{k-1}{\mu}, \frac{k}{\mu}]}$ and $Z_k := X_{\frac{k}{\mu}}$. Then again $((Z_k, \widetilde{\eta}_k))_{k \geq 0}$ is a Markov chain in a Markovian evolving environment. It is clear that $(\widetilde{\eta}_k)_{k \geq 0}$ is a Markov chain with state space $\widetilde{\mathcal{E}}$ and that for all $\zeta \in \widetilde{\mathcal{E}}$, the corresponding Markov chain p_ζ simply corresponds to doing random walk on $\mathbb{Z}_n^d$ for time $1/\mu$ during which time the bond configuration evolution is fixed to be ζ.

The following lemmas follow in exactly the same way as Lemmas 3.1 and 3.2.

Lemma 3.3. *For all $\delta > 0$, there exists $\sigma = \sigma(\delta) > 0$ so that for all $d \geq 1$, for all n, for all $\mu \leq 1/2$, for all $p \in [\delta, 1]$, for all $A \subseteq E(\mathbb{Z}_n^d)$ and for all η_0,*

$$\mathcal{P}_{\eta_0}\left(\left|a \in A : \eta_t(a) = 1 \text{ for all } t \in \left[\frac{1}{\mu} - 1, \frac{1}{\mu}\right]\right| \geq |A|\sigma\right) \geq \sigma.$$

Lemma 3.4. *For all d, there exists $c_d > 0$ so that for all n, for all $\mu \leq 1/2$, for all p, for all β, for all nonempty sets $S \subseteq \mathbb{Z}_n^d$ with $\pi(S) \leq 1/2$, for all $\eta_{[0,1/\mu]}$ satisfying*

$$\left|e \in \partial_E(S) : \eta_t(e) = 1 \text{ for all } t \in \left[\frac{1}{\mu} - 1, \frac{1}{\mu}\right]\right| \geq |\partial_E(S)|\beta,$$

we have

$$\varphi_{\eta_{[0,1/\mu]}}(S) \geq \frac{c_d\beta}{n(\pi(S))^{1/d}}.$$

Proof of Theorem 1.3. In order to prove the theorem we first consider a lazy version of the Markov chain $((Z_k, \widetilde{\eta}_k))_{k\geq 0}$ as follows. At every step, the walk remains in place with probability $1/2$ and with probability $1/2$ it jumps according to the transition matrix given by the environment at this time. When Z stays in place, then the environment at the next step also stays in place, otherwise it moves according to its transition kernel. This is the setup of Remark 2.2.

For this new chain the statement of Lemma 3.4 remains the same with an extra factor of $1/2$ in the lower bound for $\varphi_{\eta_{[0,1/\mu]}}$, which is defined to be $\varphi_{(P_{\eta_{[0,1/\mu]}}+I)/2}$ (see Remark 4.2). Also Lemma 3.3 still holds with an extra factor of $1/2$ again in the lower bound of the probability.

Remark 2.2 now shows that the statement of Theorem 2.1 remains true, and hence we obtain

$$\max_{\eta_0}\mathcal{P}_{\eta_0}\left(\eta = (\eta_t)_{t\geq 0} : t_{\text{mix}}(\varepsilon, x, \eta) \geq \frac{Cn^2\log(1/\varepsilon)}{\mu}\right) \leq \varepsilon, \tag{3.5}$$

where the mixing time refers to the lazy version of the discretised random walk.

Using (3.5) for $\varepsilon = 1/4$ and performing independent experiments immediately gives that for some constant C

$$\max_{x,\eta_0}\mathbb{E}_{x,\eta_0}[\tau_A] \leq \frac{Cn^2}{\mu}.$$

By letting $\varepsilon = \varepsilon/n^d$ in (3.5) and taking a union bound over all $x \in \mathbb{Z}_n^d$ we obtain

$$\max_{\eta_0}\mathcal{P}_{\eta_0}\left(\eta = (\eta_t)_{t\geq 0} : t_{\text{mix}}(\varepsilon, \eta) \geq \frac{Cn^2\log(n/\varepsilon)}{\mu}\right) \leq \varepsilon. \tag{3.6}$$

We will apply (3.6) when $\varepsilon = 1/100$. We break time into intervals of length $T := C'n^2\log n/\mu$ and call an interval good if the "mixing time during that interval" is at most T. Equation (3.6) yields that for all η_0, the set of good intervals dominates an i.i.d. process with density 0.99.

By the usual comparison of hitting times with mixing times, it follows that, conditioned on the past, if we enter a good interval, then, no matter where we start, we have a probability of at least $1/4$ of hitting A. It follows that if the proportion of good intervals within the first ℓ intervals is at least 0.9, then, for all starting states, the probability that we have not hit A by time ℓT is at most $(3/4)^{0.9\ell}$. Since an interval is good with probability 0.99, using the domination above, the probability that there are at most proportion 0.9 good intervals among the first ℓ decays like c^ℓ for some $c < 1$. Hence, the (environment) probability that for some starting point, the (random walk) probability that we do not hit A by time ℓT is larger than $(3/4)^{0.9\ell}$ decays like c^ℓ. It follows

9

that the (environment) probability that for some $\ell \geq k$ and some starting point, the (random walk) probability that we have not hit A by time ℓT is larger than $(3/4)^{0.9\ell}$ decays like c^k (since the tail of a geometric series also decays exponentially). On the complementary event, which has probability $1 - c^k$, we have that for all $\ell \geq k$, for all starting points, the (random walk) probability that we have not hit A by time ℓT is at most $(3/4)^{0.9\ell}$. Expressing the expectation of a random variable as a sum of the tail probabilities, we obtain that on this event, the expectation of the hitting time from any starting point is at most kT.

Therefore we obtain

$$\max_{\eta_0} \mathcal{P}_{\eta_0} \left(\eta = (\eta_t)_{t \geq 0} : \max_x \mathbb{E}_{x,\eta}[\tau_A] \geq k \cdot \frac{C' n^2 \log n}{\mu} \right) \leq c^k$$

and this concludes the proof. $\qquad\square$

Remark 3.5. We note that in the above proof we obtained an upper bound of n^2/μ for the mixing time of the lazy version of the chain as defined above. However, this does not give us a bound on the mixing time of the original chain. On the other hand, we used it in the proof to obtain an exit time result for the lazy version which does pass through to the original chain.

4 Evolving Sets

In this section, we extend the theory of evolving sets to the setting of a Markovian random environment in order to prove Theorem 2.1. The exposition here follows closely [3].

We first recall the definition of evolving sets in the context of a finite state Markov chain; see [2]. Given a Markov chain $p(x, y)$ with state space Ω and a stationary distribution π, the corresponding evolving-set process $\{S_n\}_{n \geq 0}$ is a Markov chain on subsets of Ω whose transitions are described as follows. Let Q be defined as in (2.1) (with ν being π) and let U be a uniform random variable on $[0, 1]$. If $S \subseteq \Omega$ is the present state, we let the next state $\widetilde{S}$ be defined by

$$\widetilde{S} := \left\{ y \in \Omega : \frac{Q(S, y)}{\pi(y)} \geq U \right\}.$$

Note that Ω and $\varnothing$ are absorbing states and it is immediate to check that

$$\mathbb{P}(y \in S_{k+1} \mid S_k) = \frac{Q(S_k, y)}{\pi(y)}. \tag{4.1}$$

Moreover, one can describe the evolving set process as that process on subsets which satisfies the "one-dimensional marginal" condition (4.1) and where these different events, as we vary y, are maximally coupled.

For later use we also define now

$$\psi_p(S) := 1 - \mathbb{E}\left[\sqrt{\frac{\pi(\widetilde{S})}{\pi(S)}} \right],$$

where $\widetilde{S}$ is the first step of the evolving set process started from S and when the transition probability for the Markov chain is p and the stationary distribution π.

We next define completely analogously the evolving set process in the context of a time inhomogeneous Markov chain. Consider a time inhomogeneous Markov chain with state space $\mathcal{S}$ whose

transition matrix for moving from time k to time $k+1$ is given by $p_{k+1}(x,y)$ where we assume that the probability measure π is a stationary distribution for each p_k. In this case, we say that π is a stationary distribution for the inhomogeneous Markov chain. Let Q_k be defined as in (2.1) but with respect to p_k and π. We then obtain a time inhomogeneous Markov chain $S_0, S_1, \ldots$ on subsets of $\mathcal{S}$ generated by

$$S_{k+1} := \left\{ y \in \mathcal{S} : \frac{Q_{k+1}(S_k, y)}{\pi(y)} \geq U_{k+1} \right\}$$

where $(U_i)_i$ are i.i.d. random variables uniform on $[0,1]$. We call this the evolving set process with respect to $p_1, p_2, \ldots$ and π.

We now need to consider the Doob transform of the evolving set process. If P_ζ is the transition probability for the evolving set process when the environment is ζ, then we define the Doob transform via

$$\widehat{P}_\zeta(S, S') = \frac{\pi(S')}{\pi(S)} P_\zeta(S, S').$$

We now let $\psi(\zeta, S) := \mathbb{E}_\zeta \left[\psi_{p_{\eta_1}}(S) \right]$. For $r \in [\pi_\star, \frac{1}{2}]$, we let

$$\psi(r) := \inf\{\psi(\zeta, S) : \zeta \in \mathcal{E}, \pi(S) \leq r\}$$

and $\psi(r) := \psi(\frac{1}{2})$ for $r \geq \frac{1}{2}$.

In the following, we let

$$S^\# := \begin{cases} S & \text{if } \pi(S) \leq \frac{1}{2} \\ S^c & \text{otherwise} \end{cases}$$

and

$$Z_n := \frac{\sqrt{\pi(S_n^\#)}}{\pi(S_n)}.$$

Lemma 4.1. *Let $\varepsilon > 0$ and $x \in \mathcal{S}$. If*

$$n \geq \int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u\psi(u)},$$

then $\widehat{\mathbb{E}}_{\{x\},\eta_0}[Z_n] \leq \sqrt{\varepsilon}$ for all η_0.

Proof. We fix x, η_0 and to simplify notation we do not include them in the notation. We now get that almost surely

$$\widehat{\mathbb{E}}\left[\frac{Z_{n+1}}{Z_n} \,\middle|\, S_n, \eta_0, \ldots, \eta_n \right] = \mathbb{E}\left[\frac{\pi(S_{n+1})}{\pi(S_n)} \cdot \frac{Z_{n+1}}{Z_n} \,\middle|\, S_n, \eta_0, \ldots, \eta_n \right] = \mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^\#)}{\pi(S_n^\#)}} \,\middle|\, S_n, \eta_0, \ldots, \eta_n \right].$$

Suppose first that $\pi(S_n) \leq 1/2$. Then

$$\mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^\#)}{\pi(S_n^\#)}} \,\middle|\, S_n, \eta_0, \ldots, \eta_n \right] \leq \mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1})}{\pi(S_n)}} \,\middle|\, S_n, \eta_0, \ldots, \eta_n \right].$$

The Markov property of the environment now gives

$$\mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1})}{\pi(S_n)}} \,\bigg|\, S_n, \eta_0, \ldots, \eta_n\right] = \sum_\eta R(\eta_n, \eta)(1 - \psi_{p_\eta}(S_n)) \tag{4.2}$$
$$= 1 - \psi(\eta_n, S_n) \leq 1 - \psi(\pi(S_n)).$$

Suppose next that $\pi(S_n) > 1/2$. Then

$$\mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^\#)}{\pi(S_n^\#)}} \,\bigg|\, S_n, \eta_0, \ldots, \eta_n\right] \leq \mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^c)}{\pi(S_n^c)}} \,\bigg|\, S_n, \eta_0, \ldots, \eta_n\right]$$

and using the Markov property of the environment as before (as well as the fact that $(S_n^c)_n$ is also an evolving set process) we obtain

$$\mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^c)}{\pi(S_n^c)}} \,\bigg|\, S_n, \eta_0, \ldots, \eta_n\right] \leq 1 - \psi(\pi(S_n^c)). \tag{4.3}$$

Since the function ψ is non-increasing, it follows that if $\pi(S_n) > 1/2$, then $\psi(\pi(S_n^c)) \geq \psi(\pi(S_n))$, and hence from (4.2) and (4.3) we get that in all cases

$$\mathbb{E}\left[\sqrt{\frac{\pi(S_{n+1}^\#)}{\pi(S_n^\#)}} \,\bigg|\, S_n, \eta_0, \ldots, \eta_n\right] \leq 1 - \psi(\pi(S_n)) = 1 - f_0(Z_n),$$

where following [3] we set $f_0(z) := \psi(1/z^2)$ which is non-decreasing. (Note that $\pi(S_n) = Z_n^{-2}$ when $Z_n \geq \sqrt{2}$, i.e. when $\pi(S_n) \geq 1/2$ and $\psi(x) = \psi(1/2)$ for $x \geq 1/2$.) Therefore, we conclude

$$\widehat{\mathbb{E}}[Z_{n+1} \mid Z_n] \leq Z_n(1 - f_0(Z_n))$$

and hence using [3, Lemma 11 (iii)] we get that $\widehat{\mathbb{E}}[Z_n] \leq \sqrt{\varepsilon}$ for all

$$n \geq \int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u\psi(u)}$$

and this finishes the proof. $\qquad\square$

Remark 4.2. We now explain the changes in the proof of the lemma above needed to justify Remark 2.2. The matrix R is replaced by $R(\zeta, \ell, \eta)$, where $\ell \in \{0, 1\}$ depending on whether the walk made an actual step or not, i.e. $R(\zeta, 0, \eta) = \mathbf{1}(\zeta = \eta)$ and $R(\zeta, 1, \eta) = R(\zeta, \eta)$. In the definition of ψ, the matrix p_η is replaced by $(p_\eta + I)/2$. The proof of the lemma above then remains unchanged with this new notation.

Lemma 4.3. *If $(S_k)_{k\geq 0}$ is the evolving set process relative to an inhomogeneous Markov chain (X_k) and stationary distribution π, then*

$$\mathbb{P}_x(X_k = y) = \frac{\pi(y)}{\pi(x)} \mathbb{P}_{\{x\}}(y \in S_k).$$

Proof. In the case of a homogeneous Markov chain, this is Lemma 17.12 in [2]. The proof for the inhomogeneous case goes through verbatim. $\qquad\square$

Lemma 4.4. *If $\{S_k\}_{k \geq 0}$ is the evolving set process relative to an inhomogeneous Markov chain with stationary distribution π, then starting from any initial state, $\{\pi(S_k)\}_{k \geq 0}$ is a martingale.*

Proof. In the case of a homogeneous Markov chain, this is Lemma 17.13 in [2]. The proof for the inhomogeneous case goes through verbatim. □

Lemma 4.5. *For all fixed environments $\eta = (\eta_i)_{i \geq 0}$ and all $x \in \mathcal{S}$ we have*

$$\chi \left(\mathbb{P}_{x,\eta}(Y_n = \cdot) , \pi \right) \leq \widehat{\mathbb{E}}_{x,\eta}[Z_n].$$

Proof. In the homogeneous case this is [3, equation (24)]. The proof for the inhomogeneous case goes through verbatum using Lemmas 4.3 and 4.4. □

Proof of Theorem 2.1. By Markov's inequality we obtain

$$\mathcal{P}_\zeta \left(\eta : \chi(\mathbb{P}_{x,\eta}(Y_n = \cdot) , \pi) \geq \varepsilon^{1/4} \right) \leq \varepsilon^{-1/4} \mathbb{E}_\zeta \left[\chi(\mathbb{P}_{x,\eta}(Y_n = \cdot) , \pi) \right],$$

where the last expectation is taken over the environment η started from ζ. From Lemma 4.5 we can upper bound the right hand side by $\varepsilon^{-1/4} \widehat{\mathbb{E}}_{\{x\},\zeta}[Z_n]$. From Lemma 4.1 we get that $\widehat{\mathbb{E}}_{\{x\},\zeta}[Z_n] \leq \sqrt{\varepsilon}$ for all

$$n \geq 1 + \int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u\psi(u)}.$$

Lemma 10 in [3] implies that for all p and S,

$$\psi_p(S) \geq \frac{\gamma^2}{2(1-\gamma)^2} \varphi_p^2(S).$$

Therefore, taking expectations and using Jensen's inequality we get for all ζ and S

$$\psi(\zeta, S) \geq \frac{\gamma^2}{2(1-\gamma)^2} \varphi^2(\zeta, S),$$

and hence for all r

$$\psi(r) \geq \frac{\gamma^2}{2(1-\gamma)^2} \varphi^2(r).$$

Thus this gives that for all

$$n \geq 1 + \frac{2(1-\gamma)^2}{\gamma^2} \int_{4\pi(x)}^{4/\varepsilon} \frac{du}{u\varphi^2(u)}$$

we have $\widehat{\mathbb{E}}_{\{x\},\zeta}[Z_n] \leq \sqrt{\varepsilon}$ and hence this completes the proof. □

Acknowledgements

References

[1] Luca Avena, Hakan Güldaş, Remco van der Hofstad, and Frank den Hollander. Mixing times of random walks on dynamic configuration models. *Ann. Appl. Probab.*, 28(4):1977–2002, 2018.

[2] 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.

[3] Ben Morris and Yuval Peres. Evolving sets, mixing and heat kernel bounds. *Probab. Theory Related Fields*, 133(2):245–266, 2005.

[4] Yuval Peres, Perla Sousi, and Jeffrey E. Steif. Mixing time for random walk on supercritical dynamical percolation. 2017. arXiv:1707.07632.

[5] Yuval Peres, Alexandre Stauffer, and Jeffrey E. Steif. Random walks on dynamical percolation: mixing times, mean squared displacement and hitting times. *Probab. Theory Related Fields*, 162(3-4):487–530, 2015.